International Journal of Innovative Research and Scientific Studies, 8(2) 2025, pages: 409-419

IJIRSS
International Journal of Innovative Research & Scientific Studies

ISSN: 2617-6548

URL: www.ijirss.com

Sustainable IoT-enabled predictive analytics for maternal health risk prediction: A deep learning approach

Abayomi Agbeyangi[1*], Jose Lukose[2]

[1,2]*Walter Sisulu University, Buffalo City Campus, East London, South Africa.*

Corresponding author: Abayomi Agbeyangi (*Email: aagbeyangi@wsu.ac.za*)

Abstract

Maternal health is a significant concern, especially in low-resource environments with limited healthcare infrastructure, economic constraints, and access. The rise of the Internet of Things (IoT) and deep learning presents promising solutions. This study explores the deep learning approach to create an IoT-driven predictive analytics model to evaluate maternal health risks. By using the Maternal Health Risk Dataset, the ratio of systolic to diastolic blood pressure was engineered (BP_ratio). The evaluation included random forest, support vector machine, and gradient boosting alongside the deep learning model. The deep learning model achieved a balanced performance with an accuracy of 71.17%, a precision of 72.78%, a recall of 70.29%, and an F1-score of 65.71%. These results suggest that integrating IoT with predictive analytics can enhance early detection and intervention, reducing maternal mortality and morbidity. The study offers practical insights for healthcare stakeholders and policymakers in low-resource environments to implement efficient and scalable healthcare solutions.

Keywords: Deep learning, IoT, Low-resource environments, Maternal health, Predictive analytics, Sustainability.

DOI: 10.53894/ijirss.v8i2.5189
Funding: This study received no specific financial support.
History: Received: 10 January 2025 / **Revised:** 10 February 2025 / **Accepted:** 17 February 2025 / **Published:** 7 March 2025

Competing Interests: The authors declare that they have no competing interests.
Authors' Contributions: All authors contributed equally to the conception and design of the study. All authors have read and agreed to the published version of the manuscript.
Transparency: The authors confirm that the manuscript is an honest, accurate, and transparent account of the study; that no vital features of the study have been omitted; and that any discrepancies from the study as planned have been explained. This study followed all ethical practices during writing.
Publisher: Innovative Research Publishing

1. Introduction

Maternal health remains a paramount issue globally, with approximately 295,000 women dying annually from preventable complications related to pregnancy and childbirth, predominantly in low-resource environments [1]. Ensuring timely and effective healthcare interventions in these regions is crucial to reducing maternal mortality and morbidity rates. Traditional healthcare systems in these areas often face inadequate infrastructure, limited access to medical facilities, and insufficient healthcare personnel [2, 3].

The Fourth Industrial Revolution [4] has ushered in a wave of innovative technologies, such as the Internet of Things (IoT) and deep learning, which hold the potential to revolutionize healthcare delivery [5]. IoT devices, with their ability to continuously monitor vital signs and other health indicators, provide real-time data that can be analyzed to predict and mitigate health risks. Deep learning algorithms, capable of processing large volumes of data to identify patterns and predict adverse health outcomes, enable early intervention, offering a beacon of hope in the fight against maternal mortality and morbidity [6, 7].

While existing studies, such as those by Ahmed and Kashem [8], on an IoT-based Risk Level Prediction Model for Maternal Health Care in Bangladesh, have shown the potential of IoT and machine learning in maternal health, there is a pressing need for sustainable and scalable solutions. These solutions must be effectively deployable in low-resource settings without exacerbating environmental concerns, underscoring the urgency and importance of the research. Specifically, the focus is on developing a sustainable and scalable IoT-driven predictive analytics model tailored for low-resource settings, emphasizing energy efficiency and real-time monitoring for improved maternal health outcomes.

This initial study aims to address this gap by developing a sustainable and scalable IoT-driven predictive analytics model for maternal health risk assessment using deep learning techniques. It leverages the Maternal Health Risk Data from Kaggle [8], incorporating advanced feature engineering and deep learning models to predict risk levels. The research focuses on ensuring that the proposed solutions are effective, environmentally sustainable, and adaptable to various healthcare contexts. However, it is essential to note that implementing such a model may face challenges such as data limitation concerns, technical infrastructure requirements, and the need for skilled personnel.

The proposed research makes three significant contributions, particularly in the context of IoT and deep learning, while addressing the unique challenges faced by low-resource healthcare systems:

- The study presents a deep learning model specifically tailored for predicting maternal health risks based on IoT-driven data. This model effectively categorizes risk levels (low, medium, high) by analyzing multiple health indicators, offering a more accurate and reliable tool for maternal health assessment.
- The research introduces a novel feature, BP Ratio, which captures the ratio of Systolic BP to Diastolic BP. This feature enhances the model's predictive accuracy by providing a deeper understanding of cardiovascular health, making the model more effective in assessing maternal health risks.
- The study emphasizes the model's application in low-resource settings with limited healthcare monitoring.

The research is guided by the following question: How can a sustainable and scalable IoT-driven predictive analytics model utilizing deep learning improve maternal health outcomes in low-resource settings? The remaining sections of the paper discuss related work, detailed methodology, experimental results, and a discussion of the findings. It concludes with implications for practice and recommendations for future research, highlighting the significant potential of sustainable ICT solutions in improving maternal health outcomes in low-resource settings.

2. Related Work

IoT and machine learning applications, particularly maternal health, have seen significant advancements in healthcare. Mondal, et al. [9] integrated IoT for real-time maternal health monitoring during pregnancy, focusing on vital signs like heart rate, blood pressure, fetal movements, and temperature. Their use of machine learning, particularly the Random Forest Classifier, achieved a high accuracy in predicting maternal health risks. Ahmed and Kashem [8] explored an IoT-based risk level prediction model for maternal health in Bangladesh, highlighting the potential of IoT devices to monitor and predict maternal health risks. However, their work primarily focused on conventional machine learning models, which have limitations in handling complex and high-dimensional data.

Marques, et al. [10] developed a system that uses IoT sensors to monitor maternal and fetal signals, employing a 1-D CNN classifier for diagnostics. The system tracks vital indicators such as fetal heart rate, uterine activity, and maternal metrics like blood pressure and oxygen saturation, achieving a diagnostic accuracy of 92.59% for emergencies and an F1-score between 0.74 and 0.91 for health status classification. In contrast, this study focuses on creating a sustainable and scalable deep learning model designed explicitly for low-resource settings.

Ettiyan and Geetha [5] introduced an intelligent monitoring system for maternal and fetal signals in high-risk pregnancies, leveraging an optimized single-dimensional Convolutional Neural Network (1D-OCNN) for enhanced classification and prediction of emergencies. Their evaluation demonstrated that the proposed 1D-OCNN outperformed other learning algorithms in terms of accuracy, precision, recall, and F1 score. Similarly, Krupa, et al. [6] developed an IoT-based deep-learning approach to detect fetal QRS complexes in abdominal ECG (AECG) signals without eliminating the maternal components. They employed two time-frequency techniques, Hilbert Huang Transform (HHT) and Stockwell Transform (ST), to convert AECG signals into two-dimensional images. They were then used to train pre-trained models MobileNet and ResNet18. The study utilized the 2013 challenge database for evaluation and conducted a comparative analysis between the time-frequency approaches and the deep learning architectures. Unlike Krupa et al.'s work, which centers on signal processing for fetal monitoring, this study approach aims to improve maternal healthcare outcomes by integrating IoT data with a deep learning model tailored for risk prediction, ensuring adaptability and sustainability in challenging healthcare environments.

Hossain, et al. [11] introduced a Medical Cyber-Physical System (MCPS) designed to predict maternal health risks using machine learning, specifically targeting pregnant women in developing countries. The system gathers health metrics such as temperature, blood pressure, glucose levels, and heart rate through IoT sensors. Machine learning algorithms analyze these metrics to predict potential health issues, and alerts are sent to nearby clinics when necessary. The application, built using the Streamlit framework, is accessible to users and demonstrates high accuracy with the XGB classifier, achieving 99% accuracy and low error rates. Furthermore, Venkatasubramanian [12] developed an IoT-based system using DCGAN to monitor high-risk maternal patients, reaching over 80% AUC in predicting health outcomes through advanced feature extraction, outperforming existing deep learning models in emergency diagnostics. Raza, et al. [13] proposed the DT-BiLTCN architecture to predict maternal health risks by analyzing health records from 1,218 samples, addressing class imbalance with synthetic oversampling. The study identified critical indicators such as blood pressure, heart rate, and maternal age, achieving 98% accuracy with an SVM classifier.

The related studies have examined the use of machine learning for predicting maternal health risks, but most rely on traditional algorithms. These models often lack the predictive power required for real-time applications in low-resource settings. Recent research has started to explore deep learning approaches, which offer better accuracy and scalability. However, there is still a need to address sustainability, energy efficiency, and the environmental impact of these solutions.

3. Methods

The methodology encompasses the dataset description, feature engineering, data preprocessing, model development and architecture, and evaluation metrics. The overview of the methodology is described in Figure 1.

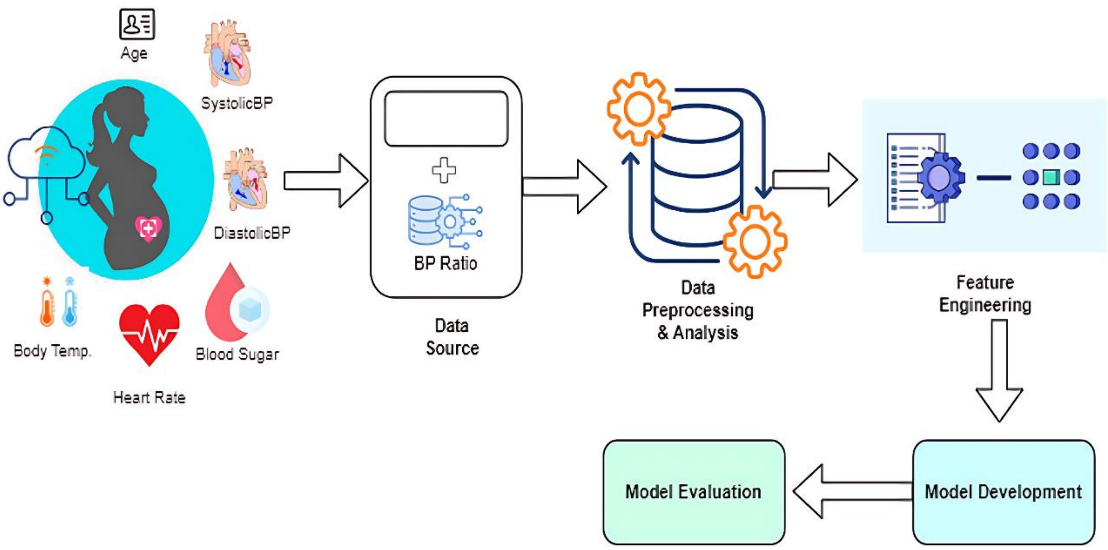

Figure 1.
Overview of the IoT-Driven Predictive Analytics Pipeline for Maternal Health (Illustrates the data flow from health monitoring to model evaluation, incorporating feature engineering and deep learning for risk prediction).

3.1. Model Formalization

This formalization provides a clear mathematical framework for the model's structure and training process, as emphasized in Sections 3.2 to 3.6.

Let $X \in R^{n\times d}$ represent the input dataset, where n is the number of instances and d is the number of features. Here, X consists of 7 features: *Age, Systolic BP, Diastolic BP, Blood Sugar (BS), Body Temperature, Heart Rate, and Blood Pressure Ratio (BP_Ratio).*

Let $y \in \{0,1,2\}^n$ be the target vector corresponding to the risk levels, where 0 represents low risk, 1 represents medium risk, and 2 represents high risk.

The input to the neural network is denoted as:

$$X_i = [x_{i1}, x_{i2},..., x_{i7}]^T \tag{1}$$

for $i = 1, 2, ..., n$

For the hidden layers, the output of the l-th hidden layer $h^{(l)}$ is computed as:

$$h(l) = f(W(l)h(l-1) + b(l)) \tag{2}$$

where:

- $W^{(l)} \in R^{ml\times ml-1}$ is the weight matrix for layer l.
- $b^{(l)} \in R^{ml}$ is the bias vector for layer l.
- $f(\cdot)$ is the ReLU activation function defined as $f(x) = \max(0, x)$.
- m_l is the number of neurons in the l-th layer.

In the model:

- First hidden layer: $m_1 = 64$
- Second hidden layer: $m_2 = 32$
- Third hidden layer: $m_3 = 16$

For the dropout layer, the dropout is applied after each hidden layer. The output of the dropout layer $d^{(l)}$ is given by:

1. $$d^{(l)} = Dropout(h^{(l)}, p) \tag{3}$$

where $p = 0.5$ is the dropout rate.

The output layer computes the final prediction vector $\hat{y}$ using the softmax activation function:

$$\hat{y}i = softmax(W(L)h(L-1) + b(L)) \quad (4)$$

where:

$$softmax(\boldsymbol{z})_j = \frac{e^{z_j}}{\sum_{k=1}^{K} e^{z_k}} \quad (5)$$

Here, $K = 3$ is the number of output classes, corresponding to the risk levels.

The model is trained by minimizing the categorical cross-entropy loss L, defined as:

$$\boldsymbol{\mathcal{L}}(\boldsymbol{y}, \boldsymbol{y}\hat{}) = -\frac{1}{n}\sum_{i=1}^{n}\sum_{k=0}^{2} y_{ik} log(\boldsymbol{y}\hat{}_{ik}) \quad (6)$$

Finally, the Adam optimizer is used to update the model parameters $W^{(l)}$ and $b^{(l)}$ during training, minimizing the loss function L. The update rules for the parameters at each iteration t are:

$$\mathbf{W}_{t+1}^{(l)} = \mathbf{W}_t^{(l)} - \eta \cdot \frac{\hat{m}_t^{(l)}}{\sqrt{\hat{v}_t^{(l)}} + \epsilon}$$
$$\mathbf{b}_{t+1}^{(l)} = \mathbf{b}_t^{(l)} - \eta \cdot \frac{\hat{m}_t^{(b)}}{\sqrt{\hat{v}_t^{(b)}} + \epsilon} \quad (7)$$

where:

- η is the learning rate.
- $m\hat{}_t$ and $\hat{}v_t$ are bias-corrected first and second moment estimates.
- ϵ is a small constant to prevent division by zero.

This formalization encapsulates the key components of the model, including input representation, layer architecture, activation functions, loss calculation, and optimization.

3.2. Dataset Description

The Maternal Health Risk Data used in this study is a publicly available dataset [8] sourced from six hospitals and maternity clinics in Bangladesh, a sample low-resource setting specifically designed to facilitate the prediction of maternal health risks. The dataset was collected between 2018 and 2020. This dataset contains 1,011 records, each corresponding to an individual case related to maternal health. The dataset includes seven critical features of a patient's health status during pregnancy. These features include age, Systolic Blood Pressure *(Systolic BP),* Diastolic Blood Pressure (*Diastolic BP*), Blood Sugar (BS), Body Temperature (*Body Temp*), and Heart Rate, all of which are continuous variables. The target variable, *Risk Level*, is categorical and classifies maternal health risk into three levels: low, medium, and high, providing a clear stratification of the health status for each record.

The dataset's features, selected for their direct relevance to maternal health, are common indicators used by healthcare professionals to assess the condition of pregnant women. For instance, blood pressure readings (*Systolic BP* and *Diastolic BP*) are crucial in diagnosing conditions like preeclampsia [14, 15] which can be life-threatening if not managed properly. Blood Sugar levels (BS) are also significant, as they can indicate gestational diabetes, a condition that requires careful monitoring [16]. Body Temperature and Heart Rate are vital signs that provide insights into the overall physiological state of the patient. The inclusion of these features underscores the dataset's value for predictive modeling, as it covers a comprehensive range of health indicators relevant to maternal health.

The Risk Level variable, being categorical, offers a structured way to evaluate maternal health risk, enabling the development of predictive models that can assist in early detection and intervention. The balanced representation of the different risk levels in the dataset is a key feature, allowing for robust training of machine learning models and reducing the likelihood of bias towards any particular class. This dataset thus provides a rich and diverse set of information, making it an ideal candidate for developing predictive analytics models to improve maternal health outcomes, particularly in low-resource settings where access to continuous healthcare monitoring is limited. A sample overview of the dataset is shown in Table 1.

Table 1.
Overview of the Maternal Health Risk Dataset, including key features such as Age, Blood Pressure, Blood Sugar, Body Temperature, Heart Rate, the derived BP Ratio, and the target variable, Risk Level. The table also shows the total value of each feature across the dataset.

Feature Name	Description	Data Type	Sample Value	Total Value in Dataset
Age	Age of the patient in years	Continuous	29	30,899
Systolic BP	Systolic Blood Pressure (in mm Hg)	Continuous	130	114,790
Diastolic BP	Diastolic Blood Pressure (in mm Hg)	Continuous	85	82,548
BS	Blood Sugar level (in mg/dL)	Continuous	120	99,837
Body Temp	Body Temperature (in °F)	Continuous	98.6	100,488.6
Heart Rate	Heart Rate (beats per minute)	Continuous	75	73,487
BP Ratio	Ratio of Systolic BP to Diastolic BP	Continuous	1.53	1,560.28
Risk Level	Categorical label indicating maternal health risk level	Categorical	Medium	N/A

3.3. Data Preprocessing

Data preprocessing is a critical step in preparing the dataset for training deep learning models, as it ensures that the data is in a clean and structured format, ready for analysis. This study's first step involved handling missing values, a common issue in real-world datasets. Fortunately, the Maternal Health Risk Data is complete with all the entries. This eliminated the need for imputation (replacing missing values with the mean, median, or mode) or deletion of records, allowing other steps to proceed without any loss of information. In addition to checking for missing values, the dataset underwent rigorous data cleaning to identify and correct errors. One such error was the presence of anomalous values in the Heart Rate feature, where specific entries were unrealistically low (e.g., a heart rate of 7). Such values were considered erroneous and removed to prevent them from negatively impacting the model's performance. Additionally, duplicate records were identified and removed, ensuring the dataset was free of redundancy and potential bias.

The next step was to encode the categorical variable Risk Level, which represents the maternal health risk in three levels: low, medium, and high. Since deep learning models require numerical inputs, straightforward Label Encoding was used to convert these categorical values into numerical values: Low risk was encoded as 0, Medium risk as 1, and High risk as 2. This simple yet effective encoding method allows the model to interpret the risk levels as an ordinal variable, where a higher number corresponds to a higher risk level, making the technical aspect of the process more approachable.

Finally, normalization was performed on the continuous features, including Age, Systolic BP, Diastolic BP, BS, Body Temp, Heart Rate, and the derived BP Ratio. Normalization is not just a step but a crucial aspect of deep learning. It ensures that all features contribute equally to the model's learning process, preventing any single feature from dominating due to its scale. Scaling these features to a similar range, typically between 0 and 1, helps the model converge faster during training and significantly improves its performance. This comprehensive preprocessing, with normalization at its core, ensures that the data fed into the deep learning model is well-prepared, consistent, and conducive to accurate predictions, providing reassurance about the model's performance.

A sample preprocessed dataset is shown in Table 2, illustrating the processed dataset after applying preprocessing steps, making it ready for model training. In the table, continuous features like Age, Systolic BP, Diastolic BP, BS, Body Temp, and Heart Rate have been normalized to a scale between 0 and 1, ensuring consistency in their contribution to the model. The BP Ratio feature was engineered to represent the relationship between systolic and diastolic blood pressure, potentially offering more profound insights into maternal health. The Risk Level has been encoded into numerical values (0 for Low, 1 for Medium, and 2 for High), allowing the model to process it as a numerical target variable during training. This ensures that the dataset is well-prepared for effective deep-learning model development.

3.4. Feature Engineering

Feature engineering is an essential step in the machine learning pipeline, as it involves creating new features or modifying existing ones to enhance the model's ability to make accurate predictions. This study focused on developing a new feature, BP Ratio, which represents the ratio of Systolic Blood Pressure (systolic BP) to Diastolic Blood Pressure (diastolic BP). The rationale behind this feature is grounded in the hypothesis that the relative relationship between these two blood pressure readings could provide deeper insights into a patient's cardiovascular condition, a critical aspect of maternal health. By capturing this relationship, the BP Ratio aims to offer additional information that might not be apparent from the individual blood pressure values alone, thus improving the model's ability to identify maternal health risks.

Incorporating the BP Ratio into the dataset and the original features adds a layer of complexity and richness to the input data for the deep learning model. This engineered feature allows the model to consider the dynamic interplay between systolic and diastolic pressures, which may correlate with conditions such as preeclampsia or other hypertensive disorders common in pregnancy. Enhancing the dataset with the BP Ratio hypothesizes that the model will better understand underlying health conditions, leading to more accurate maternal health risk level predictions.

Table 2.
Preprocessed Maternal Health Risk Dataset sample, showing normalized values for continuous features (Age, Systolic BP, Diastolic BP, BS, Body Temp, Heart Rate), the engineered BP Ratio feature, and the encoded Risk Level.

Age	Systolic BP	Diastolic BP	BS	Body Temp	Heart Rate	BP Ratio	Risk Level (Encoded)
0.28	0.33	0.45	0.3	0.29	0.55	1.53	0
0.45	0.67	0.6	0.75	0.72	0.65	1.56	2
0.15	0.18	0.25	0.2	0.15	0.45	1.71	1
0.35	0.40	0.55	0.4	0.35	0.58	1.45	1
0.50	0.72	0.62	0.85	0.70	0.68	1.56	2
0.18	0.25	0.35	0.32	0.28	0.48	1.43	0
0.40	0.55	0.50	0.65	0.65	0.60	1.10	2
0.22	0.30	0.40	0.22	0.22	0.50	1.50	0

The heatmap of the cleaned Maternal Health Risk dataset (Figure 2) reveals significant correlations between various features, particularly the strong positive relationship between Systolic Blood Pressure (Systolic BP) and Diastolic Blood Pressure (Diastolic BP), with a correlation coefficient of 0.79. This expected correlation underscores the interrelated nature of these blood pressure measures. Blood Sugar (BS) levels also show a moderate positive correlation with the Risk Level (0.55), indicating that higher blood sugar levels are linked to an increased risk of maternal health complications, consistent with medical knowledge. In contrast, features like Age, Body Temperature (Body Temp), and Heart Rate exhibit weaker correlations with Risk Level, suggesting a less direct impact on maternal health risk prediction. The heatmap underscores the importance of Systolic BP, Diastolic BP, and BS as critical predictors of maternal health risk, guiding the model development.

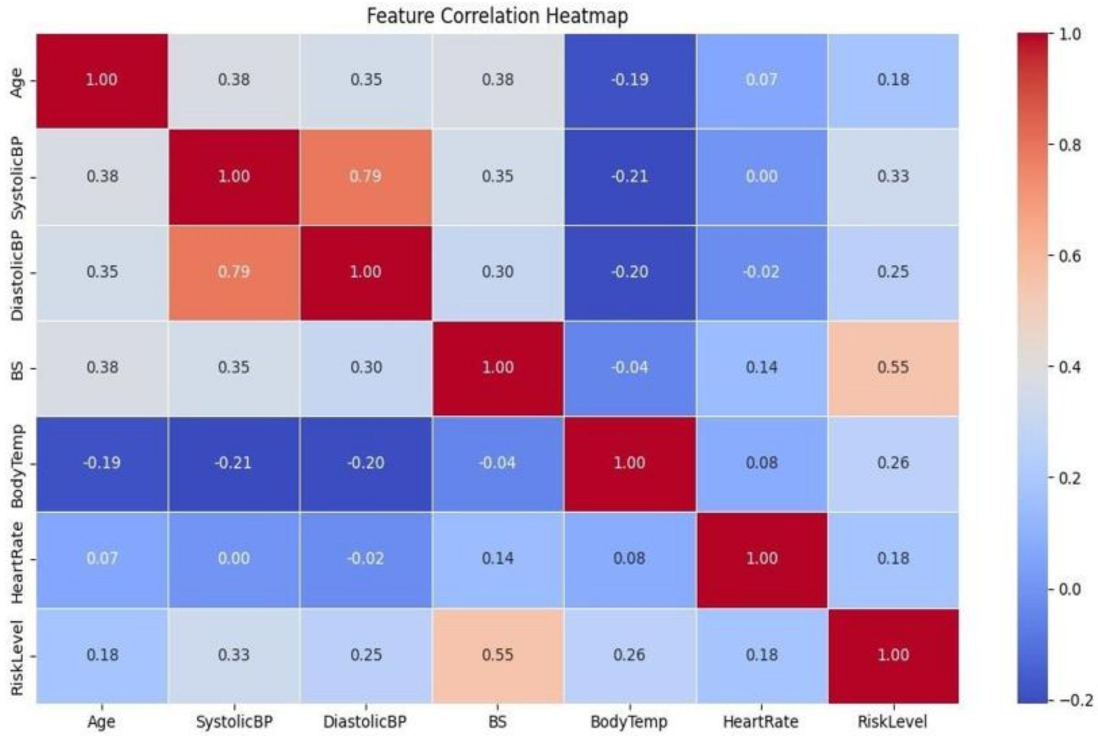

Figure 2.
Heatmap showing the correlation between various features in the Maternal Health Risk dataset.

3.5. Model development

A deep learning model, specifically a neural network, was developed to predict maternal health risk levels. Neural networks are well-suited for this task because they capture complex patterns and relationships within data. The choice of a deep learning approach allows for detecting subtle interactions between features, which traditional machine learning algorithms might miss.

The architecture of the neural network was carefully designed to balance model complexity with generalization capability, ensuring it performs well not only on the training data but also on unseen data. The network begins with an input layer comprising seven (7) neurons, each corresponding to one of the input features: Age, Systolic BP, Diastolic BP, BS, Body Temp, Heart Rate, and the engineered BP Ratio. This input layer feeds into three hidden layers, each progressively distilling the information contained in the input data.

The first hidden layer consists of 64 neurons, activated by the Rectified Linear Unit (ReLU) function. This layer captures non-linear relationships among the features, allowing the model to learn complex patterns that could indicate different risk levels. The second hidden layer, with 32 neurons utilizing the ReLU activation function, further abstracts the learned information, refining the patterns identified in the first layer. The third hidden layer, comprising 16 neurons with ReLU activation, fine-tunes these abstractions, making the model more adept at distinguishing between the different classes.

To prevent overfitting, a common challenge in deep learning, dropout layers were introduced after each hidden layer. A dropout rate of 0.5 was applied, meaning that 50% of the neurons in these layers were randomly set to zero during each training iteration. This technique helps the model generalize better by reducing reliance on any particular subset of neurons, thereby improving its performance on new data.

The output layer of the network consists of three neurons corresponding to the three possible risk levels: low, medium, and high. A softmax activation function was applied to this layer, which converts the raw output into a probability distribution over the three classes. This allows the model to output a prediction and a confidence level for each possible outcome.

The model was compiled using the Adam optimizer, a widely used optimization algorithm for efficiently handling large datasets and sparse gradients. Categorical cross-entropy was chosen as an appropriate loss function for multiclass classification tasks. The goal of this loss function is to minimize the difference between the predicted probabilities and the actual class labels. This combination of architecture, optimization, and loss function ensures that the model is both powerful and efficient, capable of delivering accurate predictions in the context of maternal health risk assessment.

3.6. Model Evaluation

The performance of the deep learning model was evaluated using several critical metrics to ensure its effectiveness in predicting maternal health risks. Accuracy was the primary metric, representing the overall proportion of correctly predicted instances out of the total. While accuracy provides a broad measure of the model's performance, it does not always capture the nuances of predictions, especially when dealing with imbalanced classes. Therefore, additional metrics like precision and recall were employed to better understand the model's effectiveness. Precision focuses on the proportion of correctly predicted positive instances (e.g., identifying high-risk cases) out of all instances predicted as positive. This metric is critical in healthcare contexts, where false positives can lead to unnecessary interventions and anxiety. Conversely, recall measures the proportion of actual positive instances correctly identified, which is crucial in minimizing the risk of missing actual high-risk cases.

The F1 − Score was also calculated to provide a balanced evaluation of the model's precision and recall. The F1 Score is the harmonic mean of precision and recall, a type of average that gives more weight to low values. This offers a single metric that balances the trade-off between the two. This is particularly valuable in healthcare applications where both the precision of predictions and the ability to capture all true positives are essential. Evaluating these metrics on a test dataset, which the model had yet to encounter during training, ensured that the assessment reflected the model's ability to generalize to new, unseen data. The combination of these metrics provided a robust evaluation framework, helping to validate the model's readiness for real-world application in predicting maternal health risks.

In addition to the deep learning model, three other machine learning models—Random Forest, Support Vector Machine (SVM) and Gradient Boosting were developed and assessed. The Random Forest model employs an ensemble of decision trees to enhance predictive accuracy, while the SVM uses a linear kernel to categorize risk levels in the feature space. On the other hand, Gradient Boosting sequentially constructs models, with each new model aiming to correct the errors of the previous ones. These models were chosen for comparison due to their proven effectiveness in classification tasks, providing a strong baseline for measuring the performance of the deep learning model. Each model underwent fine-tuning using grid search to identify the optimal hyperparameters, ensuring a fair and comprehensive comparison. Cross-validation was applied across all models to ensure reliable and stable results, reducing the risk of overfitting.

4. Experiment and Results

This section outlines the experimental setup, performance evaluation of both baseline and deep learning models, an analysis of feature importance, and a sustainability assessment of the IoT-driven solutions considered in the study.

4.1. Experimental Setup

The experiment was carefully planned to evaluate and compare the effectiveness of a deep learning model and traditional machine learning models in predicting levels of maternal health risk. The models were executed on an HP EliteBook 830 G6, equipped with an Intel Core i7 processor and 16GB of RAM, ensuring a robust environment for handling the computational demands of the analysis. This device is designed with advanced power management features, delivering high performance while minimizing energy consumption. This choice of hardware reflects a sustainable approach to computational resource use, ensuring that the model's development and execution are both effective and environmentally conscious, particularly relevant in low-resource settings where energy efficiency is crucial. The dataset was divided into training and testing sets, with 80% of the data used for training and 20% for testing. This division guarantees that the model is trained on a significant amount of data while having a separate portion for impartial evaluation. The evaluation of model efficacy in predicting maternal health risks was conducted via an analytical assessment across the pivotal metrics: accuracy, precision, recall, and F1-score, as shown in Table 3. The Random Forest algorithm demonstrated commendable performance with an average accuracy rate of 63.41%, precision at 61.77%, recall at 63.41%, and an F1-score of 62.20%. The results indicate a balanced performance across the evaluated metrics, suggesting a consistent capacity to predict the risk levels associated with maternal health accurately. Nonetheless, the observed variance across the cross-validation folds highlights a degree of sensitivity toward the training data utilized.

Contrastingly, the Support Vector Machine (SVM) model exhibited a marginally superior accuracy of 64.97% relative to the Random Forest model. Nevertheless, it registered a significantly lower precision of 52.23%, indicating that while the model efficiently identified a higher fraction of true positives, it simultaneously incurred a greater frequency of false positives. The recall metric mirrored the accuracy at 64.97%, but the F1-score was documented at 56.81%, elucidating the ramifications of diminished precision.

Surpassing both the Random Forest and SVM models, the Gradient Boosting model demonstrated superior performance with an accuracy of 67.40%, precision of 65.19%, and recall of 67.40%. The calculated F1-score of 65.74% suggests an adept balance in navigating the precision-recall trade-off. Despite a slight elevation in variance, indicative of a mild sensitivity to data distribution, the model overall delineated a more dependable performance compared to its traditional counterparts.

Among the evaluated models, the Deep Learning model proposed in this study emerged as the most productive, registering the highest performance with an accuracy of 70.96%, precision at 70.08%, and recall at 70.51%. The resultant F1-score of 64.31% affirms a remarkable consistency across these metrics, suggesting an exceptional capability in differentiating between varied levels of maternal health risk. Additionally, a lower variance was observed, implying a robust generalization capability across diverse cross-validation folds, thus rendering it the most reliable model for the specified task.

Table 3.
Performance Comparison of the Models.

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Random Forest	63.41 ± 3.28	61.77 ± 2.13	63.41 ± 3.28	62.20 ± 2.38
Support Vector Machine	64.97 ± 2.86	52.23 ± 3.33	64.97 ± 2.86	56.81 ± 1.88
Gradient Boosting	67.40 ± 3.57	65.19 ± 4.36	67.40 ± 3.57	65.74 ± 3.98
Deep Learning	70.96 ± 1.19	70.08 ± 2.53	70.51 ± 1.42	64.31 ± 1.23

Based on the results in Table 3, it is clear that the Deep Learning model surpassed the traditional machine learning models in all metrics. This suggests that the deep learning approach better captured the intricate relationships in the data, which simpler models like Random Forest or SVM might have overlooked. The Gradient Boosting model also showed strong performance, albeit slightly more variability, indicating potential reliability issues in specific scenarios.

This suggests that the deep learning model is most promising for predicting maternal health risks in low-resource settings.

The confusion matrices in Figure 3 visually represent the classification performance of the four models—Random Forest, Support Vector Machine (SVM), Gradient Boosting, and Deep Learning—on the maternal health risk prediction task. Each matrix offers insights into the number of true positive, true negative, false positive, and false negative predictions made by each model across the three risk levels: high risk, low risk, and mid risk. The Random Forest model (Figure 3(a)) shows a reasonably balanced performance, with the majority of correct predictions concentrated in the low-risk category (36 true positives). However, there are noticeable misclassifications, particularly between mid and low-risk categories. The model tends to confuse mid-risk cases with low-risk, as indicated by the 12 cases misclassified in this manner. This misclassification contributes to the relatively lower precision and recall observed for the Random Forest model. The SVM model (Figure 3(b)) demonstrates higher accuracy in identifying low-risk cases (43 true positives), but it struggles with distinguishing between the mid-risk and high-risk categories. The model frequently misclassifies mid-risk cases as low-risk (16 instances) and high-risk cases as mid-risk (7 instances). This tendency explains the lower precision observed for the SVM model, particularly in predicting the high-risk category. The Gradient Boosting model (Figure 3(c)) improves on the Random Forest and SVM models, with better accuracy in predicting low-risk and high-risk categories. The confusion between mid-risk and low-risk remains slightly reduced compared to the Random Forest model. The Gradient Boosting model's confusion matrix shows a more even distribution of errors across all categories, leading to a balanced performance as indicated by the higher F1 score.

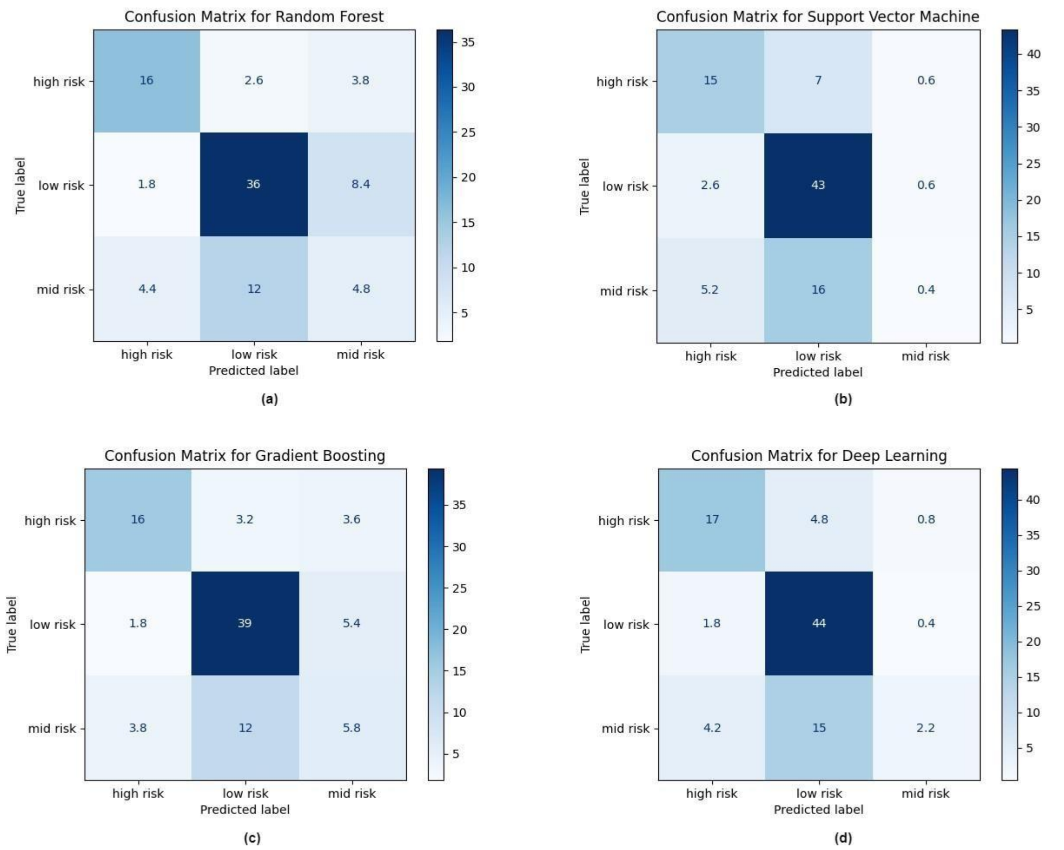

Figure 3.
Confusion Matrices for (a) Random Forest, (b) Support Vector Machine, (c) Gradient Boosting, and (d) Deep Learning Models.

The Deep Learning model confusion matrix (Figure 3(d)) outperforms the other models, particularly in correctly classifying low-risk cases (44 true positives). The model shows improved accuracy in distinguishing between high-risk and mid-risk categories, though some confusion still exists. The fewest number of misclassifications in the high-risk and mid-risk categories reflects the model's strength in handling the complexity of the data. This superior performance is consistent with the observed accuracy, precision, and recall metrics.

The ROC curves of the models in Figure 4 illustrate the trade-off between the true positive rate (sensitivity) and the false positive rate for each model across the three risk categories (low, medium, and high). The Gradient Boosting model's ROC curve (Figure 4(a)) demonstrates exceptional performance, particularly in correctly identifying low-risk cases (Class 0), with an AUC of 0.90, emphasizing its strong capability. The curves for the medium and high-risk categories also indicate satisfactory performance, with AUC values of 0.77 and 0.70, respectively, demonstrating a reasonable balance in performance across all classes. The ROC curve for the deep learning model (Figure 4(b)) exhibits solid performance with an AUC of 0.80, indicating an effective distinction between different risk levels.

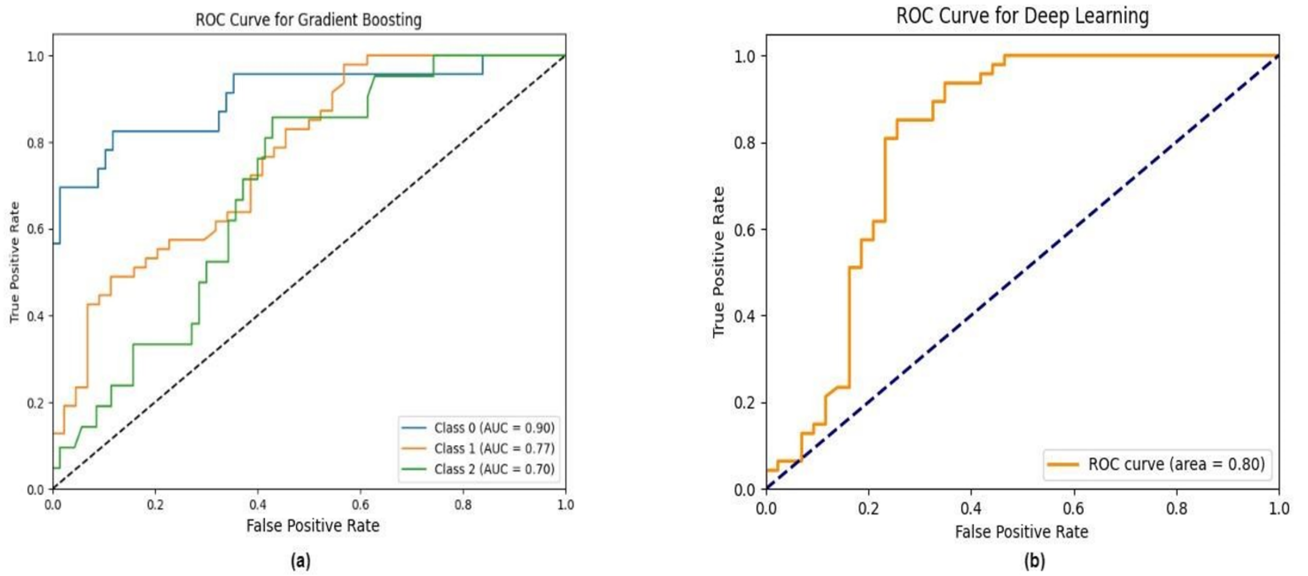

Figure 4.
ROC Curves for (a) Gradient Boosting and (b) Deep Learning Models.

4.2. Discussion

The findings of this study underscore the significant potential of IoT-driven predictive analytics, especially when combined with deep learning models, in enhancing maternal healthcare in low-resource environments. The deep learning model exhibited outstanding performance, achieving an accuracy of 70.96% and an F1-score of 64.31%, surpassing traditional machine learning models.

The findings of this study carry significant implications for healthcare providers in low-resource settings. The combination of IoT devices with deep learning models presents a potent tool for early detection and prompt intervention in maternal health risks. The enhanced accuracy and F1-score attained by the deep learning model imply that this approach has the potential to improve health outcomes, potentially lowering maternal mortality rates.

Additionally, the study's focus on sustainability is vital for ensuring that these solutions are effective, scalable, and environmentally responsible. The sustainability assessment suggests that, despite its greater computational demands, the deep learning model can be optimized to minimize energy consumption, thus establishing it as a feasible option for long-term deployment in low-resource settings.

4.2.1. Feature Importance and Sustainability Assessment

Understanding which features are most influential in predicting maternal health risks is essential. In this study, various features such as systolic and diastolic blood pressure, heart rate, age, and the BP Ratio were carefully engineered and analyzed for their significance. The findings indicated that the BP Ratio emerged as a critical feature across all models, particularly in identifying high-risk cases. This metric directly correlates to the cardiovascular health of expectant mothers, as abnormalities in blood pressure can serve as early indicators of potential complications.

Furthermore, age and heart rate consistently stood out as significant features. Advanced maternal age is a well-established risk factor, with older mothers generally facing greater health risks. Conversely, heart rate provided insights into the overall health and stress levels of the mothers, serving as potential indicators of underlying issues. Models such as Random Forest and Gradient Boosting, which naturally lend themselves to feature importance analysis, revealed the relative impact of these features. In comparison, the deep learning model further refined these insights by capturing intricate, non-linear relationships between these features and maternal health outcomes.

In addition to predicting maternal health risks, this study emphasizes the importance of sustainability in deploying IoT-driven healthcare solutions in low-resource settings. Sustainability in IoT involves considerations such as energy efficiency, device longevity, and the minimization of environmental impact. The IoT devices must be carefully selected for low power consumption and the ability to operate efficiently in scarce energy resources. By focusing on energy-efficient sensors and communication protocols, IoT solutions can be deployed in a cost-effective and environmentally friendly manner.

Despite being computationally intensive, the deep learning model was optimized to run on hardware that supports energy-efficient operations. This involved using advanced techniques such as model quantization and pruning, which reduce the computational load without significantly compromising the model's predictive accuracy. Furthermore, the scalability of the deep learning model means that it can be deployed across various settings without the need for significant hardware upgrades, thereby reducing the carbon footprint associated with large-scale implementations. The model's ability to operate effectively with smaller, energy-efficient devices ensures that the IoT solutions can be maintained over long periods with minimal environmental impact.

4.2.2. Comparison with Related Studies

The study differentiates itself from related work by emphasizing sustainability and scalability. For instance, Mondal, et al. [9] focused on real-time maternal health monitoring using IoT and machine learning, achieving high accuracy with a Random Forest Classifier. However, this approach goes beyond incorporating deep learning techniques and prioritizing sustainability in low-resource settings. Similarly, Marques, et al. [10] and Ettiyan and Geetha [5] explored IoT and machine learning for maternal health. However, this research offers a scalable solution that balances predictive accuracy with an emphasis on energy efficiency.

In further distinction, Hossain, et al. [11] and Venkatasubramanian [12] developed systems using machine learning and IoT for maternal health monitoring. While their solutions achieved high accuracy, this study uniquely addresses the challenge of sustainability, creating a model that is both effective and adaptable for long-term deployment in low-resource settings. The emphasis on energy efficiency and environmental impact sets this research apart from existing models, ensuring that technological advancements in maternal healthcare can be implemented without compromising sustainability.

The comparison of ROC curve analysis further supports these findings, with the deep learning model demonstrating a consistent area under the curve (AUC) of 0.80, indicating a strong ability to differentiate between various levels of maternal health risks. Although the Gradient Boosting model performed well, achieving an AUC of 0.90 for one class, it exhibited lower AUCs of 0.77 and 0.70 for the other two classes. This variation highlights the deep learning model's robustness and consistency across all classes, making it a more reliable choice for maternal health risk prediction.

4.2.3. Limitations

While the study demonstrates promising results, it also faces certain limitations. A key limitation is the restricted dataset size used for model training and evaluation. Although careful data preprocessing was applied, a more extensive and varied dataset could further improve the model's accuracy and generalizability. In low-resource settings, data collection challenges are common, and this limitation underscores the importance of exploring methods such as data augmentation or transfer learning to mitigate the impact of limited data.

Another notable limitation is the computational complexity associated with the deep learning model. Despite efforts to optimize energy efficiency, the model's computational demands may still be challenging in environments with limited resources. The deep learning model, while achieving higher accuracy, requires more computational power compared to the traditional models, which could pose practical challenges for deployment in extremely resource-constrained areas.

Moreover, the study primarily focused on predictive accuracy as the main performance metric, potentially overlooking other important factors such as model interpretability and ease of integration into existing healthcare systems. Although deep learning models offer superior accuracy, their "black box" nature can make them less transparent and harder for healthcare professionals to trust.

5. Conclusion

This study developed a sustainable and scalable IoT-driven predictive analytics model to improve maternal health outcomes in low-resource settings. The deep learning model demonstrated superior performance with an accuracy of 70.96%, an F1-score of 64.31%, and an AUC of 0.80, outperforming traditional machine learning models like Random Forest, SVM, and Gradient Boosting. These findings highlight the potential of deep learning to handle complex healthcare data and provide accurate maternal health risk predictions.

Despite the computational demands of the deep learning model, its alignment with sustainable practices makes it a viable option for deployment in resource-constrained environments. Future research is prioritized on expanding the datasets, reducing computational complexity, and enhancing model interpretability. This study's significant contribution to ICT for Sustainability underscores the potential of advanced analytics in addressing critical healthcare challenges in low-resource environments.

References

[1] A. Antsaklis, M. Papamichail, and P. Antsaklis, "Maternal mortality: What are women dying from?," *Donald School Journal of Ultrasound in Obstetrics and Gynecology,* vol. 14, no. 1, pp. 64-69, 2020. https://doi.org/10.5005/jp-journals-10009-1626

[2] S. Murthy and N. K. Adhikari, "Global health care of the critically ill in low-resource settings," *Annals of the American Thoracic Society,* vol. 10, no. 5, pp. 509-513, 2013. https://doi.org/10.1513/AnnalsATS.201307-246OT

[3] S. A. Spencer *et al.*, "A health systems approach to critical care delivery in low-resource settings: A narrative review," *Intensive Care Medicine,* vol. 49, no. 7, pp. 772-784, 2023. https://doi.org/10.1007/s00134-023-07084-8

[4] M. Xu, J. M. David, and S. H. Kim, "The fourth industrial revolution: Opportunities and challenges," *International Journal of Financial Research,* vol. 9, no. 2, pp. 90-95, 2018. https://doi.org/10.5430/ijfr.v9n2p90

[5] R. Ettiyan and V. Geetha, "Iod-Nets–An IoT based intelligent health care monitoring system for ambulatory pregnant mothers and fetuses," *Measurement: Sensors,* vol. 27, p. 100781, 2023. https://doi.org/10.1016/j.measen.2023.100781

[6] A. J. D. Krupa, S. Dhanalakshmi, K. W. Lai, Y. Tan, and X. Wu, "An IoMT enabled deep learning framework for automatic detection of fetal QRS: A solution to remote prenatal care," *Journal of King Saud University-Computer and Information Sciences*, vol. 34, no. 9, pp. 7200-7211, 2022. https://doi.org/10.1016/j.jksuci.2022.07.002

[7] J. R. A. Solares *et al.*, "Deep learning for electronic health records: A comparative review of multiple deep neural architectures," *Journal of Biomedical Informatics*, vol. 101, p. 103337, 2020. https://doi.org/10.1016/j.jbi.2019.103337

[8] M. Ahmed and M. A. Kashem, "IoT based risk level prediction model for maternal health care in the context of Bangladesh," in *2020 2nd International Conference on Sustainable Technologies for Industry 4.0 (STI)*, 2020: IEEE, pp. 1-6.

[9] S. Mondal, A. Nag, A. K. Barman, and M. Karmakar, "Machine learning-based maternal health risk prediction model for IoMT framework," *International Journal of Experimental Research and Review*, vol. 32, no. 1, pp. 145-59, 2023. https://doi.org/10.52756/ijerr.2023.v32.012

[10] J. A. L. Marques *et al.*, "IoT-based smart health system for ambulatory maternal and fetal monitoring," *IEEE Internet of Things Journal*, vol. 8, no. 23, pp. 16814-16824, 2020. https://doi.org/10.1109/JIOT.2020.3037759

[11] M. M. Hossain, M. A. Kashem, N. M. Nayan, and M. A. Chowdhury, "A Medical Cyber-physical system for predicting maternal health in developing countries using machine learning," *Healthcare Analytics*, vol. 5, p. 100285, 2024. https://doi.org/10.1016/j.health.2023.100285

[12] S. Venkatasubramanian, "Ambulatory monitoring of maternal and fetal using deep convolution generative adversarial network for smart health care IoT system," *International Journal of Advanced Computer Science and Applications*, vol. 13, no. 1, pp. 1-10, 2022. https://doi.org/10.14569/IJACSA.2022.0131158

[13] A. Raza, H. U. R. Siddiqui, K. Munir, M. Almutairi, F. Rustam, and I. Ashraf, "Ensemble learning-based feature engineering to analyze maternal health during pregnancy and health risk prediction," *Plos one*, vol. 17, no. 11, p. e0276525, 2022. https://doi.org/10.1371/journal.pone.0276525

[14] C. W. Redman, G. P. Sacks, and I. L. Sargent, "Preeclampsia: An excessive maternal inflammatory response to pregnancy," *American Journal of Obstetrics and Gynecology*, vol. 180, no. 2, pp. 499-506, 1999. https://doi.org/10.1016/S0002-9378(99)70239-5

[15] M. Demissie, G. Molla, A. Tayachew, and F. Getachew, "Risk factors of preeclampsia among pregnant women admitted at labor ward of public hospitals, low income country of Ethiopia; case control study," *Pregnancy Hypertension*, vol. 27, pp. 36-41, 2022. https://doi.org/10.1016/j.preghy.2021.12.002

[16] A. A. Choudhury and V. D. Rajeswari, "Gestational diabetes mellitus-A metabolic and reproductive disorder," *Biomedicine & Pharmacotherapy*, vol. 143, p. 112183, 2021. https://doi.org/10.1016/j.biopha.2021.112183