Logapriya A, 2026, 14:5
ISSN (Online): 2348-4098
ISSN (Print): 2395-4752

International Journal of Science, Engineering and Technology
An Open Access Journal

Gestational Diabetes Mellitus Prediction Using Machine Learning

(Catboost-Based Intelligent Risk Assessment System)

Logapriya A
M.Sc Data Analytics, Department Of Computer Science, Sri Ramakrishna College Of Arts & Science For Women, Coimbatore – 641044, Tamil Nadu, India

ABSTRACT- Gestational Diabetes Mellitus (GDM) is an increasingly serious maternal health concern, associated with high birth weight, premature delivery and a higher likelihood of Type-2 diabetes later in life. Conventional diagnosis relies on the Oral Glucose Tolerance Test (OGTT), which is usually performed late in pregnancy and needs laboratory infrastructure, delaying early intervention. Existing machine learning approaches, including XGBoost-based models, focus mainly on accuracy and depend heavily on preprocessing, while offering little in the way of risk categorization, visualization or patient record management. To address these limitations, this study presents a machine learning framework for early GDM prediction using the CatBoost classifier. The proposed system combines clinical and lifestyle parameters such as age, number of pregnancies, BMI, HDL, systolic and diastolic blood pressure, OGTT, hemoglobin, family history, PCOS, sedentary lifestyle and prediabetes status. Data preprocessing steps, namely missing-value imputation, removal of identifier and weak features, and standardization, are applied to improve data reliability. The trained model classifies patients as GDM or Non-GDM, estimates a probability-based risk percentage and assigns a risk category. A user-friendly Streamlit web application provides real-time prediction, a risk meter, an interactive dashboard, patient history storage and downloadable PDF reports. The model achieved an accuracy of approximately 96% on held-out data, with a confusion matrix showing only 40 misclassifications out of 3,525 records. Overall, the system offers an accurate, scalable and practical decision-support tool for early detection of gestational diabetes and improved maternal healthcare.

Keywords: Gestational Diabetes, CatBoost, Machine Learning, Risk Prediction, Healthcare Analytics

I. INTRODUCTION

Gestational Diabetes Mellitus (GDM) is a condition in which blood glucose levels become elevated during pregnancy. It is a growing health concern worldwide and, if not detected and managed early, can lead to complications for both mother and baby, including high birth weight, premature delivery and an increased risk of developing Type-2 diabetes in the future [1], [2]. In rural and resource-limited settings, access to timely laboratory diagnosis is often restricted, which further reduces the opportunity for preventive care.

Traditional diagnosis depends mainly on the Oral Glucose Tolerance Test (OGTT), performed in the later stages of pregnancy. In addition, risk factors such as age, Body Mass Index (BMI), blood

pressure, family history, lifestyle habits and conditions such as PCOS act together in complex ways. Manual analysis of these factors is time-consuming, error-prone and dependent on clinical experience [3].

Recent developments in machine learning have enabled efficient analysis of multiple clinical variables at once, leading to improved disease prediction [4]. Among these techniques, gradient boosting methods have attracted considerable attention because of their strong predictive capability on structured data. CatBoost, in particular, handles categorical and numerical features effectively and reduces overfitting [5]. Therefore, integrating CatBoost with clinical data offers a robust and scalable solution for early GDM prediction.

1. Problem Statement

Early detection of GDM is hindered by three practical problems. First, OGTT-based diagnosis occurs late in pregnancy and needs laboratory facilities, so high-risk women are often identified only after complications have begun. Second, risk depends on many interacting

© 2026 A, This is an Open Access article distributed under the terms of the Creative Commons Attribution License (http://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly credited.

factors at once, including age, BMI, blood pressure, lipid profile, family history, PCOS and lifestyle, and manual assessment of all of them is slow and inconsistent. Third, most published prediction models stop at a classification score and are never deployed as tools that clinicians can use, so they provide no risk categories, visual explanation or patient record management. There is therefore a need for an accurate, automated and accessible system that predicts GDM early, explains the risk clearly and keeps patient records for follow-up.

2. Objectives

The main objective of this work is to develop an intelligent system for early prediction of GDM using machine learning. The specific objectives are:

- To build a CatBoost classification model that predicts GDM from clinical and lifestyle features.
- To apply preprocessing, namely missing-value imputation, feature selection and standardization, to improve model performance.
- To estimate a probability-based risk percentage, a future Type-2 diabetes risk and a risk category for each patient.
- To develop a Streamlit web application for real-time prediction, a risk meter and an analysis dashboard.
- To store patient records in a database and generate downloadable PDF and CSV reports for clinical documentation.
- To evaluate the model using accuracy, precision, recall, F1-score and confusion matrix analysis.

3. Scope Of The Work

The scope covers structured clinical data of pregnant women, including age, BMI, OGTT, blood pressure, HDL, hemoglobin, family history, PCOS, sedentary lifestyle and prediabetes status. The system provides real-time prediction, risk categorization, visualization, patient history management and report generation through a web application. It is intended as a decision-support

tool for healthcare professionals and patients and does not replace professional medical diagnosis. Its predictions are limited to the patterns present in the training dataset.

4. Organization Of The Paper

Section 2 reviews related work. Section 3 describes the proposed system, including the dataset, preprocessing, feature selection, classification, model configuration, system modules and testing. Section 4 presents and discusses the results. Section 5 concludes the paper and outlines future work, and Section 6 lists the references.

II. LITERATURE REVIEW

Research in GDM prediction has increasingly shifted toward machine learning and data-centric methodologies. Foundational contributions have played a key role in this transition. Breiman (2001) developed the Random Forest algorithm, which improves generalization by aggregating multiple decision trees [6]. Friedman (2001) proposed Gradient Boosting Machines, which build models sequentially to correct earlier errors and form the basis of modern boosting algorithms such as XGBoost and CatBoost [7].

Building upon these techniques, several studies have applied machine learning to GDM prediction. Smith, Johnson and Lee (2020) used Logistic Regression, Decision Trees and Random Forest on clinical data such as BMI, age and glucose levels, and showed that ensemble models provide better accuracy than traditional methods [8]. Li, Chen and Zhou (2022) evaluated Logistic Regression and Random Forest using accuracy and sensitivity and reported that ensemble methods handle complex medical datasets better [9]. Sharma and Gupta (2021) proposed a data mining approach using classification and clustering on patient health records to identify high-risk patients [10].

More recent work has focused on boosting-based models. Zhang, Wang and Li (2022) compared SVM, KNN and XGBoost for diabetes prediction and concluded that boosting algorithms achieve higher accuracy because they handle complex data patterns effectively [11]. Zhang, Wang and Li (2025) applied XGBoost to clinical and demographic features of GDM and observed better accuracy and feature-importance interpretation than traditional models [12]. Kumar and Patel (2023) introduced an AI-based risk assessment system using blood pressure, OGTT and lifestyle factors that produced risk scores and risk levels for clinicians [13].

Beyond modeling, interpretability and clinical guidance are critical. Lundberg and Lee (2017) introduced SHAP for explaining model decisions, which is important for trust in medical applications [14]. Clinical guidelines from the American Diabetes Association (2022) identify age, obesity, family history and glucose levels as key risk factors [15], while the World Health Organization (2021) highlights the rising prevalence of GDM and the importance of early detection [16]. Tools such as Streamlit enable interactive deployment of such models [17].

Despite these advancements, certain limitations remain. Most studies focus only on accuracy, depend on extensive preprocessing, and lack risk categorization, visualization and patient history management. Therefore, integrating a CatBoost-based model with a real-time web application and record management presents a promising direction for practical GDM screening.

2.1 Existing Methods and Research Gap

Existing GDM prediction systems rely mainly on Logistic Regression, Decision Trees, Random Forest and, more recently, XGBoost. These methods analyze clinical features such as age, BMI, blood pressure and glucose levels, and boosting models capture complex relationships better than traditional ones. However, they typically require extensive preprocessing, including imputation and encoding of categorical variables, and some are sensitive to class imbalance. Real clinical datasets are often incomplete and inconsistent, which makes such pipelines harder to apply in practice.

A second limitation is the gap between model and application. Most studies report accuracy and related metrics but do not integrate the model into a real-time interface that healthcare professionals can use. Table 1 summarizes the capabilities of existing approaches against the proposed system, based on the studies reviewed above.

Table 1: Comparison of existing GDM prediction approaches and the proposed system

Approach	Representative study	Real-time web application	Risk categorizatio n	Visualization / risk meter	Patient history and reports
Logistic Regression, Decision Tree, Random Forest	Smith et al. (2020); Li et al. (2022)	Not reported	Not reported	Limited	Not reported
XGBoost	Zhang et al. (2025)	Not reported	Not reported	Feature importance	Not reported
AI-based risk scoring	Kumar and Patel (2023)	Partial	Yes	Limited	Not reported
Proposed CatBoost system	This work	Yes (Streamlit)	Yes (low, moderate, high)	Yes (gauge, dashboard)	Yes (database, PDF, CSV)

The key gaps are therefore a focus on accuracy alone, missing risk categorization and visualization, and no patient history management. The proposed work addresses all three.

III. PROPOSED WORK

This work presents a Gestational Diabetes Prediction System using CatBoost, developed to support early detection, risk assessment and monitoring of GDM through a unified, data-centric framework. The system integrates a machine learning model, risk analysis and an interactive web interface within a single platform, addressing the shortcomings of conventional diagnosis and earlier prediction studies.

The framework begins with data collection and preprocessing, where the clinical dataset is cleaned, identifier and weak features are removed, missing values are imputed and numerical features are

standardized. The processed data is divided into training and testing sets, and a CatBoost classifier is trained to predict whether a patient has GDM. Beyond

the binary prediction, the system computes a probability-based risk percentage, estimates future Type-2 diabetes risk and assigns a risk category (low, moderate or high).

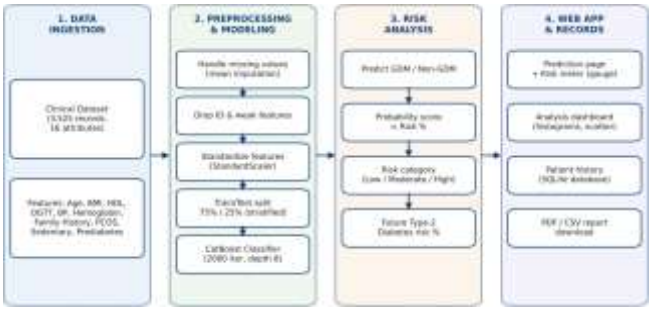


Figure 1: Overall Workflow of the GDM Prediction System

The trained model is deployed through a Streamlit application with four modules: Home, Prediction, Dashboard and Patient History. The Prediction module accepts patient details and returns instant results with a risk meter (gauge chart). The Dashboard presents histograms and scatter plots of patient trends. The Patient History module stores records in an SQLite database and supports CSV download, while a report module generates a downloadable PDF for clinical documentation. The modular design allows future integration of hospital data, IoT-based monitoring and mobile access.

IV. METHODOLOGY

4.1 Dataset

This study uses a Gestational Diabetes dataset containing 3,525 clinical records of pregnant women, each described by 16 attributes plus a binary class label. Of these records, 2,153 (61.1%) are Non-GDM and 1,372 (38.9%) are GDM, indicating a moderate class imbalance. The dataset was compiled from publicly available healthcare datasets and research-based sources such as Kaggle and clinical study references. The attributes include age, number of pregnancies, gestation in previous pregnancy, BMI, HDL, family history, unexplained prenatal loss, large child or birth defect, PCOS, systolic and diastolic blood pressure, OGTT, hemoglobin, sedentary lifestyle and prediabetes status.

Table 2: Summary of key dataset attributes

Attribute	Type	Mean	Range
Age (years)	Numeric	32.6	20 – 45
BMI	Numeric	27.9	13.3 – 45.0
HDL (mg/dL)	Numeric	46.5	15 – 70
Systolic BP (mmHg)	Numeric	135.8	90 – 185
Diastolic BP (mmHg)	Numeric	81.5	60 – 124
OGTT (mg/dL)	Numeric	170.7	80 – 403
Hemoglobin (g/dL)	Numeric	14.0	8.8 – 18.0
Family history, PCOS, sedentary lifestyle, prediabetes	Binary	–	0 / 1

4.2 Data Preprocessing

Data preprocessing ensures the quality and consistency of the input data. The dataset contained missing values in BMI (1,081), HDL (1,001), systolic BP (1,705) and OGTT (513), which were handled using mean imputation. The identifier column (Case Number) was removed because it carries no predictive information. The attributes unexplained prenatal loss and large child or birth defect were found to be weak predictors and were dropped. The class label was encoded numerically (GDM = 1, Non-GDM = 0), and numerical features were standardized using StandardScaler, which was saved for reuse during real-time prediction. Finally, the data was split into 75% training and 25% testing subsets using stratified sampling (random state 42) to preserve the class ratio.

4.3 Feature Extraction

Feature selection identifies the most influential variables for GDM. The final model uses age, number of pregnancies, gestation in previous pregnancy, BMI,

HDL, family history, PCOS, systolic and diastolic blood pressure, OGTT, hemoglobin, sedentary lifestyle and prediabetes as inputs. Weak and identifier features were eliminated to reduce dimensionality, improve computational efficiency and enhance interpretability. CatBoost's built-in feature importance is used to rank the contribution of each attribute, and the results show that OGTT, HDL, blood pressure, BMI and prediabetes have the greatest influence on the prediction outcome.

4.4 Classification

The classification stage predicts whether a patient has GDM or not. The CatBoost classifier was selected as the primary technique because it is a gradient boosting algorithm that handles numerical and categorical data effectively, reduces overfitting and delivers high accuracy on structured data. The model was configured with 2,000 iterations, a tree depth of 8, a learning rate of 0.03 and the Logloss loss function. Class weights of [1, 3] were applied to give greater importance to GDM cases and reduce missed diagnoses. The trained model outputs a probability score, which is

converted into a risk percentage and a risk category. Performance is evaluated using accuracy, precision, recall, F1-score and confusion matrix analysis.

4.5 CatBoost Algorithm and Model Configuration

CatBoost is a gradient boosting algorithm that builds an ensemble of decision trees sequentially, where each new tree corrects the errors of the previous ones. It uses ordered boosting to reduce target leakage and prediction shift, which lowers overfitting on small and medium-sized tabular datasets. For a feature vector x, the ensemble prediction after M trees is the sum of the tree outputs:

$$FM(x) = \sum_{m=1}^{M} \eta\, hm(x)$$

where h_m is the m-th tree and η is the learning rate. For binary classification the raw score is converted into the probability of GDM by the logistic function, and the model is trained by minimizing the Logloss:

$$P(y = 1 \mid x) = \frac{1}{1 + e^{-FM(x)}}$$

$$L = -\frac{1}{N} \sum_{i=1}^{N} [yi \log pi + (1 - yi) \log(1 - pi)]$$

The predicted probability is used directly as the risk percentage shown on the risk meter and is mapped to a risk category (low, moderate or high). Table 3 lists the model configuration used in this work.

Parameter	Value
Algorithm	CatBoostClassifier
Iterations	2000
Tree depth	8
Learning rate	0.03
Loss function	Logloss
Evaluation metric	Accuracy
Class weights	[1, 3] (Non-GDM, GDM)
Train / test split	75% / 25% (stratified, random state 42)
Feature scaling	StandardScaler

Table 3: CatBoost model configuration

4.6 System Modules

The system is organized into eight modules that together cover the full path from raw data to a clinical report.

1. Data Collection Module: gathers the clinical records of pregnant women that serve as the basis for training and prediction.

2. Data Preprocessing Module: handles missing values, removes irrelevant attributes and applies feature scaling so that the data is consistent.

3. Model Training Module: trains the CatBoost classifier on the preprocessed dataset and saves the trained model and scaler for reuse.

4. Prediction Module: accepts patient details from the web interface, scales them with the saved scaler and returns the GDM or Non-GDM result with a probability score.

5. Risk Analysis Module: converts the probability into a risk category and estimates the future Type-2 diabetes risk to support preventive care.

6. Visualization Module: draws the risk meter (gauge chart) and the dashboard charts that explain the result.

7. Web Application Module: provides the Streamlit interface with Home, Prediction, Dashboard and Patient History pages.

8. Database Management Module: stores patient names, key measurements and results in an SQLite database for later retrieval and download.

The model was developed in Google Colab, where preprocessing, feature selection and training were carried out, and the trained artifacts were then loaded by the Streamlit application. Patient input is validated, scaled with the same scaler used in training and passed to the model, so training and deployment follow an identical pipeline.

4.7 Testing

Five levels of testing were applied to verify the correctness and reliability of the system. Unit testing checked individual functions such as preprocessing, scaling, prediction and report generation, each against known inputs. Integration testing confirmed the data flow from the Streamlit input form through the model to the displayed output and the database. System testing evaluated the complete application with different patient inputs. Model testing measured performance on unseen test data to confirm generalization. Finally, user interface testing verified input fields, buttons and result displays for ease of use.

V. RESULTS

The proposed Gestational Diabetes Prediction System was assessed using the integrated clinical dataset described above. After preprocessing and stratified splitting, the CatBoost model achieved an accuracy of approximately 96% on unseen test data, indicating that it can effectively predict the presence of GDM from clinical input features.

To evaluate the model across all records, a confusion matrix was also computed over the complete dataset of 3,525 patients (Figure 2). The model correctly identified 1,351 GDM cases (true positives) and 2,134 Non-GDM cases (true negatives), with only 19 false positives and 21 false negatives. These values correspond to an overall accuracy of 98.87%, precision of 98.61%, recall (sensitivity) of 98.47%, specificity of 99.12% and an F1-score of 98.54%. As this matrix includes records used during training, the hold-out accuracy of approximately 96% is the more conservative estimate of real-world performance.

Metric	Formula	Value
Accuracy	(TP+TN)/(TP+TN+FP+FN)	98.87%
Precision	TP/(TP+FP)	98.61%
Recall	TP/(TP+FN)	98.47%
Specificity	TN/(TN+FP)	99.12%
F1-Score	2×P×R/(P+R)	98.54%

Table 4: Performance metrics derived from the confusion matrix (n = 3,525)

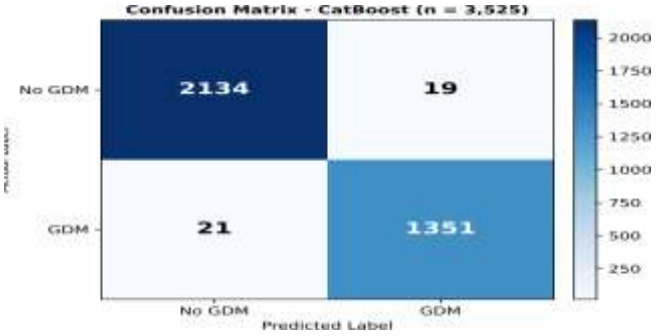


Figure 2: Confusion Matrix

The low number of false negatives (21) is particularly valuable in a medical setting, because a missed GDM case can delay diagnosis and treatment and lead to complications for mother and baby. The low number of false positives (19) also reduces unnecessary stress and follow-up procedures for healthy individuals. The use of class weights contributed to this balanced behavior despite the moderate class imbalance.

Feature importance analysis showed that OGTT, HDL, blood pressure, BMI and prediabetes status have a significant impact on prediction, which agrees with established clinical risk factors [15]. Probability-based predictions are converted into risk categories, supporting early diagnosis and clinical decision-making. For example, a patient with age 30, BMI 31, OGTT 144, family history, PCOS, sedentary lifestyle and prediabetes was classified as GDM with high risk, whereas a patient with age 19, BMI 21 and OGTT 117 and no risk factors was classified as Non-GDM.

The Streamlit application provides real-time prediction with a risk meter, an analysis dashboard of age, BMI and age-versus-OGTT distributions (Figure 3), a patient history module backed by a database, and downloadable PDF and CSV reports. The dashboard

shows that the dataset has an average age of 32.6 years, average HDL of 46.5 and average systolic BP of 135.8, and that OGTT values tend to be higher in older age groups.

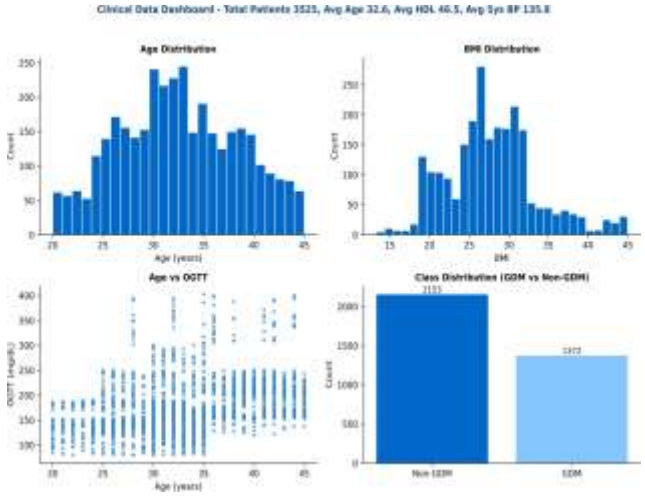


Figure 3: Analysis Dashboard

5.1 Evaluation Metrics

The classifier was evaluated using four standard metrics derived from the confusion matrix, where TP, TN, FP and FN denote true positives (correctly predicted GDM), true negatives (correctly predicted Non-GDM), false positives (Non-GDM predicted as GDM) and false negatives (missed GDM cases). Accuracy measures overall correctness, precision measures how many predicted GDM cases are real, recall measures how many real GDM cases are found, and the F1-score balances precision and recall.

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}$$

$$\text{Precision} = \frac{TP}{TP + FP}$$

$$\text{Recall} = \frac{TP}{TP + FN}$$

$$F1 = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

In medical screening, recall is especially important because a missed GDM case can delay diagnosis and treatment, while precision limits unnecessary anxiety and follow-up tests for healthy women. The F1-score is therefore used to confirm that the model keeps both errors low at the same time.

5.2 Confusion Matrix Analysis

Table 5 gives the confusion matrix over all 3,525 records. Of the 2,153 Non-GDM patients, 2,134 were classified correctly and 19 were flagged as GDM. Of the 1,372 GDM patients, 1,351 were detected and 21 were missed. The two error types are almost equal in number, which shows that the class weights of [1, 3] did not push the model toward over-predicting GDM.

Actual \ Predicted	Non-GDM	GDM	Total
Non-GDM	2134 (TN)	19 (FP)	2153
GDM	21 (FN)	1351 (TP)	1372
Total	2155	1370	3525

Table 5: Confusion matrix of the CatBoost model (n = 3,525)

5.3 Feature Importance and Risk Categorization

CatBoost feature importance showed that OGTT, HDL, blood pressure, BMI and prediabetes status contribute most to the prediction. This is consistent with clinical knowledge, since impaired glucose tolerance, abnormal lipid levels, hypertension and excess weight are established risk factors for GDM [15]. The agreement between the learned importance and clinical guidelines increases confidence that the model relies on meaningful medical signals rather than artifacts of the dataset.

For every prediction the system also reports a risk percentage, a risk category and a gauge-style risk meter whose zones are green (0–30%), orange (30–60%) and red (60–100%). This turns a raw probability into an easily understood message for patients and clinicians, so that high-risk cases can be referred for early testing and lifestyle counselling.

5.4 Patient Records and Reports

Every prediction is stored in an SQLite database together with the patient name, age, BMI, OGTT value and result. The Patient History page lists these records, shows a bar chart of GDM versus Non-GDM counts and allows the full history to be downloaded as a CSV file, while a PDF report can be generated for each individual prediction. These features support long-term monitoring and clinical documentation, which existing prediction studies rarely provide.

1. Comparison With Existing Work

Previous studies based on Logistic Regression, Random Forest and XGBoost demonstrate that ensemble and boosting models outperform traditional approaches, and the present results agree with this finding. The contribution of this work lies in combining a boosting model that handles mixed data types and limits overfitting with a complete application layer. As Table 1 shows, the proposed system adds risk categorization, visualization, patient history and report generation that were not reported in the reviewed studies. A direct numerical comparison with other studies is not made here because the datasets and evaluation protocols differ.

2. Limitations

The study has several limitations. The dataset comes from public and research-based sources and may not represent every population, so external validation on hospital data is needed before clinical use. The confusion matrix covers the full dataset, including records used for training, so the hold-out accuracy of approximately 96% is the safer estimate of generalization. Missing values were filled using mean imputation, which can reduce variability in features such as BMI, HDL and systolic blood pressure. Finally, the system supports clinical decisions but does not replace laboratory diagnosis by a qualified physician.

3. Summary

In summary, the proposed framework demonstrates strong potential for accurate and practical GDM screening. Compared with existing XGBoost-based studies that focus mainly on accuracy, the proposed system additionally provides risk categorization, visualization, patient history management and automated reporting.

VI. CONCLUSION

This work presents a machine learning-based framework for early prediction of Gestational Diabetes Mellitus using clinical and lifestyle data. The application of the CatBoost classifier enables accurate classification of patients into GDM and Non-GDM categories, with an accuracy of approximately 96% on unseen data and very few misclassifications overall. Data preprocessing, feature selection and standardization enhance model stability and reliability.

The accompanying Streamlit web application provides real-time prediction, risk percentage, risk categorization, a visual risk meter, an interactive dashboard, patient history storage and automated PDF reports, making the system practical for healthcare use. By identifying high-risk patients at an early stage, the system supports timely medical care, lifestyle changes and preventive measures, helping to reduce risks for both mother and baby.

The system is intended as a decision-support tool and does not replace professional medical diagnosis. As a future enhancement, it can be extended by integrating real-time hospital data, IoT-based health monitoring devices, mobile applications, larger multi-center datasets and explainability methods such as SHAP to further improve accuracy and clinical trust.

REFERENCES

1. World Health Organization (WHO), Gestational Diabetes Mellitus – Guidelines and Information, 2022.
2. American Diabetes Association (ADA), Standards of Medical Care in Diabetes, 2022.
3. Kumar, R., & Patel, S., AI-Based System for Gestational Diabetes Risk Prediction, 2023.
4. Pedregosa, F., et al., Scikit-learn: Machine Learning in Python. Journal of Machine Learning Research, 12, 2011.
5. Dorogush, A. V., Ershov, V., & Gulin, A., CatBoost: Gradient Boosting with Categorical Features Support, 2018.
6. Breiman, L., Random Forests. Machine Learning, 45(1), 5–32, 2001.
7. Friedman, J. H., Greedy Function Approximation: A Gradient Boosting Machine. Annals of Statistics, 29(5), 2001.
8. Smith, P., Johnson, A., & Lee, R., Machine Learning Techniques for Gestational Diabetes Prediction, 2020.
9. Li, K., Chen, L., & Zhou, H., Prediction of Gestational Diabetes Using Machine Learning Models, 2022.
10. Sharma, K., & Gupta, M., Data Mining Approach for Early Detection of Gestational Diabetes, 2021.

11. Zhang, L., Wang, Y., & Li, X., Comparison of Machine Learning Algorithms for Diabetes Prediction, 2022.
12. Zhang, Y., Wang, J., & Li, X., Machine Learning Approach for Gestational Diabetes Prediction using XGBoost, 2025.
13. Kumar, R., & Patel, S., AI-Based Risk Assessment of Gestational Diabetes Using Multiple Health Parameters, 2023.
14. Lundberg, S. M., & Lee, S., A Unified Approach to Interpreting Model Predictions. Advances in Neural Information Processing Systems, 2017.
15. American Diabetes Association (ADA), Classification and Diagnosis of Diabetes, 2022.
16. World Health Organization (WHO), Gestational Diabetes – Prevalence and Early Detection Report, 2021.
17. Streamlit Documentation, Streamlit Web Application Framework, 2023.