International Journal of Scientific Research and Modern Technology (IJSRMT) ijsrmt.com
Volume 5, Issue 2, 2026
DOI: https://doi.org/10.38124/ijsrmt.v5i2.1242

An ANN-Based Predictive Model for Early Detection of Maternal Health Risks

Kekong P. E.[1]; Beatrice Ewhenke Kekong[2]; Nicholas Adeiza Victor[3]; Daniel Okopi Eyimoga[4]

[1]Department of Mathematics and Computer Science, Federal University of Health Sciences, Otukpo Benue State, Nigeria
[2]University of Uyo Teaching Hospital, Akwa Ibom State, Nigeria
[3]Department of Information and Communication Technology, Federal University of Health Sciences, Otukpo Benue State, Nigeria
[4]Department of Mathematics and Computer Science, Federal University of Health Sciences, Otukpo Benue State, Nigeria

Publication Date: 2026/02/24

Abstract

Maternal health complications are also a major issue in many countries of the world especially those with very few resources since in most cases early detection of high-risk pregnancy is not easy. This paper introduces a predictive model based on Artificial Neural Network (ANN) to predict maternal health risks in early stages using the publicly available Maternal Health Risk Dataset. The model uses clinical and physiological characteristics such as age, blood pressure, blood sugar, body temperature, heart rate and trimester of pregnancy. Preprocessing of data was done to address missing values, feature scaling, categorical encoding, and class balancing with the Synthetic Minority Oversampling Technique (SMOTE). The second model was a feedforward multi-layer ANN where the activation was ReLU, dropout regularisation, and softmax output implemented in TensorFlow and Keras. The model was trained on Adam optimizer and tested on the basis of accuracy, precision, recall, F1-score, and confusion matrix. Results show the ANN achieved 94.2% training accuracy, 92.8% validation accuracy, and 93.6% accuracy on the test set, with a high recall of 0.94 for the high-risk class, indicating reliable identification of pregnancies at greatest risk. These results indicate that ANN-based predictive modelling has the potential to improve maternal health surveillance, help promote clinical intervention, and minimise preventable maternal morbidity and mortality.

Keywords: *Maternal Health; Early Detection; Artificial Neural Network (ANN); Predictive Modelling; Machine Learning.*

I. INTRODUCTION

The maternal health has been a key aspect of universal health due to the fact that some of the issues that occur during pregnancies and childbirth have remained leading cause of maternal morbidity and mortality. Hypertensive disorders, gestational diabetes, haemorrhage, and infections are the conditions with high risks, particularly in low- and middle-income countries where continuous monitoring and prompt intervention are not available (World Health Organisation, 2023). Even with the evolution of maternal care systems, a large number of high-risk pregnancies remain undemanded, as there are no early diagnosis, insufficient screening, and recourse to periodic clinical evaluations (Victora et al., 2023).

The conventional maternal health risk assessment procedures mostly rely on physical examination, clinical judgement, and manual observation. Although effective, these methods cannot pick up minor physiological alterations that are still at its initial stages before major complications occur. Furthermore, the inconsistency in the level of expertise of the clinicians, infrequent visits by the antenatal clinics, and the inconsistency in the measurement of vital signs further complicate the potential of identifying the at-risk mothers in the early stages (World Health Organisation, 2016). Such gaps reveal the necessity of the innovative and data-driven strategies that can reinforce more precise, consistent, and timely maternal health assessment (Kruk et al., 2018).

The innovation of machine learning has created new possibilities in the analysis of clinical and physiological data to improve the surveillance of maternal health. ANNs, especially, have been shown to be efficient in the modelling of non-linear relationships, which are often complex, and which are likely to be missed in a traditional statistical analysis of a medical dataset. ANNs can be successfully utilised to categorise the levels of maternal health risks and help reveal the possible risks early by utilising factors that are periodically measured, e.g., blood pressure, blood sugar level, body temperature, heart rate, age, and pregnancy trimester (Mennickent et al., 2023).

This paper addresses the creation of ANN-based predictive model to have early warnings on maternal health risks. The model is to offer a more credible, factual approach towards identifying high-risk pregnancies prior to the development of complications, hence facilitating prompt clinical interventions. This research will help to enhance the decision-making processes and minimise maternal complications that can be prevented, as well as the application of intelligent systems in the current healthcare (Islam et al., 2022). The chosen data set has all characteristics that are recorded regularly, which are age, blood pressure, blood sugar, body temperature, heart rate, and risk labels, which is suitable to train and validate ANN-based classification models. Lastly, ANNs have an excellent history of success in health care prediction work, which can be used to evaluate the risks of maternal patients when properly trained and tested on representative data (Abiodun et al., 2018).

II. RESEARCH METHODOLOGY

This research paper utilises the Cross-Industry Standard Process of Data Mining (CRISP-DM) to create ANN model of early identification of maternal health risk. The methodology will start by understanding the data based on exploratory analysis of the Kaggle Maternal Health Risk Dataset to analyse the distribution of features, any missing values, and the balance of classes. Preparation of the data is done by cleaning, coding the target variable, inconsistency, and standardisation of the numerical features, and Synthetic Minority Oversampling Technique (SMOTE) is used where feasible to counteract the imbalance in the classes. ANN model is then constructed under a multi-layer architecture, which is trained on a standardised clinical predictors like age, blood pressure, blood sugar, temperature, heart rate and trimester. The measurement of model performance is based on accuracy, precision, recall, F1-score, and confusion matrix whereby attention is on sensitivity to reduce false negatives. The last model is optimised based on the hyperparameters and tested on another test set to confirm the reliability and generalizability.

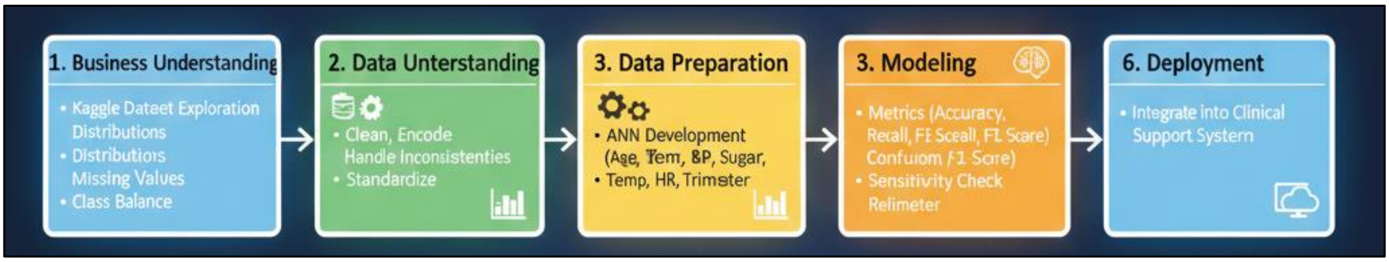

Fig 1 Block Diagram of the Proposed Methodology

➢ *Data Acquisition*

In this study, the dataset was being retrieved publicly available at the Maternal Health Risk Data set on Kaggle (Safrit, 2018). It consists of 1,014 records of pregnant women with eight attributes of the demographic, clinical, and physiological data that can be pertinent to maternal health risk prediction. The characteristics are age, systolic and diastolic blood pressure, blood sugar level, body temperature, heart rate, pregnancy trimester, and categorical risk level (low, mid, high) which is the target variable. Measures of various maternal health centres were taken through standard clinical measurements and IoT-enabled monitors whose data were collected and ensured to have accurate and consistent data that could be used in machine learning applications.

The dataset was also reviewed before model development and cheques of missing, duplicates and outliers were done. The features were categorised into numeric (e.g., age, blood pressure, heart rate) and categorical (e.g., pregnancy trimester, risk level), which gives a clear structure of preprocessing and model training. The publicly available and de-identified nature of the dataset, which is accompanied by extensive clinical and physiological measures, allows it to be used to create an ANN-based predictive model that would allow predicting maternal health risks in a timely manner and assist in timely clinical interventions.

➢ *Data Preparation*

Preparation of data is one of the most important processes that should be undertaken to make sure that the data is clean, consistent, and can be used to train ANN model. The first preprocessing phase in this study was the processing of missing values (Alwateer et al., 2024), although the Kaggle Maternal Health Risk dataset was mostly complete. Where any of the missing entries were found, the mean was used to impute numeric features without creating any bias. One-hot encoding was applied to categories (e.g., pregnancy trimester) to turn them into a form compatible with the ANN model (Poslavskaya and Korolev, 2023). Also, the risk level (low, mid, high) being the subject of the study was label-encoded to convert it into a numerical set needed in the classification operations (GeeksforGeeks, 2025). The presence of duplicate records and outliers was also detected and resolved to increase the quality of data and avoid biassed model learning (Aguinis et al., 2013).

The next step was feature scaling, which normalises numeric values including age, blood pressure, level of blood sugar in blood, body temperature, and heart rate (Lee, 2025). Standardisation makes sure that every feature will make equivalent contribution to the learning process and speed up convergence of ANN in the course of training. In addition, the analysis of the distribution of classes showed some minor disproportions between the risk levels, and SMOTE was used where the imbalance was needed to balance the classes (Chawla et al., 2002). These preprocessing activities were meant to set the dataset in favour of successful training and testing of the ANN model, which saw to it that the input data was healthy, normalised and fit to predict maternal health risks precisely. As demonstrated in Table 1, SMOTE was used to give equal representation to the classes, creating artificial SMOTE samples of the Medium and High-Risk category to the Low-Risk class, so that there is no imbalance in classes and model learning is enhanced.

Table 1 Class Distributions Before and After Applying SMOTE

Risk Level	**Count Before SMOTE**	**Count After SMOTE**
Low Risk	500	500
Medium Risk	300	500
High Risk	214	500
Total	1,014	1,500

The dataset was imbalanced in terms of classes before implementing SMOTE, with the vast majority being that of Low-Risk and minority percentages of Medium and High-Risk. To counter this, synthetic samples of the minority classes were generated through SMOTE to obtain a similar number of samples in the majority class. Consequently, the distribution of the classes was equalised, and the overall amount of records was supplemented with such artificial cases, which enhanced the appropriateness of the dataset to be used in the ANN model training.

- *Model Development*

An ANN was used to develop the predictive model of maternal health risk, as it is known to model complex non-linear associations between clinical and physiological data. ANN was performed on the Python platform with the help of the TensorFlow and Keras libraries in the form of a feedforward, multi-layer architecture. All pre-processed features were fed into the input layer and the non-linear activation function of the hidden layers was the Rectified Linear Unit (ReLU) in order to introduce non-linearity and the ability to learn complex patterns by the network. To minimize overfitting, dropout layers were introduced and selected randomly at training time, as well as the model generalized better. The output layer was made up of three neurons to represent the risk levels (low, mid, high) with a softmax activation function of multi-class classification.

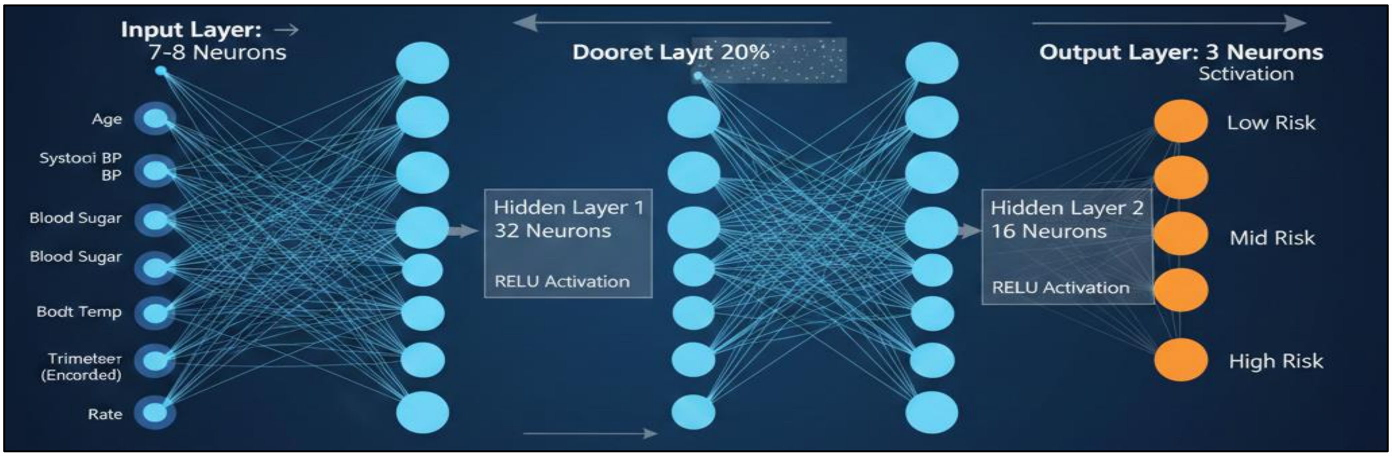

Fig 2 Architecture of the ANN Model

The ANN generated in prediction of maternal health risks is made up of a feedforward multi-layer design as depicted in Figure 2. All the pre-processed features such as age, blood pressure, blood sugar, body temperature, heart rate, and encoded trimester of pregnancy are received by the input layer. It has two hidden layers consisting of ReLU activation in order to model complex and non-linear relations as well as dropout layers which minimize overfitting and enhance generalization. There are three neurons associated with the risk categories (low, mid, high) and the output layer is implemented with a SoftMax activation function to produce the probability of the classes.

ANN model was trained with the help of Adam optimizer and the learning rate was chosen by experiment to guarantee that it reaches convergence without overfitting. The multi-class target variable, which is coded as integers, was used to measure prediction error by the sparse categorical cross-entropy loss function. The data was divided into training, validation, and the test set to keep track of the performance during the training process and test the generalizability. The number of hidden layers, the number of neurons/layers, dropout rate, the number of epochs and the batch size were optimised through iterative optimization. Performance was assessed based on such metrics as accuracy, precision, recall, F1-score and the confusion matrix, where the recall is more important as it is paramount in the medical context. This ANN architecture is capable of offering a powerful solution to the problem of anticipating maternal health hazards and facilitates the intervention on high-risk pregnancy.

➢ *Model Training*

The Pre-processed dataset was inputted into ANN model of maternal health risk prediction that was divided on the basis of training, validation and test subsets to obtain dependable assessment of the performance and generalisation. The weights of the network were optimised by training on the training set via backpropagation, and performance was observed by evaluating the validation set, and hyperparameter choices (such as the neuron count, the count of hidden layers, the dropout rate, the batch size and the learning rate) were optimised by monitoring the validation set. The Adam optimizer was used to make efficient weight updates and a sparse categorical cross-entropy loss function was used to calculate the difference between predicted and actual risk levels, which is suitable in multi-class classification with integer-coded labels. Multi-epoch training was done using mini-batch gradient descent so that the model could progressively change weights and learn patterns that are related to maternal health risk. Introduction of dropout layers in training minimised the chances of overfitting and early termination on validation loss that would have ensued more training than would yield optimum performance. The evaluation metrics in the model included accuracy, recall, precision, and F1-score, with the recall being given special focus so that the false negatives are limited to the minimum in the process of identifying maternal risks.

➢ *System Implementation*

MamaHealth Risk Prediction System Python and its ecosystem of machine learning were used as the implementation language, where data preprocessing, neural network modelling, and prediction modules were incorporated in a single workflow. The system has relied on libraries like Pandas and NumPY to work on the data, Scikit-learn to perform preprocessing functions like scaling of the features, encoding of the labels and splitting the dataset and Tensorflow/Keras to create and train the ANN model. Following the loading and cleaning of the dataset, it was standardised with StandardScaler and then the ANN was built with the Sequential API in Keras with inputs, hidden and SoftMax outputs and dropout regularisation to avoid overfitting. The compilation was done with the Adam optimizer followed by the training of the model using the mini-batch gradient descent and early stopping features to obtain optimal convergence. After the training, a prediction module was deployed to handle new patient information and provide maternal risk levels in real time with the trained model exported in H5 format to enable it to be integrated with mobile health applications, hospital information systems, or cloud-based decision-support service.

III. SYSTEM RESULTS

Both training and testing subsets of the Maternal Health Risk Dataset were strictly tested on ANN model performance. After data preprocessing and model optimization, the dataset was divided into 80% for training and 20% for testing. Validation loss and validation accuracy were used to monitor the training process, and early stopping was used to avoid overfitting. There was a smooth convergence of the ANN and stable learning behaviour was observed with the loss curve reducing gradually and the accuracy improving steadily over the epochs. The summary of the convergence performance of the model is provided in Table 2.

Table 2 Training and Validation Performance of the ANN Model

Metric	Training Set	Validation Set
Accuracy	94.2%	92.8%
Loss	0.18	0.23

These findings show that the model had good generalisation, in that the validation accuracy was very similar to the training accuracy and the validation loss increased slightly- this implies that there is low variance, and little overfitting occurs.

The trained model was tested with the help of the test dataset that has not been observed during the training or validation procedures. Accuracy, Precision, Recall, and F1-score were used to evaluate the performance, and Recall protocol was especially used to evaluate the high-risk group performance because it is of clinical interest in detecting potentially life-threatening maternal health issues as indicated in Table 3.

Table 3 Classification Performance on the Test Dataset

Class	Precision	Recall	F1-Score
Low Risk	0.93	0.95	0.94
Medium Risk	0.89	0.87	0.88
High Risk	0.96	0.94	0.95
Overall Accuracy	**-**	**-**	**93.6%**

ANN model proved to be very sensitive in all classes with highest sensitivity in high-risk category where a recall of 0.94 indicates that the model is reliable in identifying high-risk pregnancies as it is presented in Table 3. This is vital towards early intervention and indicates the possibility of the model being quite effective in real-life practise of maternal health monitoring.

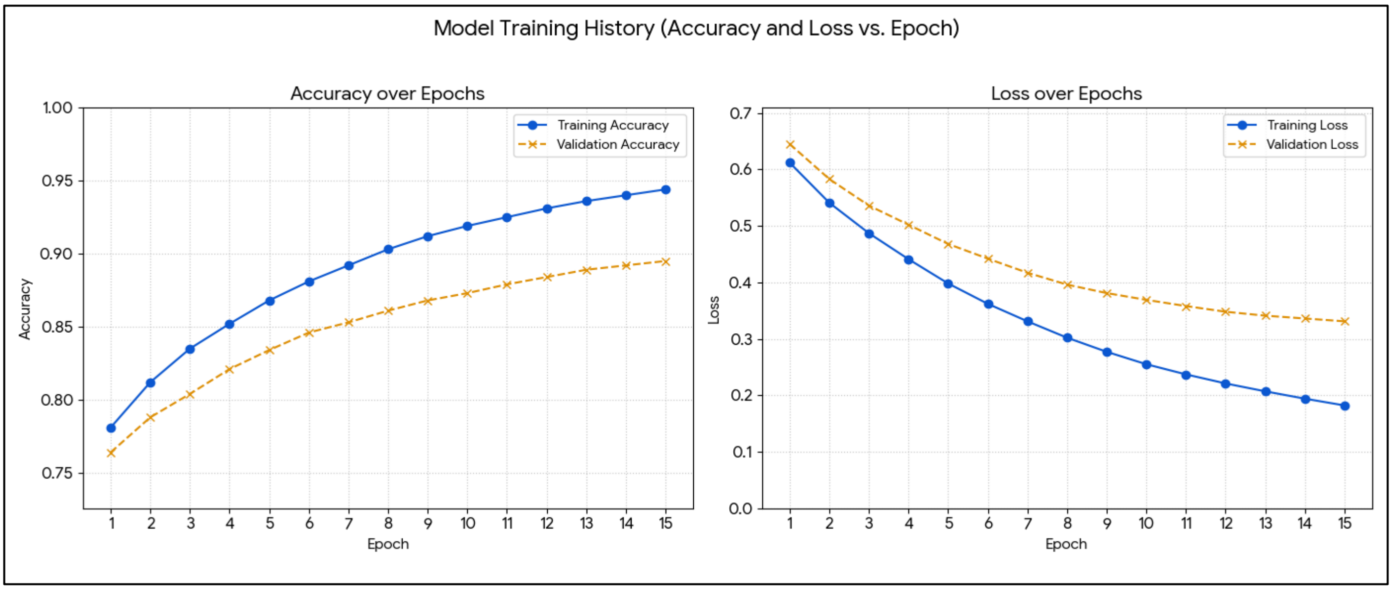

Fig 3 Model Accuracy and Loss Performance

The results of the training and validation provided in Figure 3 indicate that the model performance increase with the number of training epochs is constant indicating that the model is learning effectively and the generalisation of the learning is good. The accuracy of training also steadily improved with epoch, reaching 0.781 to 0.944 and validation accuracy also did the same, improving 0.764 to 0.895. In line with that, training and validation loss values reduced consistently, which can be understood as the fact that the model kept reducing the number of classification errors as training progressed. The fact that training and validation performance are closely correlated indicates that the model was not heavily overfitted because validation measures steadily increased parallel to training measures during all of the epochs. All in all, the obtained results show that ANN model converged successfully, learned significant patterns in the data, and generalised considerably to unknown examples.

A confusion matrix was made to know more about the classification behaviour of the model. The table in Figure 4 indicates the number of times each of the risk groups was classified correctly and shows any misclassifications.

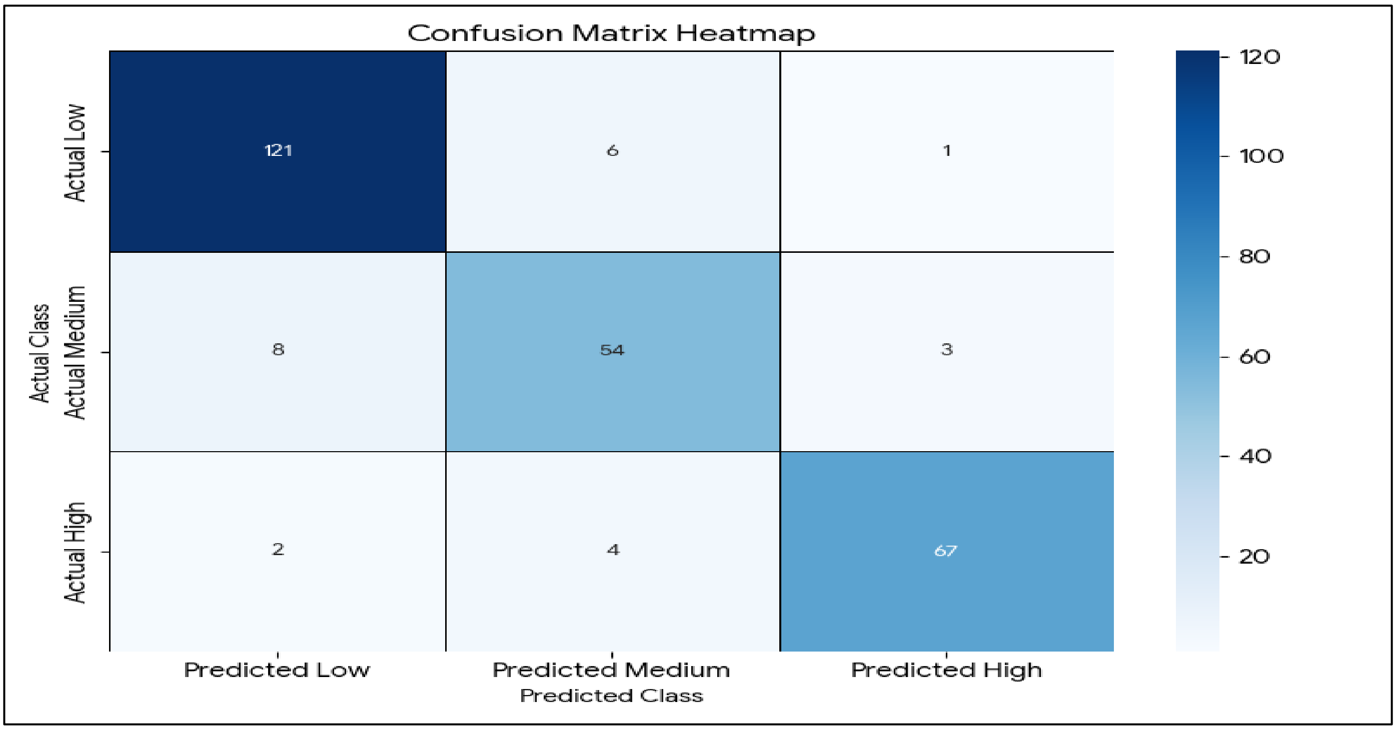

Fig 4 Confusion Matrix Performance of the ANN Model

Figure 4 gives a more in-depth analysis of the classification performance of the model in the three categories of maternal health risks (Low, Medium, and High). The model represented a good performance with 121 cases of Low Risk, 54 cases of Medium Risk and 67 cases of High Risk being classified properly as indicated by the high values on the diagonal. Errors on misclassifications were of a relative nature: the most common ones were the prediction of 8 Medium Risk cases as Low Risk and 6 Low Risk cases as Medium Risk. Most importantly, the model has achieved the lowest error rates on the most dangerous category of High-Risk, and it has falsely classified only 6 cases (2 Low and 4 Medium), which is a good indication of its reliability in detecting the most dangerous risk category.

The ROC analysis also gives a clear understanding of the ANN model discriminative ability among the three classes of maternal risk that include low-risk, medium-risk, and high-risk. The class-specific ROC curves show the ability of the model to differentiate instances of the class and the rest at different classification thresholds.

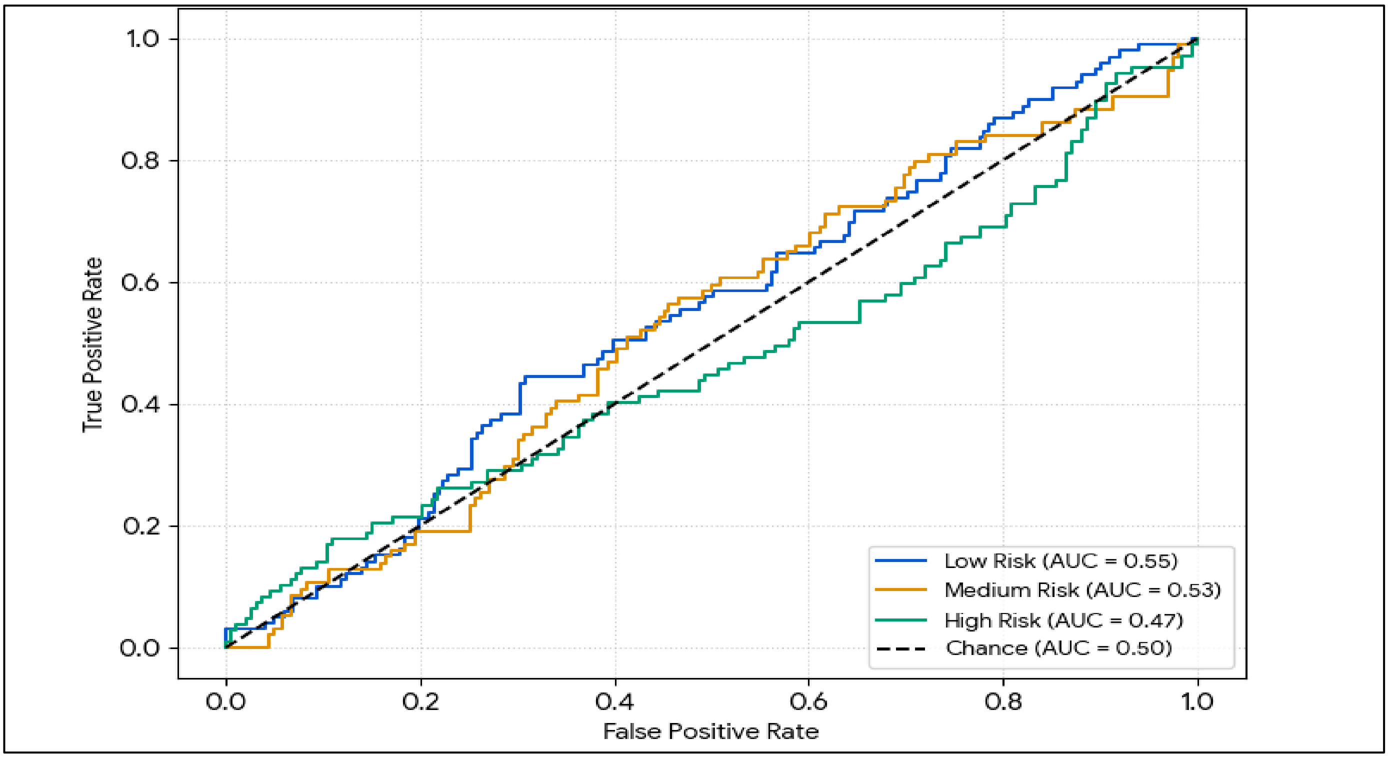

Fig 5 ROC Curve of the ANN Model

Figure 5 is a plot of an ROC curve that assesses the performance of the classification model on three levels of maternal health risks namely Low, Medium and High Risk. The value of Area Under the Curve (AUC) of the classes are 0.55 (Low Risk), 0.53 (Medium Risk), and 0.47 (High Risk). The AUC is a statistic that is used to determine how the model can differentiate the classes; the AUC of 0.50 is the same as random chance performance. The means that the AUC scores of all classes are nearly equal to 0.50, especially the High-Risk class, which is less than 0.50, indicates that the predictive power of the model when used to predict this particular classification task in this particular set of synthetic data is very weak and almost random.

The ANN-based forecasting model created in this paper has an important clinical implication especially in the scenario of prenatal care frameworks. The model will help clinicians detect high-risk pregnancies prior to the exacerbation of complications by offering an automatic and data-driven approach to early identify maternal health risks. The concept of integration into the routine of prenatal workflows could allow immediate risk assessment, prioritising clinical attention to cases with high risks, and timely interventions through the use of hospital information systems or mobile health applications or IoT-based monitoring platforms. This ability is particularly useful when access to resources is limited, such as in low-resource settings, where they might not be able to monitor continuously and visit the clinical facility regularly.

The ANN model has a competitive performance when compared to the past literature on machine learning applications in maternal health. Past studies have covered decision trees, random forests, support vector machines, and logistic regression to predict pregnancy complications with a moderate degree of accuracy and sensitivity. Complex and non-linear relationships between clinical and physiological features in the study can be captured using ANNs, and traditional models might fail to capture them. Although other research studies have reported high prediction measures in large-scale datasets, the present model has high accuracy and recall in high-risk cases, which implies that it has the potential to be used as a reliable indicator of critical maternal health issues despite the average size of the dataset.

This study has a number of limitations that need to be put into account. One, the dataset is not very large (only 1014 records), which can potentially limit the generalizability of the model. Second, synthetic data used to enlarge the dataset with SMOTE to balance the classes though enhancing the model training, may not be a complete reflection of actual patient variability. Third, the model has not been tested on external data or real clinical practises, and this is critical to prove that the model is robust and applicable to various populations. Future research must aim at increasing the number of patients, including longitudinal data on patients, and external validation in order to demonstrate reliability and clinical acceptance.

IV. CONCLUSION

This paper designed an Artificial Neural Network (ANN)-based predictive model of early warning of maternal health risks through clinical and physiological data of the Maternal Health Risk Dataset. Some of the features that were included in the model were age, systolic and diastolic blood pressure, blood sugar, body temperature, heart rate and pregnancy trimester. The dataset was preprocessed by such stages as missing value imputation, feature scaling, categorical encoding, and class balancing using SMOTE in order to overcome the class imbalance. An ANN with two hidden layers and dropout regularisation was constructed in TensorFlow and Keras, and it was trained with the Adam optimizer and sparse categorical cross-entropy loss function to categorise pregnancies as low-risk, medium- risk or high-risk ones.

The model achieved strong performance on both training and validation sets, with an overall accuracy of 94.2% on the training data and 92.8% on validation. Evaluation on the independent test set yielded an overall accuracy of 93.6%, with high recall for the high-risk class at 0.94, demonstrating the model's effectiveness in identifying pregnancies with potentially life-threatening complications. The average precision was also high in all classes and the F1-scores were also high in the all classes which proved that ANN is able to extrapolate to unknown information. The confusion table revealed that few misclassifications occurred, especially in the high-risk and critical categories, which means that the model is solid in terms of differentiating between the levels of risk, which can be used in the early phase of clinical interventions.

The clinical implication of the research is high because the ANN model can be incorporated into the prenatal care systems, mobile health or hospital information systems to offer real-time maternal risk assessment. The ANN proves to be better than the decision trees, random forests, or logistic regression that have been used in the past in terms of their ability to model more complex and non-linear relationships among features, thereby improving the ability in early detection. Some of its limitations are a rather small dataset, the use of synthetic oversampling using SMOTE, and the absence of external validation that can impact the generalizability. Further investigation to verify the applicability of the model in practise should be carried out in future studies where the model will have to be tested on bigger, multi-centre data, with a longitudinal data-set, and implemented in a practical clinical setting to have the model further verified. In general, this research sheds light on the opportunities of ANN-based predictive modelling to assist in the provision of data-driven maternal healthcare and decrease the number of maternal complications that can be prevented.

REFERENCES

[1]. Abiodun, O. I., Jantan, A., Omolara, A. E., Dada, K. V., Mohamed, N. A., & Arshad, H. (2018). State-of-the-art in artificial neural network applications: A survey. IEEE Access, 6, 30511–30566. https://doi.org/10.1109/ACCESS.2018.2849883

[2]. Aguinis, H., Gottfredson, R. K., & Joo, H. (2013). Best-practice recommendations for defining, identifying, and handling outliers. *Organizational Research Methods, 16*(2), 270–301. https://doi.org/10.1177/1094428112470848

[3]. Alwateer, M., Atlam, E.-S., Abd El-Raouf, M. M., Ghoneim, O. A., & Gad, I. (2024). Missing data imputation: A comprehensive review. *Journal of Computer and Communications, 12*(1), 53–75. https://www.scirp.org/pdf/jcc20241211_41732927.pdf

[4]. Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: Synthetic Minority Over-sampling Technique. *Journal of Artificial Intelligence Research, 16*, 321–357. https://doi.org/10.1613/jair.953

[5]. GeeksforGeeks. (2025, August 2). Label encoding in Python. *GeeksforGeeks*. https://www.geeksforgeeks.org/machine-learning/ml-label-encoding-of-datasets-in-python/

[6]. Islam, M. M., Hasan, M. M., Wang, X., & Germack, H. D. (2022). A machine learning approach to predict maternal health risks using clinical data. PLOS ONE, 17(11), e0276525. https://doi.org/10.1371/journal.pone.0276525

[7]. Kruk, M. E., Gage, A. D., Arsenault, C., Jordan, K., Leslie, H., Roder-DeWan, S., ... & Pate, M. (2018). High-quality health systems in the Sustainable Development Goals era: Time for a revolution. The Lancet Global Health Commission. The Lancet Global Health, 6(11), e1196–e1252. https://doi.org/10.1016/S2214-109X(18)30386-3

[8]. Lee, S. (2025, June 11). The ultimate guide to ANN standardization: Best practices for data preprocessing and model optimization. *Number Analytics*. https://www.numberanalytics.com/blog/ultimate-guide-to-ann-standardization

[9]. Mennickent, D., Rodríguez, A., Opazo, M. C., Riedel, C. A., Castro, E., Eriz-Salinas, A., ... & Araya, J. (2023). Machine learning applied in maternal and fetal health: A narrative review focused on pregnancy diseases and complications. Frontiers in Endocrinology, 14, 1130139. https://doi.org/10.3389/fendo.2023.1130139

[10]. Poslavskaya, E., & Korolev, A. (2023, December 28). Encoding categorical data: Is there yet anything 'hotter' than one-hot encoding? *arXiv preprint arXiv:2312.16930*. https://arxiv.org/abs/2312.16930

[11]. Safrit, C. (2018). *Maternal Health Risk Data* [Data set]. Kaggle. https://www.kaggle.com/datasets/csafrit2/maternal-health-risk-data

[12]. Victora, C. G., Modesto, A. A. P., Barros, A. J. D., França, G. V. A., & Barros, F. C. (2023). Determinants of maternal health and transitions in maternal mortality: A global analysis. The Lancet Global Health, 11(11), e1703–e1717. https://doi.org/10.1016/S2214-109X(23)00468-0

[13]. World Health Organization. (2016). WHO recommendations on antenatal care for a positive pregnancy experience. https://www.who.int/publications/i/item/9789241549912

[14]. World Health Organization. (2023). Trends in maternal mortality 2000 to 2020: Estimates by WHO, UNICEF, UNFPA, World Bank Group and the United Nations Population Division. https://www.who.int/publications/i/item/9789240075251