A Resource-Efficient Boosting Ensemble Architecture for Continuous Gestational Diabetes Risk Assessment via Wearable IoT Bio-Sensors and Mobile Edge Computing

Authors

Nelson Ifeakachukwu, Goodman Aze, Okila Mark, Somdy Ugmu

Date; June 28, 2026

Abstract

Gestational Diabetes Mellitus (GDM) presents a significant global health challenge, with its prevalence rising steadily and posing substantial risks to both maternal and fetal well-being. Current diagnostic paradigms, reliant on the Oral Glucose Tolerance Test typically administered between 24 and 28 weeks of gestation, are reactive, failing to enable early intervention. While wearable Internet of Things (IoT) bio-sensors and machine learning offer promising avenues for continuous health monitoring, existing predictive models are often computationally intensive, unsuitable for resource-constrained edge devices, and lack validated architectures for continuous GDM risk assessment. This study proposes a novel resource-efficient boosting ensemble architecture, deployed at the mobile edge, for the continuous and interpretable risk assessment of GDM using data from wearable IoT bio-sensors. The methodology integrates a Light Gradient Boosting Machine (LightGBM) as the core ensemble learner, optimized for low-latency inference and minimal memory footprint, with a stream-based data processing pipeline for real-time feature extraction from physiological signals. The framework was validated using a simulated dataset representing continuous sensor streams from a cohort of 500 pregnant individuals. The proposed architecture achieved a predictive accuracy of 89.4% in identifying high-risk GDM cases, outperforming baseline models such as AdaBoost (84.2%) and Gradient Boosting (86.1%) while significantly reducing inference time. This architecture offers a practical, scalable, and resource-aware solution for integrating continuous GDM risk assessment into routine antenatal care via mobile health platforms.

Keywords: Gestational Diabetes Mellitus, Boosting Ensemble, Wearable IoT, Mobile Edge Computing, Resource-Efficient Machine Learning

1. Introduction

1.1 Background

The global burden of diabetes is escalating at an alarming rate, with a significant proportion of cases linked to pregnancy. Gestational Diabetes Mellitus (GDM), defined as glucose intolerance first recognized during pregnancy, affects approximately one in six live births globally, a figure that continues to rise in tandem with increasing rates of obesity and maternal age . GDM is not merely a transient condition of pregnancy; it is associated with a spectrum of short-term and long-term adverse health outcomes for both the mother and the child. These include an elevated risk of pre-eclampsia, cesarean delivery, and the development of type 2 diabetes mellitus later in life for the mother . For the fetus and neonate, GDM increases the likelihood of macrosomia, shoulder dystocia, neonatal hypoglycemia, and a predisposition to obesity and metabolic disorders in childhood and adulthood.

Current clinical guidelines recommend universal screening for GDM, typically between 24 and 28 weeks of gestation, using the Oral Glucose Tolerance Test (OGTT) . The OGTT, while a well-established diagnostic tool, has several limitations. It is a point-of-care test that provides only a snapshot of glycemic status, failing to capture the dynamic nature of postprandial glucose excursions and the subtle, progressive glycemic dysregulation that may occur earlier in pregnancy. The 24–28 week timeframe is often too late for interventions that could prevent the most significant fetal programming effects . This reactive approach necessitates a paradigm shift toward proactive, continuous monitoring to enable early identification and timely intervention.

The rapid proliferation of wearable Internet of Things (IoT) bio-sensors presents an unprecedented opportunity to transform GDM risk assessment . Devices capable of continuously monitoring physiological parameters such as glucose levels, physical activity, heart rate variability, and sleep patterns can provide a high-resolution, longitudinal view of an individual's health status . However, the voluminous and continuous nature of data generated by these sensors creates significant challenges in data processing, storage, and analysis . Transmitting all raw data to centralized cloud infrastructure can introduce latency, increase bandwidth costs, and raise privacy concerns.

1.2 Problem Statement

Prior work on predicting diabetic conditions using machine learning has largely focused on centralized cloud-based processing using high-resource computational power and static datasets. For instance, the use of deep learning models such as LSTM for pregnancy risk monitoring, while achieving high accuracy, requires significant computational resources, making them unsuitable for real-time, low-latency applications on edge devices . Furthermore, the application of ensemble learning for diabetes prediction, such as the work by Bhatta using stacking and XGBoost, has shown promise but has been primarily applied to static, pre-collected clinical datasets . These approaches have not adequately addressed the architectural requirements for continuous deployment on mobile edge devices, which are constrained in terms of processing power, memory, and battery life. Recent studies combining wearable devices and machine learning models like Random Forest Regressor for early GDM biomarker prediction have yielded encouraging results but face computational bottlenecks and rely on medical background information requiring centralized processing . These models also lack explainability—a crucial requirement for clinical decision support.

A critical gap exists in the literature for a validated, resource-efficient predictive architecture that integrates the continuous data streams from wearable IoT bio-sensors, processes them locally at the mobile edge, and applies a robust ensemble learning framework to provide real-time, interpretable GDM risk assessment. The challenge lies not only in building an accurate predictive model but also in architecting a system that is sustainable, scalable, and practical for use in routine antenatal care settings, particularly in resource-constrained environments.

1.3 Objectives of the Study

General objective:
To design, implement, and validate a resource-efficient boosting ensemble architecture for continuous and interpretable GDM risk assessment, leveraging data from wearable IoT bio-sensors processed at the mobile edge.

Specific objectives:

1. To identify the most informative set of physiological and behavioral predictors from wearable sensors for early GDM risk stratification.
2. To design and optimize a LightGBM-based ensemble architecture with a minimal memory footprint and low inference latency for deployment on mobile edge devices.

3. To validate the predictive performance of the proposed architecture against benchmark ensemble models and existing literature using simulated continuous sensor data.
4. To evaluate the resource-efficiency and real-time processing capability of the proposed architecture in a simulated mobile edge environment.

1.4 Research Questions

1. What combination of continuous physiological and behavioral features, derived from wearable IoT bio-sensors, provides the most accurate early prediction of GDM risk?
2. How does the predictive performance and computational resource utilization of a LightGBM-based boosting ensemble compare to other boosting algorithms (AdaBoost, XGBoost) when deployed in a resource-constrained mobile edge environment?
3. Can a stream-based data processing pipeline at the mobile edge effectively transform raw sensor data into a feature set suitable for continuous and low-latency GDM risk classification?

1.5 Significance of the Study

This research is significant for multiple stakeholders. For clinicians and healthcare administrators, the proposed architecture offers a practical tool for moving from reactive to proactive GDM management. Continuous risk assessment can enable earlier lifestyle interventions, more targeted glucose monitoring, and potentially reduce the incidence of adverse pregnancy outcomes. For patients, it promises a non-invasive, low-burden method for monitoring their own health, increasing engagement and empowerment. For policymakers, this architecture provides a scalable and potentially cost-effective framework for improving antenatal care, especially in regions where access to traditional healthcare infrastructure is limited. From an academic perspective, this study contributes to the growing body of knowledge on real-time health informatics by providing a validated blueprint for resource-efficient machine learning at the edge, addressing a critical gap between high-accuracy models and practical healthcare deployment.

1.6 Scope and Limitations

This study focuses on the design and validation of a resource-efficient boosting ensemble architecture. The scope is defined as follows:

- **Time period:** The simulation and validation are conducted on a static dataset representing continuous sensor streams. Real-time prospective validation is beyond the scope of this initial investigation.
- **Population:** The study uses a simulated cohort representing a general pregnant population. The framework is not validated on a specific demographic or clinical subgroup.

- **Data sources:** Analysis is limited to a suite of simulated wearable sensor data, including continuous glucose monitoring and actigraphy. Real-world sensor data from multiple devices is not used due to practical constraints.
- **Limitations:** Key limitations include the use of simulated data, which may not fully capture the complexity and noise of real-world sensor data, and the absence of validation against clinical outcomes from a prospective cohort study. The architecture assumes stable internet connectivity for model updates from a central repository, though inference is performed locally.

2. Literature Review

2.1 Conceptual Review

The conceptual framework of this study is built on three core constructs: continuous health monitoring, resource-efficient machine learning, and mobile edge computing.

Gestational Diabetes Mellitus (GDM) is a condition of carbohydrate intolerance resulting in hyperglycemia of variable severity with onset or first recognition during pregnancy. The pathophysiology involves insulin resistance due to placental hormones, leading to a relative insulin deficiency. Early detection and management are crucial to mitigate risks to both mother and fetus .

Wearable IoT Bio-Sensors are non-invasive devices that continuously monitor and record physiological and behavioral data. In this study, they are conceptualized as sources for continuous glucose levels, physical activity (steps, movement intensity), and heart rate variability . The continuous stream of data from these sensors enables a temporal understanding of health, contrasting with the snapshot provided by the OGTT.

Boosting Ensembles are a powerful machine learning meta-algorithm that builds a strong classifier by combining many weak learners sequentially. Algorithms like AdaBoost, Gradient Boosting, and XGBoost have demonstrated superior performance in classification tasks . The key strength lies in their ability to reduce bias and variance, often leading to higher accuracy than single models, albeit at a potential increase in computational cost.

Resource-Efficient Machine Learning refers to the development and optimization of ML models that are computationally light, have a small memory footprint, and can perform inference with low latency and low energy consumption . This is particularly relevant for deployment on edge and mobile devices. Techniques include model pruning, quantization, and using inherently efficient algorithms like LightGBM.

Mobile Edge Computing (MEC) is a network architecture that brings computational and storage capabilities closer to the data source (the user). By processing data at the edge (e.g., on a smartphone or a local gateway), MEC reduces latency, decreases bandwidth usage, and improves data privacy . This paradigm is ideal for healthcare applications requiring real-time, privacy-preserving data analysis.

2.2 Theoretical Framework

This study is guided by two primary theoretical frameworks: the principles of predictive analytics and the practical constraints of ubiquitous computing.

Predictive Analytics in Healthcare: This framework posits that by analyzing historical and current data, one can forecast future events or risk levels. This study applies this theory to GDM by using past and present physiological data from wearables to predict a future risk score. It aligns with the principles of early intervention and preventive medicine, where identifying at-risk individuals allows for proactive and personalized care.

Resource-Constrained Computing Theory: This theoretical perspective emphasizes designing and optimizing systems to operate effectively under hardware, energy, and network limitations. It guides the architectural choices in this study, ensuring the machine learning model is not only accurate but also practical for deployment on devices with limited computing power and battery life. This framework is fundamental to realizing the MEC vision, where intelligence is pushed to the edge to overcome the limitations of cloud-centric architectures .

2.3 Empirical Review

Wearable Technology and GDM Prediction:
Recent research has explored the potential of wearable devices for early GDM prediction. For instance, a pilot study by Kolozali et al. demonstrated that combining sensory measurements from a continuous glucose monitor and a wrist-worn accelerometer with medical background information could predict GDM-associated biomarker values such as LDL, HbA1c, and OGTT results up to 16 weeks before diagnosis . While promising, the study utilized a limited cohort and relied on cloud-based processing, highlighting the need for validation in larger, more diverse populations and for more efficient architectures. Another study introduced a Heap Optimization-based Multi-LSTM for maternal-fetal health monitoring, achieving high accuracy (97.4% for fetal health, 94.5% for maternal risk), but this approach is computationally demanding, making it less ideal for continuous edge processing .

Ensemble Learning in Diabetes Prediction:
Ensemble methods have been extensively studied for static diabetes prediction. For example, Kathiria et al. developed a fused ensemble approach integrating SVM, ANN, RF, and Gradient

Boosting with a fuzzy logic mechanism, achieving exceptional accuracy (99.35%) on a symptom-based dataset . However, this high performance was achieved on a curated dataset and did not address real-time edge deployment. Bhatta compared multiple boosting and stacking ensemble models on the PIMA Indian Diabetes dataset, finding that a stacking ensemble performed best with an accuracy of 86%, outperforming XGBoost (82%) . While the study highlights the superiority of ensemble models, it was also conducted on a static dataset.

Edge Computing for Predictive Healthcare:
The application of edge computing in healthcare is an emerging field. G. Thilagavathi and N.K. Karthikeyan proposed an Optimized Ensemble Boosting Fuzzified Neural Network (OEBFNN) for early-stage diabetes detection at the edge, demonstrating improved speed and performance (95.5% accuracy) over previous systems . This research confirms the viability of edge computing for real-time health data analysis but uses a neural network approach that may still be more resource-intensive than a tree-based ensemble. Moreover, a study by Raza et al. proposed an Ensemble Extreme Learning Machine (EN-ELM) integrated with edge computing, designed to enhance healthcare decision support systems, achieving high accuracy on chronic disease datasets . This further validates the move toward edge-based ensemble models.

2.4 Research Gap

Despite these advances, a clear gap exists. Existing studies have focused on either high-accuracy but resource-intensive models (e.g., LSTMs, deep ensembles) or on static datasets for diabetes prediction . While some have explored edge computing, they often use complex neural network architectures or are not specialized for the continuous, multimodal data streams of a pregnancy monitoring scenario. A comprehensive, validated architecture that combines a resource-efficient boosting ensemble with a dedicated stream-based pipeline for continuous GDM risk assessment on mobile edge devices is currently lacking. This study aims to fill this gap by designing and evaluating a LightGBM-based architecture that balances predictive performance and computational efficiency for a practical, real-time mHealth application.

3. Methodology

3.1 Research Design

This study adopts a design-based research (DBR) and simulation methodology, integrating retrospective data analysis with prospective system architecture design. The design involved creating and simulating a continuous data stream from wearable sensors to mimic a real-world scenario. The aim was to evaluate and compare the performance and resource consumption of different algorithms within a defined resource-constrained environment. This approach allows for rigorous testing and validation of the proposed architecture without the complexity and cost of a prospective clinical trial, serving as an essential proof-of-concept.

3.2 Study Area / Population

The study targets the population of pregnant women undergoing routine antenatal care. The simulated population was designed to mirror a diverse demographic with varying baseline risk factors for GDM, including age, BMI, and family history of diabetes. This approach ensures that the architecture is evaluated under realistic and heterogeneous conditions.

3.3 Sample Size and Sampling Technique

A synthetic dataset was generated representing continuous sensor data from a cohort of **500 virtual pregnant individuals**. These virtual patients were designed with baseline characteristics (age, BMI, family history) sampled from demographic distributions to create a varied population. Each virtual patient provided a 30-day simulation period, resulting in **15,000 individual data samples**. A stratified split was used to ensure representative subgroups, with **80% of the cohort (400 patients)** used for training and **20% (100 patients)** used for testing the models.

3.4 Data Collection Methods

All data were generated synthetically due to the unavailability of a public dataset with the required high-frequency, multi-modal sensor streams from a large cohort of pregnant women. The simulation was designed to model continuous data streams from two key wearable sensors:

1. **Continuous Glucose Monitor (CGM):** Provided continuous glucose readings (in mmol/L) sampled every 5 minutes over a 24-hour period.
2. **Actigraphy Monitor (Accelerometer-based Wristband):** Provided physical activity data (step count, activity intensity) sampled per minute.

To these sensor streams, baseline clinical information (age, BMI, gestational week) and a ground-truth GDM diagnosis label were added for each virtual patient. GDM risk trajectories were modeled based on literature, where increasing insulin resistance leads to higher glucose variability and postprandial excursions.

3.5 Research Instruments

The research architecture and analysis were implemented using a suite of Python-based tools and libraries:

- **Core Language:** Python 3.9
- **Machine Learning Libraries:** Scikit-learn (v1.0) was used for baseline models and preprocessing, while LightGBM (v3.3) and XGBoost (v1.5) were used for the ensemble models. A custom implementation of an AdaBoost model was used as a baseline .
- **Data Analysis Libraries:** Pandas (v1.4) and NumPy (v1.21) were used for data manipulation and feature engineering.
- **Preprocessing:** Missing values were handled using a combination of interpolation for CGM data and MIDA-based imputation for other features to ensure data integrity .
The *slice adaptive normalization* technique, as proposed by Geetha et al., was employed to normalize the sensor data in a streaming fashion, ensuring the model is robust to variations across different times of day or patient activity levels .

3.6 Validity and Reliability

Content Validity: The features selected for inclusion (e.g., glucose variability, mean glucose, activity levels, BMI) were derived from established clinical knowledge and validated in peer-reviewed literature, ensuring that the features are representative of GDM pathophysiology .
Predictive Validity: The validity of the model was assessed using standard performance metrics against the simulated labels. The model's ability to generalize was evaluated using k-fold cross-validation.
Reliability: The entire preprocessing and modeling pipeline was implemented in a scripted, reproducible manner. To ensure reliability in a streaming context, the feature extraction functions were designed to operate on fixed-length sliding windows of sensor data, providing a consistent input format to the model.

3.7 Data Analysis Techniques

A range of boosting algorithms was benchmarked against the proposed LightGBM model:

- **AdaBoost:** A classic boosting algorithm that adjusts instance weights to focus on misclassified samples .
- **XGBoost:** A highly performant gradient boosting implementation known for its speed and regularization .
- **Gradient Boosting:** A standard gradient boosting algorithm as implemented in Scikit-learn.

- **Proposed LightGBM:** A gradient boosting framework that uses tree-based learning and its Gradient-based One-Side Sampling (GOSS) and Exclusive Feature Bundling (EFB) to make it significantly faster and more memory-efficient than other boosting methods.

The models were evaluated on the following performance metrics:

1. **Accuracy:** Percentage of correctly classified samples.
2. **Precision:** True positives / (True Positives + False Positives).
3. **Recall:** True positives / (True Positives + False Negatives).
4. **F1-Score:** Harmonic mean of precision and recall.
5. **ROC-AUC:** Area Under the Receiver Operating Characteristic Curve.
6. **Computational Resource Utilization:** Inference time per sample, model memory footprint (MB), and CPU usage during inference were measured to assess resource-efficiency.

All models were validated using **5-fold cross-validation** on the training set to ensure robust performance estimates and avoid overfitting. Hyperparameters were tuned to optimize the F1-score on the validation set.

3.8 Ethical Considerations

This study is entirely based on synthetically generated data. As such, it did not involve the collection, use, or analysis of any real patient health information (PHI). Therefore, no ethical approval from an Institutional Review Board (IRB) was required. The simulation parameters were based on published literature, ensuring that the analysis is representative of realistic scenarios while avoiding any privacy or confidentiality risks.

4. Results

4.1 Data Presentation

The synthetic dataset successfully simulated continuous sensor streams. Table 1 summarizes the descriptive statistics of key features across the entire cohort.

Table 1. Descriptive Statistics of Key Features

Feature	Mean (SD)	Min	Max	GDM Group (n=250) Mean (SD)	Non-GDM Group (n=250) Mean (SD)
Age (years)	31.2 (5.2)	20	44	33.8 (4.5)	28.6 (4.9)
BMI (kg/m²)	28.4 (6.1)	19.5	42.0	31.2 (5.2)	25.6 (5.8)
Mean Glucose (mmol/L)	5.6 (0.9)	4.1	8.5	6.4 (0.6)	4.8 (0.4)
Glucose Variability (SD)	1.2 (0.5)	0.6	2.8	1.5 (0.4)	0.9 (0.3)
Activity (steps/day)	6200 (1800)	1500	12000	5500 (1500)	6900 (1900)

As expected, individuals in the GDM group exhibited higher mean glucose, greater glucose variability, and lower physical activity levels compared to the non-GDM group.

4.2 Analysis of Results

The proposed LightGBM model demonstrated superior or comparable performance to the benchmark algorithms on the test set. The results are presented in Table 2.

Table 2. Performance Comparison of Boosting Models

Model	Accuracy (%)	Precision	Recall	F1-Score	ROC-AUC	Inference Time (ms)	Model Size (MB)
AdaBoost	84.2	0.84	0.82	0.83	0.90	4.5	5.2
Gradient Boost	86.1	0.86	0.85	0.85	0.92	8.1	12.8
XGBoost	87.5	0.87	0.86	0.86	0.93	6.4	9.7
LightGBM (Proposed)	**89.4**	**0.89**	**0.88**	**0.88**	**0.94**	**2.1**	**3.8**

The LightGBM model outperformed all baselines, achieving the highest accuracy (89.4%) and ROC-AUC (0.94). The results of the statistical significance test (paired t-test) showed a p-value of 0.03 compared to the next best model (XGBoost), indicating that the performance difference is statistically significant.

Feature Importance: The analysis of feature importance in the LightGBM model identified the most critical predictors of GDM risk. The top five most important features were:

1. **Glucose Variability (Standard Deviation of CGM readings over 2 hours)**
2. **Postprandial Glucose Spike (Maximum glucose 1-hour after a meal)**
3. **Fasting Glucose (Mean of glucose readings from 12 AM to 6 AM)**
4. **Mean Daily Steps**
5. **Age**

These findings align with clinical understanding, where glucose variability and postprandial excursions are considered key indicators of glycemic control.

5. Discussion

5.1 Interpretation

The results demonstrate that the proposed resource-efficient LightGBM architecture is a highly effective and practical solution for continuous GDM risk assessment. The model achieved an accuracy of 89.4% on a simulated test set, confirming its strong predictive capabilities. This performance is comparable to, and in some cases exceeds, the results of studies using more complex, cloud-based models . However, the key contribution lies in its resource-efficiency. The LightGBM model's inference time of 2.1 ms per sample is significantly lower than that of other algorithms like Gradient Boosting (8.1 ms). More importantly, its model size is only 3.8 MB, making it feasible for deployment on mobile devices with limited storage capacity.

This finding directly answers the research question regarding resource utilization: the proposed architecture successfully balances high predictive performance with low computational demand, making it ideal for mobile edge deployment. The model's excellent recall (0.88) and precision (0.89) suggest that it is effective at correctly identifying both high-risk patients (minimizing false negatives) and not over-alarming (minimizing false positives), a critical balance for clinical decision support. The statistically significant improvement over XGBoost validates the selection of LightGBM as the core ensemble learner, which is optimized for training speed and memory efficiency through its proprietary techniques.

5.2 Implications

Academic Implications: This research contributes to the burgeoning field of real-time health informatics by presenting a validated methodological framework for deploying complex machine learning at the edge. It provides empirical evidence that resource-efficient ensemble models can achieve state-of-the-art predictive performance for continuous health monitoring, bridging the gap between high-accuracy models and practical deployment. This study also demonstrates the efficacy of using sliding-window feature engineering to capture the temporal dynamics of sensor data for predictive health analytics.

Practical Implications: For clinicians, this architecture offers a tool to move from a reactive to a proactive model of care. By continuously monitoring at-risk patients via their wearable devices, physicians can receive early warnings of escalating GDM risk, enabling timely advice on lifestyle modifications or early glucose monitoring, potentially preventing the progression of the condition. For healthcare administrators, the scalability of an edge-based solution is highly promising. It reduces the burden on central IT infrastructure and can be integrated into existing mobile health platforms, making it a cost-effective solution for population-level antenatal care. Key metrics such as a glucose variability exceeding a dynamic threshold could be used to trigger an automated alert for the patient and their care team.

5.3 Limitations

1. **Synthetic Data:** The primary limitation is the reliance on simulated data. While modeled on real-world patterns, synthetic data may not capture the full complexity, noise, and artifacts present in real sensor data.
2. **Limited Validation:** The model was validated on a simulated cohort. Its performance on a prospective, diverse patient population in a real-world clinical setting remains unknown.
3. **Assumption of Stability:** The current architecture assumes the model's feature relationships remain stable over time. However, the physiological changes of pregnancy may necessitate periodic model retraining.

5.4 Future Research Directions

1. **Prospective Clinical Validation:** Conduct a prospective cohort study to validate the architecture's performance on real pregnant women using commercially available wearable devices.
2. **Federated Learning:** Develop a federated learning framework where models are trained across multiple edge devices without centralizing the raw patient data, enhancing privacy and personalization.
3. **Real-time Intervention Integration:** Integrate the risk assessment output with a clinical decision support system to provide automated, real-time recommendations for lifestyle interventions to the patient.
4. **Feature Expansion:** Incorporate additional sensor data streams, such as heart rate variability and sleep quality, to see if they can further enhance predictive accuracy.

6. Conclusion

This paper presented the design, implementation, and validation of a novel resource-efficient boosting ensemble architecture for continuous GDM risk assessment. By integrating a LightGBM-based ensemble with a stream-based data processing pipeline and deploying it conceptually at the mobile edge, we addressed a critical gap in real-time health monitoring. The architecture demonstrated high predictive performance, achieving an accuracy of 89.4% in identifying high-risk GDM cases, while significantly outperforming baseline models in inference speed and memory footprint. This work provides a replicable, scalable, and practical framework for future mHealth applications, demonstrating that high-quality, real-time health analytics can be performed directly on personal devices without requiring intensive cloud computing. For clinicians and patients, this architecture offers a pathway toward a more proactive and personalized approach to antenatal care, potentially transforming the management and outcomes of GDM globally.

References

1. Geetha, V., & Ettiyan, R. (2026). Heap Optimization Based Multi LSTM for enhancing maternal and fetal health monitoring through wearable devices. *Optik, 346*, 172621.
2. Kolozali, Ş., White, S. L., Norris, S., Fasli, M., & van Heerden, A. (2024). Explainable Early Prediction of Gestational Diabetes Biomarkers by Combining Medical Background and Wearable Devices: A Pilot Study With a Cohort Group in South Africa. *IEEE Journal of Biomedical and Health Informatics, 28*(4), 1860–1871.
3. Bhatta, R. P. (2026). Ensemble based machine learning model for prediction of diabetes. *NPRC Journal of Multidisciplinary Research, 3*(2), 116–129.
4. Thilagavathi, G., & Karthikeyan, N. K. (2025). Edge computing for diabetic early disease prediction based on optimized ensemble boosting Fuzzified neural network using recursive support vector feature selection. *Biomedical Signal Processing and Control, 108*, 107913.
5. Raza, M., et al. (2024). Edge computing-based ensemble learning model for health care decision systems. *Scientific Reports, 14*, 26997.
6. Kathiria, P., Chavda, N., Bhatt, M., et al. (2026). Integration of machine learning and fuzzy logic for diabetes prediction using a fused ensemble approach. *Discover Artificial Intelligence*.
7. Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794).
8. Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., ... & Liu, T. Y. (2017). LightGBM: A highly efficient gradient boosting decision tree. *Advances in Neural Information Processing Systems, 30*, 3146–3154.
9. Zhu, T., Li, K., Herrero, P., & Georgiou, P. (2023). IoMT-enabled real-time blood glucose prediction with deep learning and edge computing. *IEEE Internet of Things Journal, 10*(4), 3706–3718.
10. He, H., Bai, Y., Garcia, E. A., & Li, S. (2008). ADASYN: Adaptive synthetic sampling approach for imbalanced learning. In *2008 IEEE International Joint Conference on Neural Networks* (pp. 1322–1328).
11. Liu, F. T., Ting, K. M., & Zhou, Z. H. (2008). Isolation forest. In *2008 Eighth IEEE International Conference on Data Mining* (pp. 413–422).

12. Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software, 67*(1), 1–48.

13. Pedregosa, F., et al. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research, 12*, 2825–2830.

14. 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.

15. Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction* (2nd ed.). Springer.