

# Advancing Type II Diabetes Predictions with a Hybrid LSTM-XGBoost Approach

Ayoub Djama Waberi<sup>1</sup>, Ronald Waweru Mwangi<sup>2</sup>, Richard Maina Rimiru<sup>3</sup>

<sup>1</sup>Department of Mathematics, Pan African University Institute for Basic Sciences, Technology and Innovation (PAUSTI), Nairobi, Kenya

<sup>2</sup>Department of Computing, Jomo Kenyatta University of Agriculture and Technology, Nairobi, Kenya

<sup>3</sup>School of Computing and Information Technology (SCIT), Jomo Kenyatta University of Agriculture and Technology (JKUAT), Nairobi, Kenya

Email: ayoubdjama18@gmail.com, waweru\_mwangi@jkuat.ac.ke, rimiru@jkuat.ac.ke

**How to cite this paper:** Waberi, A.D., Mwangi, R.W. and Rimiru, R.M. (2024) Advancing Type II Diabetes Predictions with a Hybrid LSTM-XGBoost Approach. *Journal of Data Analysis and Information Processing*, 12, 163-188.  
<https://doi.org/10.4236/jdaip.2024.122010>

**Received:** February 29, 2024

**Accepted:** April 27, 2024

**Published:** April 30, 2024

Copyright © 2024 by author(s) and Scientific Research Publishing Inc. This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

## Abstract

In this paper, we explore the ability of a hybrid model integrating Long Short-Term Memory (LSTM) networks and eXtreme Gradient Boosting (XGBoost) to enhance the prediction accuracy of Type II Diabetes Mellitus, which is caused by a combination of genetic, behavioral, and environmental factors. Utilizing comprehensive datasets from the Women in Data Science (WiDS) Datathon for the years 2020 and 2021, which provide a wide range of patient information required for reliable prediction. The research employs a novel approach by combining LSTM's ability to analyze sequential data with XGBoost's strength in handling structured datasets. To prepare this data for analysis, the methodology includes preparing it and implementing the hybrid model. The LSTM model, which excels at processing sequential data, detects temporal patterns and trends in patient history, while XGBoost, known for its classification effectiveness, converts these patterns into predictive insights. Our results demonstrate that the LSTM-XGBoost model can operate effectively with a prediction accuracy achieving 0.99. This study not only shows the usefulness of the hybrid LSTM-XGBoost model in predicting diabetes but it also provides the path for future research. This progress in machine learning applications represents a significant step forward in healthcare, with the potential to alter the treatment of chronic diseases such as diabetes and lead to better patient outcomes.

## Keywords

LSTM, XGBoost, Hybrid Models, Machine Learning, Deep Learning

## 1. Introduction

Type II diabetes is a chronic disease that affects millions of individuals worldwide. The disease can cause serious damage to the body, especially nerves and blood vessels, and is often preventable. Type II Diabetes Mellitus is a serious public health concern with significant impacts on human life and health. It affects individuals' functional capacities and quality of life, leading to significant morbidity and premature mortality [1]. The sudden increase in the number of Type II Diabetes cases has raised serious public health concerns. The multifactorial nature of Type II Diabetes Mellitus poses a challenge for early detection, as symptoms can be mild and take years to manifest. Additionally, the complexity of the disease and its interactions with other factors make it difficult to predict with high accuracy using traditional methods. Current predictive models have limitations in capturing complex patterns in patient data, and there are concerns about suboptimal control of blood glucose and other targets for many patients [2].

Type II diabetes is a prevalent and serious health condition that affects a diverse range of individuals globally. It is characterized by the body's ineffective use of insulin, with around 90% of all diabetes diagnoses being type II. This chronic disease can lead to various health complications, including kidney disease, amputations, blindness, cardiovascular disease, obesity, hypertension, hypoglycemia, dyslipidemia, and an increased risk of heart attack or stroke. Notably, diabetes claims more lives annually than breast cancer and AIDS combined.

The prevalence of type II diabetes is on the rise, with more young people being diagnosed. In America alone, expenditures related to diabetes healthcare costs have significantly increased over the years. Lifestyle factors such as obesity and lack of exercise contribute to the development of type II diabetes. Genetics also plays a significant role in increasing the risk of this condition, especially for individuals with close relatives who have diabetes [3].

Moreover, people from certain ethnic backgrounds are at a higher risk of developing type II diabetes. For instance, individuals of South Asian, Chinese, African-Caribbean, and black African origin are more likely to develop this condition. Regular exercise and maintaining a healthy weight can significantly reduce the risk of developing type II diabetes by more than 50%.

Early diagnosis and treatment are crucial in managing type II diabetes effectively. Regular check-ups and blood tests are essential for early detection to prevent severe complications associated with the disease. Individuals at risk or those with pre-diabetes need to take preventative steps to avoid the progression to type II diabetes.

The importance of accurately predicting Type II Diabetes cannot be emphasized. Early detection and action can improve disease and reduce the risk of serious consequences. However, predicting Type II Diabetes is difficult due to the complexity of the components involved, which include genetic, behavioral, and environmental influences. Traditional techniques of prediction frequently rely

on a custom knowledge base using graphs, frames, first-order logic, etc., which may not always capture the correct patterns found in patient data [4].

To overcome this issue, we offer a hybrid model that incorporates Long Short-Term Memory (LSTM) networks and Extreme Gradient Boosting (XGBoost). The hybrid LSTM-XGBoost model represents an advancement over traditional methods, offering improved accuracy in predicting Type II Diabetes Mellitus and its complications, thereby contributing to early intervention and better patient outcomes.

This model tries to combine the strengths of LSTM and XGBoost to process and analyze complex medical data. We will discuss the LSTM model, this network is a sort of recurrent neural network that is noted for its capacity to process sequential data, making it perfect for dealing with time-series data, which is common in medical records. They can detect patterns over time, providing detailed insights into patient history and trends. In contrast, we will discuss the XGBoost, this model is a sophisticated implementation of gradient boosting techniques noted for its excellent efficiency, adaptability, and efficacy in classification tasks. By combining these two methods, our approach tries to capture both the temporal dynamics and complex correlations in the data, enhancing diabetes prediction accuracy.

The objectives of the study are to develop a hybrid model that leverages LSTM for temporal data analysis and XGBoost for robust classification, to validate the model's effectiveness in predicting diabetes using comprehensive datasets, and to contribute to the field of predictive healthcare by introducing a model with high accuracy, precision, recall, and F1 score. This research is significant because it advances the field of medical data analysis and predictive healthcare. Our work aims to improve prediction accuracy, allowing for earlier diagnosis and more effective therapies. This has the potential to enhance patient outcomes while also lowering the overall strain on healthcare systems. The findings of this study are likely to provide useful insights into the application of advanced machine learning techniques in healthcare [5].

## 2. Related Works

Several studies have been conducted on diabetes prediction using traditional statistical methods and machine learning algorithms. Traditional statistical methods such as logistic regression, decision trees, and k-means clustering have been used to predict diabetes with varying degrees of accuracy.

In recent years, many researchers have been using the concept of machine learning to predict Diabetes Mellitus disease. Some of the commonly used algorithms include logistic regression (LR), XGBoost (XGB), gradient boosting (GB), decision trees (DTs), ExtraTrees, random forest (RF), and light gradient boosting machines (LGBM). Each classifier has its advantages over the other classifiers.

Another recent development in machine learning is the so-called Extreme gradient boosting (XGBoost), which was introduced by [6]. XGBoost is an effi-

cient implementation of gradient boosting that is based on parallel tree learning and efficient proposal calculation and caching for tree learning. The XGBoost algorithm has found a wide variety of use cases, also in the context of energy systems research.

As the area evolved, researchers began to investigate more complicated algorithms and various datasets, recognizing the multiple nature of diabetes and its data. This shift is evident in studies such as those conducted by [7], who not only predicted diabetes but also classified its types using a variety of machine learning methods such as Random Forest, Light Gradient Boosting Machine (LGBM), Gradient Boosting Machine, SVM, Decision Tree, and XGBoost. Their approach, which included data augmentation and sampling, yielded a high accuracy rate with the LGBM Classifier. This work represents the trend of using advanced methodologies and comprehensive data processing to improve forecast accuracy and illness knowledge.

Here shows the evolution of research advancements in machine learning and healthcare for predicting diabetes. Initially, the study was to establish the viability of applying machine learning to medical predictions. Initially simple knowledge base is used to predict the disease meaning predefined rules. With time, machine learning models replace this knowledge base because they capture the semantic meanings. The simple machine learning models alone are used for classification tasks as I have discussed in previous paragraphs. Over time, the emphasis shifted to increasingly sophisticated challenges, such as distinguishing between diabetes kinds and incorporating diverse data formats, including clinical and demographic data.

[8] and others proposed a short-term traffic flow prediction model based on the CNN-XGBoost hybrid model. Although this model studies the temporal and spatial characteristics of traffic flow, the disadvantage of the CNN prediction model compared to the LSTM model is that it is difficult to perform traffic flow multi-step prediction. The grey prediction model can predict traffic flow and real-time and dynamic data.

In [9], an adaptive decomposition method is used together with an XGBoost-based regression model to forecast loads of industrial customers in China and Ireland. The authors of [10] separately forecast day-ahead loads through an LSTM neural network and XGBoost. Subsequently, an error-reciprocal method is used to combine the forecasts. However, both methods are used for a general load forecast, instead of focusing the XGBoost forecast on peak loads. Previous works like [7] have shown that XGBoost outperforms neural networks for regression and classification tasks on tabular data.

[11] proposed a Type II Diabetes Mellitus prediction model using machine learning techniques. Their dataset consisted of 1939 records with 11 biological and lifestyle parameters. Various machine learning algorithms such as Bagged Decision Trees, Random Forest, Extra Trees, AdaBoost, Stochastic Gradient Boosting, and Voting (Logistic Regression, Decision Trees, Support Vector Machine) were

employed. The greatest rate of accuracy among these classifiers was 99.14%, which was achieved by Bagged Decision Trees.

[12] implemented a machine learning system for Type I and Type II Diabetes Mellitus that employs an ensemble learning technique to track glucose levels based on independent features. They used data from 27,050 cases and 111 attributes gathered from patients at 10 different Slovenian healthcare facilities that focused on preventative medicine. For this framework, 59 variables were selected after preprocessing and feature engineering. When compared to other classifiers, LightGBM achieved better results across the board. This included better accuracy, precision, recall, AUC, AUPRC, and RMSE.

Using a variety of machine learning classifiers such as k-nearest neighbors, decision trees, AdaBoost, naive Bayes, XGBoost, and multi-layer perceptrons, 15 created a solid framework for Type II Diabetes Mellitus. They used EDA to do tasks including outlier detection, missing value completion, data standardization, feature selection, and result validation. With a sensitivity of 0.789, a specificity of 0.934, a false omission rate of 0.092, a diagnostic odds ratio of 66.234, and an AUC of 0.950, the ensembling classifiers AdaBoost and XGBoost performed the best.

### Theoretical Frameworks for Advancing Diabetes Prediction

Long Short-Term Memory (LSTM) is a type of recurrent neural network (RNN) that is particularly effective for processing sequential data, such as time series data. Recurrent Neural Networks (RNN) are a type of neural network that is particularly effective for processing sequential data, such as text, speech, and time series data. RNNs contain loops that enable them to maintain a memory of past inputs, making them suitable for tasks like language translation, speech recognition, and predicting time series data [13].

LSTMs are designed to overcome the vanishing gradient problem that occurs in traditional RNNs, which can make it difficult to learn long-term dependencies in sequential data [14]. LSTMs contain memory cells that can maintain a memory of past inputs, making them suitable for tasks like predicting time series data. The LSTM model can capture time-dependent patterns in diabetes progression and treatment response, making it a suitable model for diabetes prediction [15].

Gradient Boosting is a machine learning technique that combines multiple weak predictive models to create a strong predictive model. It is an iterative process that fits each new model to the residuals of the previous model, thereby reducing the overall error. Gradient Boosting is particularly effective in classification tasks and has been used in various applications, including diabetes prediction [16].

The Gradient Boosting framework consists of the following steps:

- 1) Start with an initial weak predictive model, such as a decision tree.
- 2) Calculate the residuals, which are the differences between the actual and predicted values.
- 3) Fit a new weak predictive model to the residuals.

4) Combine the new weak predictive model with the previous models to create an updated model.

5) Repeat steps 2 - 4 until a stopping criterion is met, such as a maximum number of iterations or a minimum reduction in error.

Gradient Boosting is effective in classification tasks because it can handle non-linear relationships and interactions between features, and it can be used with various types of weak predictive models, such as decision trees, linear regression, and neural networks [6].

The integration of Long Short-Term Memory (LSTM) with XGBoost represents a novel contribution to diabetes prediction. This integration is expected to capture time-dependent patterns in diabetes progression and treatment response while addressing the challenges posed by high-dimensional patient data. By leveraging the strengths of LSTM for temporal data analysis and XGBoost for robust classification, the hybrid model is anticipated to significantly improve the accuracy of diabetes prediction, thereby enabling more effective early intervention and patient care.

### 3. Methodology

The methodology section of this study outlines the comprehensive approach undertaken to develop and evaluate a hybrid predictive model that synergizes the capabilities of Long Short-Term Memory (LSTM) networks and eXtreme Gradient Boosting (XGBoost) for the prediction of Type II Diabetes Mellitus. This innovative model leverages the sequential data processing strength of LSTM to capture temporal dependencies and intricate patterns within patient data, alongside the robust classification and predictive power of XGBoost, to effectively identify potential diabetes cases. This section delineates the step-by-step process, from data collection and preprocessing to the final evaluation of the model's performance, establishing a clear and structured pathway toward achieving the goal of improved diabetes prediction.

#### 3.1. Data Collection and Preprocessing

The foundation of our predictive model is anchored in the meticulously curated datasets obtained from the Women in Data Science (WiDS) Datathons for the years 2020 and 2021. These datasets are integral to our research, providing a comprehensive array of patient information crucial for the accurate prediction of Type II Diabetes Mellitus.

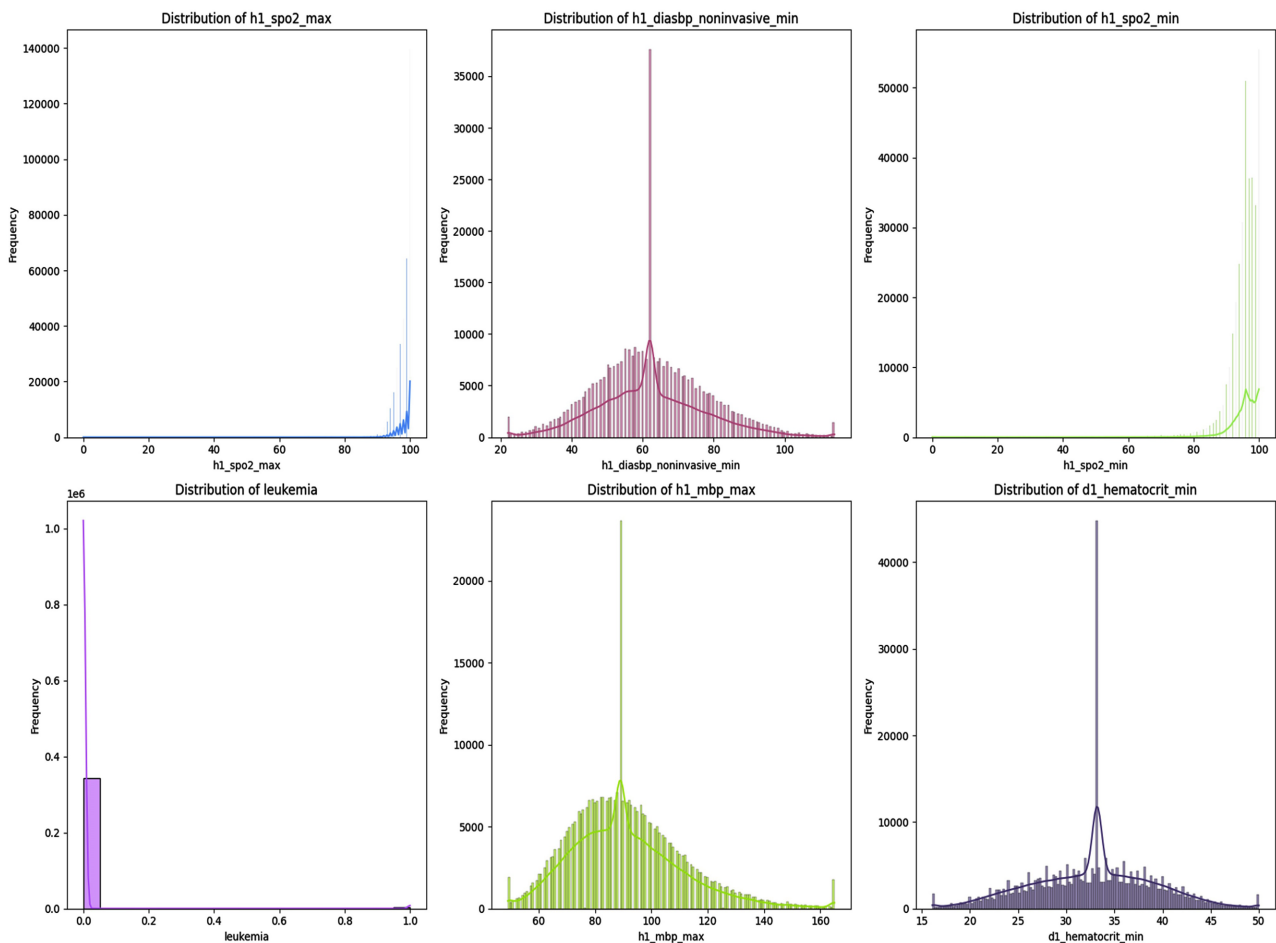
The 2020 and 2021 WiDS datasets encompass a broad spectrum of patient information, including but not limited to demographics, medical histories, and laboratory results. The 2020 dataset comprises 91,713 entries, while the 2021 dataset contains 130,157 entries, cumulatively offering a rich dataset of 221,870 patient records. This extensive collection of data points serves as a robust basis for our model, reflecting the multifaceted nature of diabetes onset and progression.

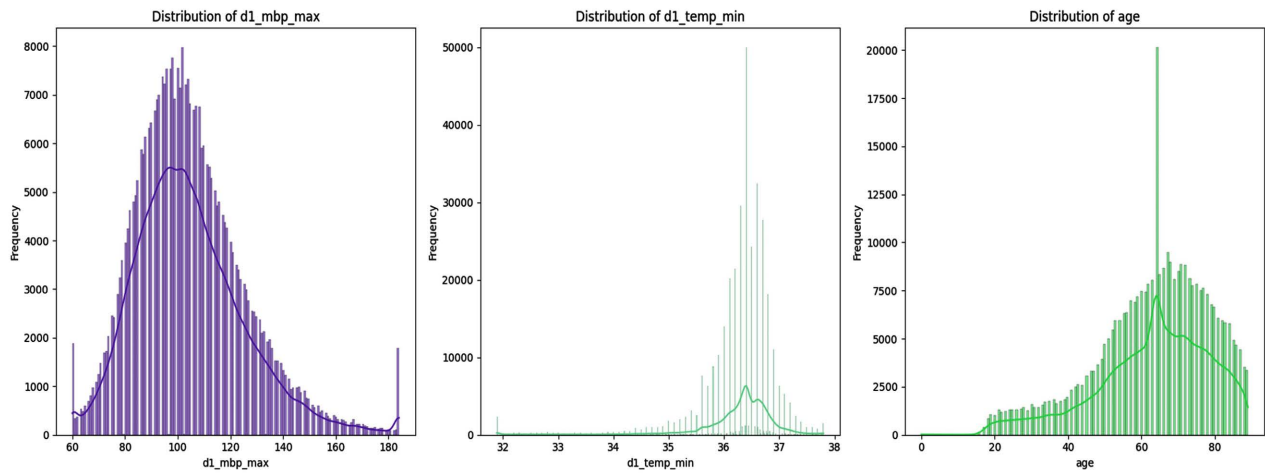
Each dataset includes critical features such as patient identifiers, hospital in-

formation, BMI, age, gender, ethnicity, blood pressure measurements, blood test results, and pre-existing health conditions, including diabetes mellitus status. To ensure the integrity and applicability of our model, we conducted a thorough preprocessing routine. This involved the elimination of columns with more than 30% missing values and identifier columns, which do not contribute to the predictive analysis. The resulting dataset was further refined to address residual missing values, with medians imputed for numerical data and modes for categorical data, ensuring a dataset devoid of null values.

Feature engineering played a pivotal role in enhancing the predictive capability of our model. This step involved the creation of new variables from existing data points, designed to uncover underlying patterns and relationships indicative of diabetes risk. Additionally, categorical variables were encoded to facilitate their integration into the machine learning models, which necessitate numerical input.

**Figure 1** illustrates the varied distributions of selected clinical features from the WiDS Diabetes Prediction Dataset. Each subplot highlights the different patterns and ranges for features such as maximum oxygen saturation ( $h1\_spo2\_max$ ), minimum noninvasive diastolic blood pressure ( $h1\_diasbp\_noninvasive\_min$ ), and patient age.





**Figure 1.** Distribution of clinical measurements for diabetes prediction.

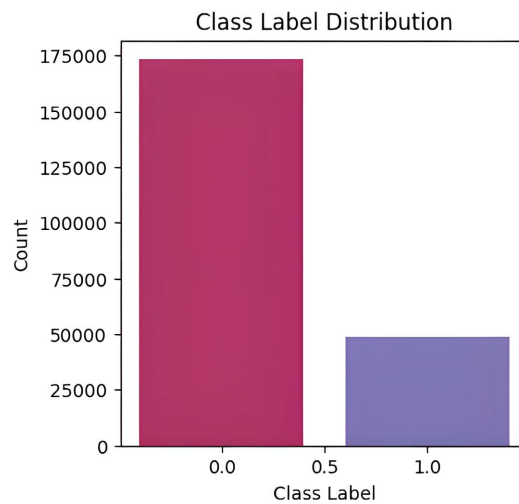
To address dataset imbalance where diabetes cases are fewer than non-diabetes cases, the study uses Random Over-Sampling. This method duplicates the diabetes cases to balance the dataset, which helps prevent model bias toward the more common non-diabetes cases.

The final stage of preprocessing involved standardizing the dataset using a Standard Scaler. This procedure adjusted the data to have a mean of zero and a standard deviation of one, a critical step to ensure uniformity in feature contribution and to foster model convergence.

**Table 1** illustrates the status before and after over-sampling:

**Table 1.** The number of instances before and after applying Random Over-Sampling to balance the dataset.

Status	Before Over-sampling	After Over-sampling
Diabetes Mellitus	48,643	173,227
Non-Diabetes	173,227	173,227

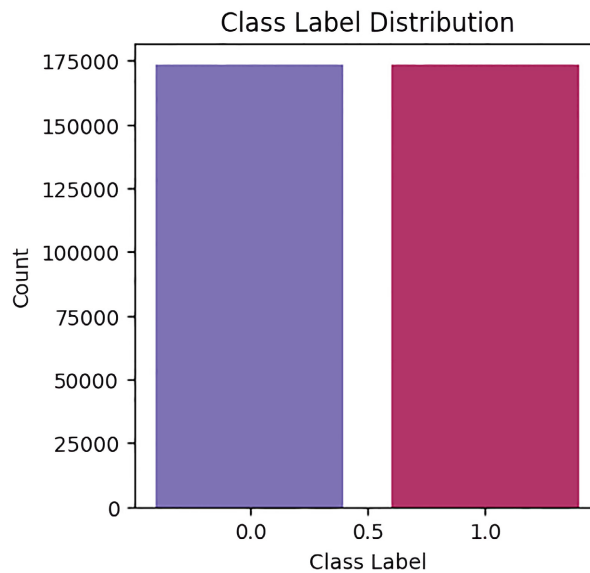


**Figure 2.** Class label distribution before over-sampling.



**Figure 2** shows the disparity between the cases with and without diabetes, indicating the necessity for over-sampling.

**Figure 3** demonstrates a balanced number of cases for both classes, achieved by Random Over-Sampling to correct the imbalance in the dataset.



**Figure 3.** Class label distribution after over-sampling.

### 3.2. LSTM Model Description

In this study, a hybrid approach based on deep learning and machine learning is proposed. This deep learning structure is based on the Recurrent Neural Network (RNN) structure. The value to be estimated in RNN structures is not only analyzed based on the current value but also based on historical data. Therefore, RNN structures are frequently used in time series data [17]. RNN structures do not delete old data, such as the work of the human brain. Classical neural network structures delete old data after using it in the weight adjustment [18]. This structure is formed by chaining the same networks. The input of each network is connected to the output of the previous RNN cells. Among the varieties of RNN structures, Long Short Term Memory (LSTM) is used in this study to create a hybrid algorithm for the detection of atrial fibrillation. The LSTM structure has begun to be widely used in estimation processes based on historical data. While RNN has a single-layer network structure, the LSTM structure has a four-layer network structure with gate mechanisms that manage the flow of information to the neural cell. The sigmoid function used in the neural network layer yields values between 0 and 1, determining the extent of the signal that is allowed to pass. This value, varying between 0 and 1, is used as a ratio.

The forget gate  $f_t$  layer decides which information to discard from the cell state. It looks at  $h_{t-1}$  and  $x_t$ , and outputs a number between 0 and 1 for each number in the cell state  $C_{t-1}$ . A 1 represents “completely keep this” while a 0 represents “completely get rid of this”.

$$f_t = \text{sigmoid}(W_f [h_{t-1}, x_t] + b_f) \tag{1}$$

The input gate decides which new values to store, using both sigmoid and tanh functions to produce the updated value and an intermediate value  $C_{ix}$ , respectively:

$$i_t = \text{sigmoid}(W_i [h_{t-1}, x_t] + b_i) \tag{2}$$

$$C_{ix} = \tanh(W_c [h_{t-1}, x_t] + b_c) \tag{3}$$

These values are combined to generate  $C_t$ , which incorporates old data with new inputs:

$$C_t = f_t \cdot C_{t-1} + i_t \cdot C_{ix} \tag{4}$$

The cell output is then calculated, using the sigmoid function to decide which data will be output from the cell, and the tanh function to scale this output:

$$o_t = \text{sigmoid}(W_o [h_{t-1}, x_t] + b_o) \tag{5}$$

$$h_t = o_t \cdot \tanh(C_t) \tag{6}$$

The Multivariate LSTM structure used in this study is similar to the classical LSTM structure but is specifically tailored for time series analysis in diabetes prediction. It captures the dynamic changes in health indicators over time, contributing to the risk of diabetes [19].

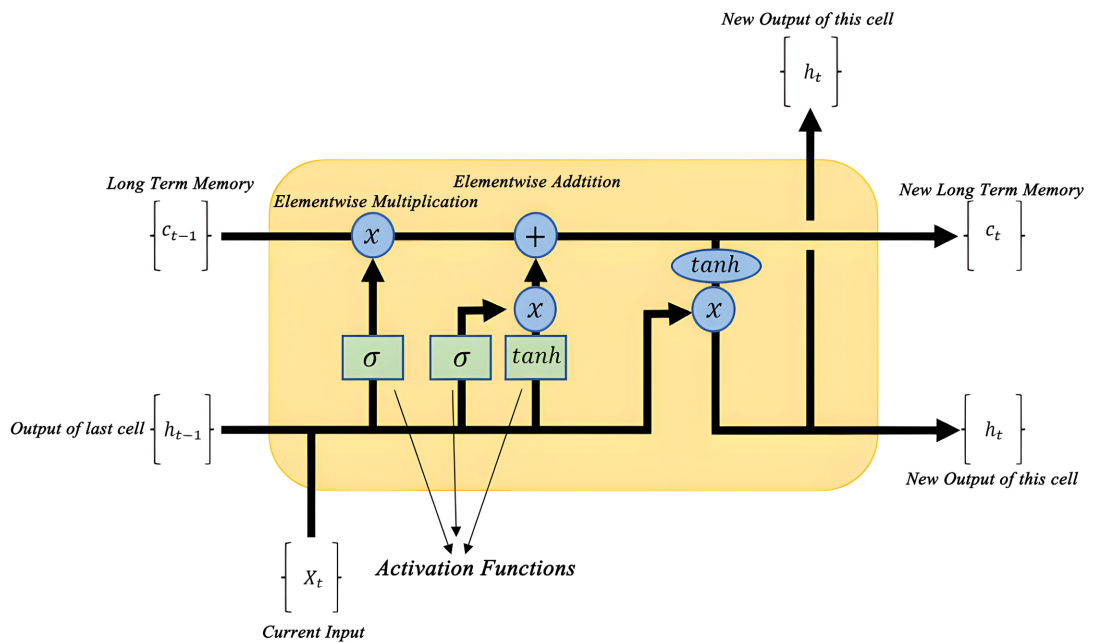


Figure 4. LSTM structure diagram.

In Figure 4, we illustrate the intricate architecture of the LSTM cell which is pivotal in the feature extraction phase of our hybrid model. This diagram depicts the flow of information through an LSTM cell, detailing the interaction between the cell state and the gates responsible for regulating the long-term and

short-term memory of the network. It is through this mechanism that the LSTM can retain relevant information over long sequences of data, a capability that is leveraged in our model to predict the progression of Type II Diabetes Mellitus effectively.

### 3.3. XGBoost Model Description

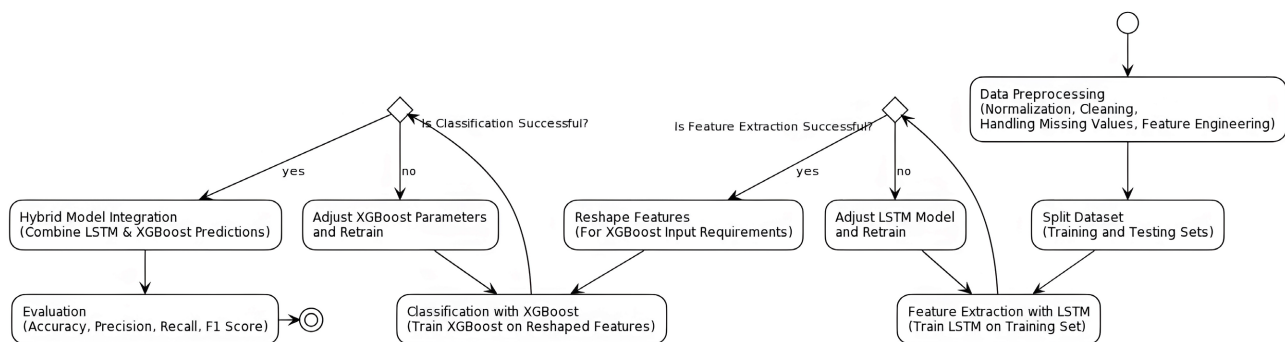
Extreme Gradient Boosting (XGBoost) is a machine learning framework that uses parallel processing to achieve high efficiency, flexibility, and portability. It is an advanced implementation of gradient-boosted decision trees designed for speed and performance. XGBoost builds upon the principles of gradient boosting by optimizing the objective function and employing regularization techniques to prevent overfitting. In our study, we utilize the XGBoost model for the classification of diabetes mellitus [20].

XGBoost operates on the principle of ensemble learning, specifically boosting, where multiple decision trees are constructed in succession to correct the errors made by prior trees [21]. The addition of the “Gradient” aspect implies the use of gradient descent to minimize the loss when adding new models.

### 3.4. Proposed Hybrid Model

We propose a hybrid LSTM-XGBoost model, aiming to combine LSTM’s ability to process sequential data and capture temporal dependencies with XGBoost’s robust classification performance. This integration seeks to address the complexities of diabetes prediction by harnessing both models’ strengths.

As illustrated in **Figure 5**, the hybrid model integration process begins with data preprocessing, followed by feature extraction using the LSTM network. The subsequent steps involve reshaping the features for compatibility with the XGBoost model, classification, and then the final integration of the two models’ predictions. This integration seeks to combine the distinct advantages of LSTM’s temporal pattern recognition and XGBoost’s classification accuracy.



**Figure 5.** Flowchart of hybrid LSTM-XGBoost model development for diabetes prediction.

#### 3.4.1. Feature Extraction with LSTM

LSTM networks are utilized for their proficiency in handling sequential data, enabling the extraction of meaningful temporal features from patient records.

This process is mathematically represented as:

$$F = \text{LSTM}(x) \quad (7)$$

where  $x$  denotes the sequential input data, and  $F$  represents the extracted features.

### 3.4.2. Reshaping LSTM Features for XGBoost

To ensure compatibility with XGBoost, LSTM-extracted features are reshaped:

$$F_{\text{reshaped}} = \text{reshape}(F, (m, -1)) \quad (8)$$

This step adapts the feature set for efficient processing by the XGBoost classifier.

### 3.4.3. Classification with XGBoost

The reshaped features are then used to train the XGBoost classifier, optimized through parameter tuning:

$$\text{Obj}(\text{XGB}) = \sum L(y_i, \hat{y}_i) + \sum \Omega(f_k) \quad (9)$$

where  $L$  denotes the loss function, and  $\Omega$  represents the regularization component.

### 3.4.4. Hybrid Model Integration

The final model integrates predictions from both LSTM and XGBoost, employing a weighted approach:

$$y_{\text{hybrid}} = \alpha \cdot \text{LSTM}(x) + (1 - \alpha) \cdot \text{XGB}(F_{\text{reshaped}}) \quad (10)$$

where  $\alpha$  is a weight parameter balancing the contributions from each model.

---

#### Algorithm Hybrid LSTM-XGBoost Model

---

**Require:** Training dataset  $D = \{(x_i, y_i)\}_{i=1}^N$ , where  $x_i$  is the input feature vector and  $y_i$  is the corresponding label.

**Ensure:** Predicted outcomes  $\hat{y}_i$  for the test dataset.

- 1: **Preprocessing:**
  - 2: Normalize and clean the dataset  $D$ .
  - 3: Handle missing values and perform feature engineering.
  - 4: Split the dataset into training and testing sets.
  - 5: **Feature Extraction with LSTM:**
  - 6: Initialize LSTM model with specified architecture.
  - 7: Train the LSTM model on the training dataset to learn temporal features.
  - 8: Extract features  $F = \text{LSTM}(x_i)$  for each instance  $x_i$  in the dataset.
  - 9: **Reshape Features for XGBoost:**
  - 10: Reshape the extracted features  $F$  to fit the input requirements of XGBoost:  
 $F_{\text{reshaped}} = \text{reshape}(F, (m, -1))$ .
  - 11: **Classification with XGBoost:**
  - 12: Initialize XGBoost model with specified parameters.
  - 13: Train the XGBoost model on the reshaped features  $F_{\text{reshaped}}$  and labels  $y_i$ .
  - 14: **Hybrid Model Integration:**
  - 15: Integrate the predictions from LSTM and XGBoost models, potentially using a weighted scheme or another integration technique.
  - 16: **Evaluation:**
  - 17: Evaluate the hybrid model on the test set using metrics such as accuracy, precision, recall, and F1 score.
  - 18: **return** The predicted outcomes  $\hat{y}_i$  for each instance in the test dataset. =0
-

The hybrid LSTM-XGBoost model merges LSTM's feature extraction from sequential data with XGBoost's classification strength, enhancing diabetes prediction by understanding temporal patterns and employing a robust classification framework. This innovative approach aims to surpass traditional models in accuracy, marking a significant advancement in analyzing complex health data.

### 3.5. Model Architecture and Training

The hybrid LSTM-XGBoost model's architecture is a critical component of our study, designed to harness the strengths of both LSTM for sequential data processing and XGBoost for robust classification. Below we detail the architecture and training process:

#### 3.5.1. LSTM Architecture

- The LSTM network is composed of several layers, each with a specific number of units: 256, 128, 64, 32, and 16 units respectively.
- The LSTM layers have additional configurations like 'return sequences' set to true or false, ensuring the sequential output is passed correctly between layers.
- Dropout layers with a rate of 0.5 are interspersed between LSTM layers to prevent overfitting.
- A dense layer with a single unit is used at the output to provide the final prediction.

#### 3.5.2. XGBoost Architecture

- The XGBoost model employs a binary: logistic objective function for binary classification.
- The learning rate is set at 0.1, with 500 estimators and a random state of 42 to ensure reproducibility.
- Key hyperparameters include a learning rate of 0.01, a max depth of 12, and 550 estimators, with a regularization term (reg alpha) of 0.001 to enhance model generalization.

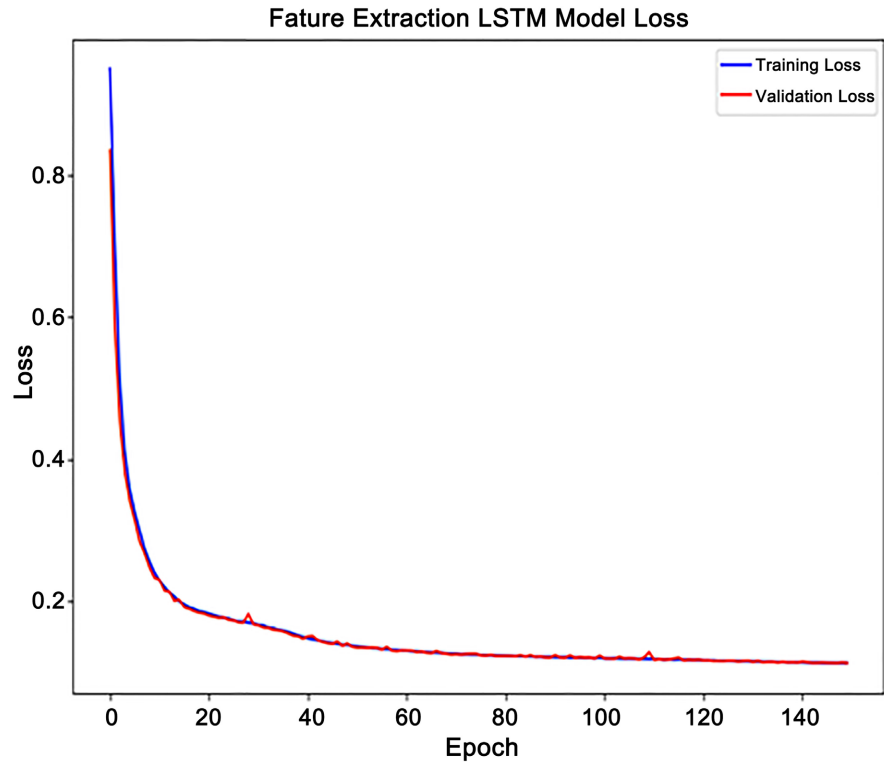
#### 3.5.3. Training Process

- The LSTM network is trained on sequential patient data, learning to capture temporal dependencies and extract meaningful features.
- The extracted features are then reshaped and fed into the XGBoost model for classification.
- Both models are integrated, leveraging LSTM's feature extraction capabilities and XGBoost's classification efficiency to predict diabetes mellitus effectively.

In **Figure 6**, the blue line depicts the training set loss and the red line delineates the loss on the validation set. This illustrates the model's learning progression and its convergence over successive epochs.

### 3.6. Evaluation Criteria

To evaluate our machine learning model's performance, we'll use key metrics



**Figure 6.** Training and validation loss of the LSTM model over epochs.

such as Accuracy, Precision, Recall, and the F1-Score. Additionally, the Confusion Matrix will provide a detailed view of the model's classification accuracy across different categories.

### 3.6.1. Accuracy

This metric evaluates the total number of instances correctly predicted by the trained model relative to all possible instances. Accuracy is defined as the proportion of images accurately classified to the total number of images provided.

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}}, \quad (11)$$

where TP refers to true positive, TN refers to true negative, FP refers to false positive, and FN refers to false negative values.

### 3.6.2. Precision

This metric measures the proportion of true positive cases among all predicted positive instances. For instance, it is mathematically represented as follows:

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}, \quad (12)$$

where TP refers to true positive and FP refers to false positive values.

### 3.6.3. Recall

This metric assesses the model's ability to correctly detect diabetes patients out of all actual cases of diabetes. Recall becomes an important measure when the

consequences of false negatives outweigh those of false positives. It is defined mathematically by the subsequent equation:

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}, \quad (13)$$

where TP refers to true positives and FN refers to false negative values.

### 3.6.4. F1-Score

The F1 score offers a combined metric of classification accuracy, taking into account both precision and recall. It is the harmonic mean of the two, providing a balance between them. The F1 score reaches its maximum value when precision and recall are equal. This measure effectively gauges the model's comprehensive performance by integrating the results of both precision and recall.

$$\text{F1 Score} = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (14)$$

### 3.6.5. Confusion Matrix (CM)

A confusion matrix presents algorithm performance in a tabular format. It offers a visual representation of key predictive metrics like recall, specificity, accuracy, and precision. This matrix is a table used to describe the performance of a classification model. It provides insight into the types of errors made by the model, showing the number of True Positives, False Positives, True Negatives, and False Negatives. These metrics provide a full assessment of the model's performance. Together, these criteria allow for a full evaluation of the model's ability to accurately forecast diabetes, ensuring its dependability and effectiveness in real-world applications.

## 4. Experiment and Results

In the "Experiment Results" section, we scrutinize the efficacy of the LSTM, XGBoost, and hybrid LSTM-XGBoost models. Our comprehensive analysis delineates the performance of each model across a range of metrics, including accuracy, precision, recall, and the F1 score. A comparative examination is also presented, elucidating their respective performances in a side-by-side assessment.

### 4.1. LSTM Model Performance

The LSTM model, designed to capture temporal dependencies within the data, exhibited a training accuracy of 0.8220. Its testing accuracy was slightly superior at 0.83, which is noteworthy considering the complexity of the sequential data being processed. The precision of the model stood at 0.80, while recall was remarkably high at 0.89, suggesting the model's proficiency in identifying true positive cases. The F1 score, a critical measure in medical diagnostics, was 0.84, reflecting a robust balance between precision and recall.


#### Architecture:

**Table 2** illustrates the configuration details for both the LSTM and XGBoost

components within our hybrid model. The LSTM part encompasses a sequence of layers with “relu” activations, tuned to capture the temporal dynamics of the data. The “return sequences” parameter is carefully adjusted to ensure the output feeds appropriately into subsequent layers. For the XGBoost classifier, a precise selection of hyperparameters balances the model’s learning complexity with performance, incorporating a binary: logistic objective and regularization to optimize classification tasks.

**Table 2.** Model configuration for the LSTM and XGBoost components.

Component	Layer Type	Configuration/Size	Additional Parameters
<b>LSTM Part</b>	LSTM Layer	128 units, activation = “relu”	return_sequences = True
	LSTM Layer	64 units, activation = “relu”	return_sequences = True
	LSTM Layer	32 units, activation = “relu”	return_sequences = False
	RepeatVector	-	-
	LSTM Layer	32 units, activation = “relu”	return_sequences = True
	LSTM Layer	64 units, activation = “relu”	return_sequences = True
	TimeDistributed	Dense layer	size equal to X_train_scaled.shape[2]
<b>XGBoost Part</b>	XGBoost Classifier	binary:logistic objective	colsample_bytree = 0.3
		learning_rate = 0.1	max_depth = 12
		n_estimators = 500	reg_alpha = 0.001
		random_state = 42	-



**Figure 7.** LSTM confusion matrix.

**Confusion Matrix**

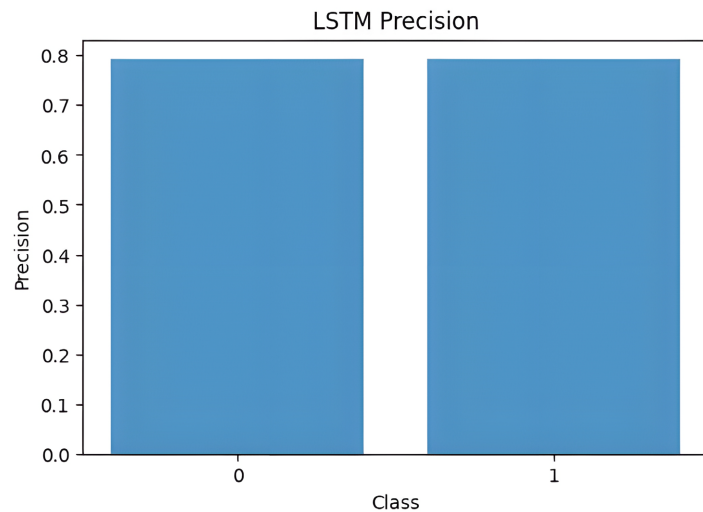
**Figure 7** illustrates the LSTM model’s classification performance, with the confusion matrix providing a clear visual representation. Darker shades indicate higher numbers of correctly predicted cases, delineating the model’s true positive



and true negative rates. This visualization is key in evaluating the model's ability to distinguish between diabetic and non-diabetic instances accurately.

### Precision

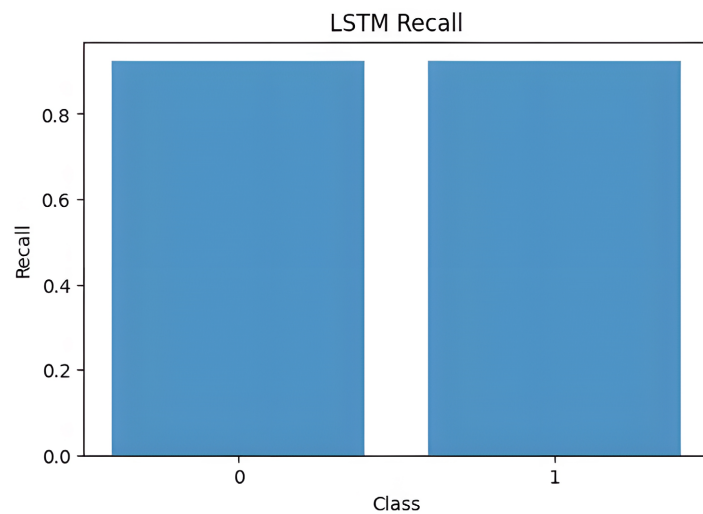
**Figure 8** reflects the model's precision, indicating the proportion of true positive predictions out of all positive predictions. High precision relates to a low false positive rate, crucial for medical diagnostic tools.



**Figure 8.** LSTM precision.

### Recall

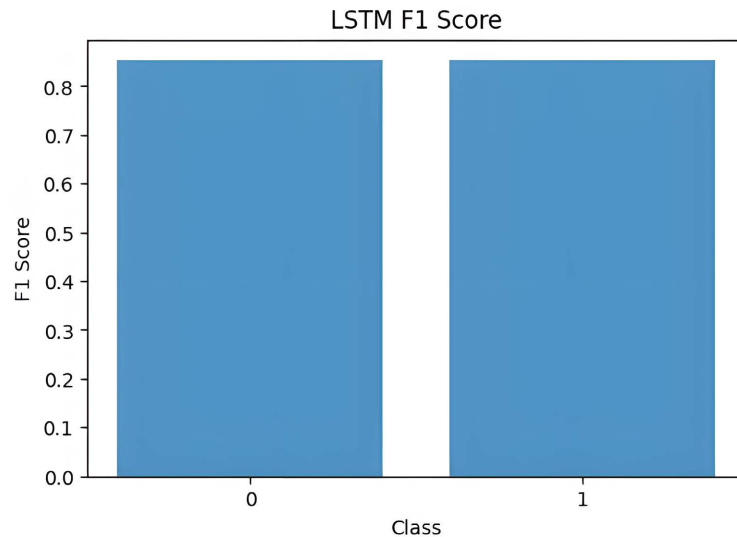
**Figure 9** shows the model's recall, reflecting its capability to identify all actual positives accurately. High recall indicates minimal false negatives, a vital factor in medical diagnosis, where overlooking a true condition could have significant consequences.



**Figure 9.** LSTM recall.

### F1 Score

**Figure 10** presents the F1 score, amalgamating precision and recall into a solitary measure that offers an equitable perspective on the LSTM model's classification efficacy. A high F1 score suggests a balanced classification capability.



**Figure 10.** LSTM F1 score.

## 4.2. XGBoost Model's Performance

The XGBoost model exhibited exemplary performance. It achieved a remarkable training accuracy of 0.98, indicative of its proficiency in learning from the training data. The test accuracy stood at 0.93, affirming the model's generalization capabilities. Precision was high at 0.92, reflecting the model's ability to identify positive cases correctly. At the same time, the recall was even more impressive at 0.95, suggesting that it successfully recognized the vast majority of true positive instances. The F1 score, balancing precision and recall, was an excellent 0.93, signifying a well-rounded predictive model.

**Table 3.** Architectural parameters of the XGBoost model with detailed descriptions, highlighting the model's complexity and regularization strategies to ensure effective learning without overfitting.

Parameter	Value	Description
Objective	binary: logistic	The objective function for binary classification.
colsample_bytree	0.3	Fraction of features used per tree, combating overfitting.
Learning Rate	0.01	Shrinks feature weights to improve model robustness.
Max Depth	12	Limits tree depth to prevent over-complex models.
reg_alpha	0.001	L1 regularization on weights, encouraging sparsity.
n_estimators	550	Total count of boosting trees to be constructed.
Random State	42	Ensures reproducibility with a fixed seed.


### Architecture:

**Table 3** delineates the architectural parameters of the XGBoost model, detailing

the specific values and their functions. It sheds light on the model's complexity and the implemented regularization strategies, such as feature fraction selection and weight penalization, which are pivotal in fostering effective learning and averting overfitting.

### Confusion Matrix

**Figure 11** illustrates the model's proficiency in classifying true positives and true negatives, which are pivotal for appraising the performance of a binary classifier.



**Figure 11.** XGBoost confusion matrix.

**Table 4** shows the precision, recall, and F1 score for the XGBoost model, showcasing its reliable performance across both classes. The scores indicate the model's balanced accuracy in classifying both the negative and positive instances, essential for medical diagnostics.

**Table 4.** XGBoost metrics.

Metric	Class 0	Class 1
Precision	0.90	0.91
Recall	0.92	0.93
F1 Score	0.91	0.92

### 4.3. Hybrid LSTM-XGBoost: Model Performance

The hybrid model, employing LSTM for feature extraction and XGBoost for classification, exhibited stellar performance. It achieved an impeccable training accuracy of 0.99, demonstrating flawless learning and fitting to the training data. The model also posted a commendable test accuracy of 0.98, signifying its outstanding generalization capabilities. With a precision of 0.98 and a recall of 0.99, the model showed exceptional proficiency in identifying positive cases while mi-

nimizing false negatives. The near-perfect F1 score of 0.98 underscores an optimal balance between precision and recall.

### Architecture

**Table 5** details the hybrid model's intricate architecture, showcasing a multi-layered LSTM configuration replete with regularization and dropout strategies to refine feature learning and mitigate overfitting, ultimately converging to a singular dense layer for binary classification output.

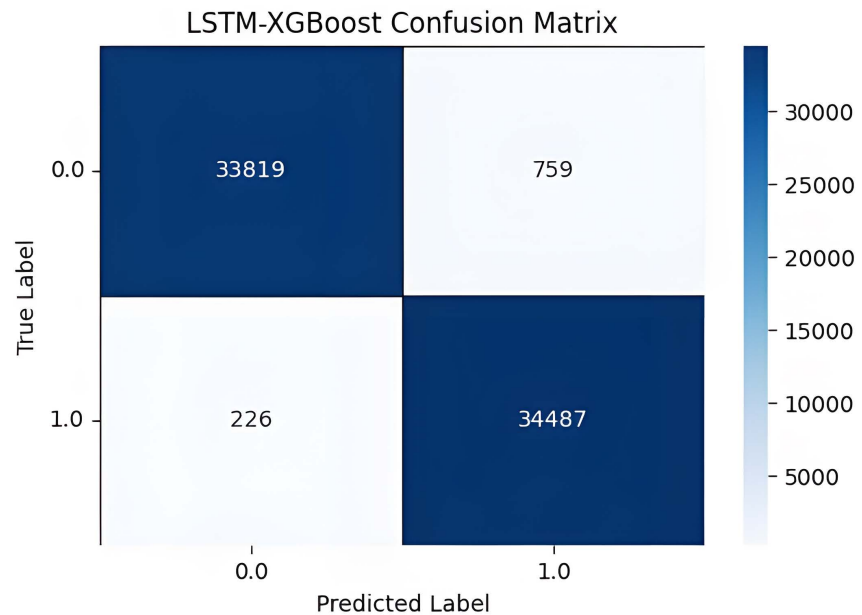
**Table 5.** Hybrid architecture.

Layer Type	Size/Configuration	Parameters
LSTM	256 units	activation = "relu", return sequences = True, kernel and recurrent regularizer = l1_l2 (1e-5, 1e-4)
Batch Normalization	-	Normalizes the activations from the previous layer
Dropout	0.5	Randomly sets input units to 0 at each step to prevent overfitting
LSTM	128 units	activation = "relu", return sequences = True, kernel and recurrent regularizer = l1_l2 (1e-5, 1e-4)
Batch Normalization	-	-
Dropout	0.5	-
LSTM	64 units	activation = "relu", return sequences = True, kernel and recurrent regularizer = l1_l2 (1e-5, 1e-4)
Batch Normalization	-	-
Dropout	0.5	-
LSTM	32 units	activation = "relu", kernel and recurrent regularizer = l1_l2 (1e-5, 1e-4)
Batch Normalization	-	-
Dropout	0.5	-
Dense	16 units	activation = "relu", kernel regularizer = l1_l2 (1e-5, 1e-4)
Batch Normalization	-	-
Dropout	0.5	-
Dense	1 unit	activation = "sigmoid", kernel regularizer = l1_l2 (1e-5, 1e-4)

### Confusion Matrix

**Figure 12** shows the hybrid model's true positive and true negative rates, with the top left and bottom right cells displaying the counts of accurately predicted negative (0) and positive (1) classes, respectively. The off-diagonal cells denote the instances of misclassification.

**Table 6** presents a concise summary of the LSTM-XGBoost model's performance, detailing the precision, recall, and F1 score metrics for both classes. Precision values demonstrate the model's accuracy in predicting positive cases, while recall figures reflect its effectiveness in identifying all positive samples. The F1 scores indicate a well-balanced harmony between precision and recall for both classes.



**Figure 12.** LSTM-XGBoost confusion matrix.

**Table 6.** Summary of the LSTM-XGBoost model's performance.

Metric	Class 0	Class 1	Comments
Precision	0.99	0.98	Demonstrate the model's exactness in positive predictions.
Recall	0.98	0.99	Measure the model's success at capturing all positive samples.
F1 Score	0.98	0.98	Shows a harmonious balance of precision and recall.

#### 4.4. Comparative Analysis

The comparative examination of these models demonstrates the advantages of each strategy. The XGBoost model, noted for its resilience and efficiency, has excellent balance across all criteria. The LSTM model, while slightly lacking in accuracy and precision, excels in recall, making it useful in situations where missing a positive case could be crucial. However, the hybrid model stands out in every aspect, combining the benefits of both LSTM and XGBoost to attain near-perfect scores across all metrics. This demonstrates the efficacy of integrating LSTM's feature extraction capabilities with the predictive power of XGBoost.

**Table 7** compares the XGBoost, LSTM, and a hybrid model that combines the two models for diabetes prediction. The XGBoost model has a high level of overall efficacy, with a training accuracy of 0.98 and a test accuracy of 0.93, demonstrating that it can learn from training data and apply that knowledge to new data. Its precision of 0.92 and recall of 0.95 demonstrate its capacity to reliably and fully identify positive cases of diabetes. The F1 Score of 0.93 demonstrates a balanced approach, taking into account both precision and recall criteria. Meanwhile, the LSTM model, albeit somewhat lower in training (0.8220) and test accuracy (0.83), excels in recall (0.89), demonstrating its ability to detect the majority of genuine positive diabetes patients. However, its accuracy score of 0.80 indicates

room for growth in reliably diagnosing non-diabetic cases, while its F1 Score of 0.84 indicates a decent but not ideal combination of precision and recall.

**Table 7.** Comparative performance of LSTM, XGBoost, and hybrid models.

Model Type	Metric	Train Accuracy	Test Accuracy	Precision	Recall	F1 Score
LSTM	-	0.8220	0.83	0.80	0.89	0.84
XGBoost	-	0.98	0.93	0.92	0.95	0.93
Hybrid (LSTM-XGBoost)	-	0.99	0.98	0.97	0.99	0.98

In comparison, the hybrid model, which includes the properties of both LSTM and XGBoost, outperforms the separate models by scoring near-perfect on all criteria. It achieves an impressive training accuracy of 0.99 and a test accuracy of 0.98, demonstrating great learning and generalization abilities. The model achieves a high precision score of 0.97 and a flawless recall score of 0.99, demonstrating its outstanding ability to reliably identify all positive diabetes cases with no false negatives. The hybrid model has a considerably higher F1 Score (0.98) than the standalone LSTM and XGBoost models, indicating a better balance of precision and recall. The hybrid model's comprehensive and high-performing nature demonstrates the usefulness of combining LSTM's sequential data processing capacity with XGBoost's powerful classification, resulting in the most robust and dependable model for predictive tasks in this study.

## 5. Discussion

Our study's findings suggest that combining LSTM (Long Short-Term Memory) and XGBoost models, known as a hybrid model, is good at predicting diabetes. This hybrid model has demonstrated good levels of accuracy, precision, recall, and F1 scores, all of which indicate how well the model predicts diabetes. The rationale for this success is that LSTM excels at interpreting and processing patient data over time, whereas XGBoost excels at categorizing it (such as "has diabetes" or "does not have diabetes"). They work better together than they would individually. The LSTM detects crucial trends and patterns in the patient's health data over time, and XGBoost uses these discoveries to reliably forecast whether a patient has diabetes.

### 5.1. Advantages of the Hybrid Approach

The main advantage of employing this hybrid strategy is that it combines the greatest features of two modern machine learning algorithms. LSTM excels at working with data that change over time, such as a patient's health records, whereas XGBoost is extremely efficient and accurate at classifying data, which is critical for determining whether a patient has a condition like diabetes. The model leverages LSTM to effectively capture and analyze time-dependent features in the data, which are crucial for predicting the progression of Type II Diabetes Mellitus.

By incorporating XGBoost, the model benefits from a powerful classification algorithm that improves prediction accuracy, especially on large and complex datasets. This combination improves the model's ability to analyze complicated health data while also understanding the finer specifics of each patient's circumstance. This could lead to more personalized healthcare, as the model can recommend therapies based on individuals' distinct health patterns and demands.

## 5.2. Limitations and Ethical Considerations

This research acknowledges the inherent limitations associated with the hybrid LSTM-XGBoost model in predicting Type II Diabetes Mellitus. While the model demonstrates promising results, its reliability and interpretability in healthcare settings are crucial areas for further scrutiny. The opacity of machine learning models, especially in complex healthcare scenarios, necessitates ongoing efforts to enhance model transparency and understandability.

Furthermore, ethical implications, including patient data privacy and the potential consequences of model decisions, are paramount. It is essential to continually evaluate the model against these factors to ensure its ethical deployment in real-world healthcare environments. This study underscores the need for a multidisciplinary approach, incorporating insights from healthcare professionals, data scientists, and ethicists, to advance the field of predictive healthcare responsibly.

## 5.3. Future Work

Our study's technique might be improved by using more forms of data, such as photographs and written patient records. For example, adding images from medical scans (such as retinal scans) could aid in the early detection of diabetes problems. Including this type of information could provide us with a more accurate and full picture of a patient's health. Similarly, using modern language processing tools to evaluate what patients write about their symptoms and sentiments could help us better comprehend their diseases. This could help doctors diagnose diabetes more correctly and recommend treatments that are better suited to each patient's individual needs.

The idea of incorporating this type of data into our model is a promising step forward in healthcare. This means that we may utilize machine learning not only to crunch data but also to comprehend the nuances of human language and visual clues. This combination of technology and healthcare may lead to new methods of predicting, diagnosing, and treating diseases such as diabetes. It's a move towards healthcare that's more in tune with each patient's individual needs, potentially transforming the way we approach medical care.

## 6. Conclusion

### 6.1. Summary of Key Findings

Our research finds some significant findings to predict diabetes using deep

learning and machine learning techniques. The key achievement was the creation and validation of a hybrid LSTM-XGBoost model, which outperformed standalone LSTM and XGBoost models. This model correctly predicted diabetes by efficiently processing patient data, particularly identifying temporal trends with LSTM and robust classification with XGBoost. The strong accuracy, precision, recall, and F1 scores suggest that this model has the potential to be a trustworthy diabetes prediction tool in healthcare.

The hybrid approach's effectiveness stems from its ability to combine the benefits of LSTM's sequential data processing with XGBoost's excellent categorization capabilities. This synergy has proven especially useful when working with complicated datasets common in healthcare, where variables are numerous and interdependent.

## 6.2. Future Research Directions

Looking ahead, there are several promising areas for future research. One significant aim is to include multimodal data sources, such as medical imaging and textual patient records, in prediction models. This technique has the potential to improve the model's diagnostic capabilities by detecting nuanced indicators of diabetes-related problems that would otherwise go undetected in normal clinical data.

## 6.3. Concluding Remarks

Finally, our findings represent a substantial advancement in the use of machine learning in healthcare, notably in the field of diabetes prediction. The success of the hybrid LSTM-XGBoost model opens up new avenues for early and accurate diagnosis, which is critical for effective diabetes management and treatment. This technique has the potential to go beyond diabetes prediction, with implications for healthcare diagnostics and tailored medicine. As we continue to research and improve these technologies, we get closer to a future in which healthcare is more predictive, personalized, and accessible.

## 7. Experimental Setup

Our research utilized Jupyter Notebooks via Anaconda and Google Colab's cloud-based platform to develop and evaluate the hybrid LSTM-XGBoost model for diabetes prediction.

### 7.1. Software and Tools

- Development was done in Python 3.x within Anaconda's Jupyter Notebooks, utilizing TensorFlow for LSTM implementation, XGBoost for classification, and pandas and NumPy for data handling.

### 7.2. Computational Resources

- The project leveraged Google Colab for its GPU acceleration and up to 16GB



of RAM, providing a robust and accessible environment for model training and testing.

### 7.3. Cloud Computing Advantages

Google Colab's cloud-based platform was instrumental in:

- Facilitating scalable and flexible computational resources.
- Enabling seamless collaboration and accessibility to the project from various locations.
- Offering a cost-effective approach by providing free access to high-performance computing resources.

This setup highlights our approach to integrating cutting-edge computational resources and data science tools to advance diabetes prediction methodologies.

### Conflicts of Interest

The authors declare no conflicts of interest regarding the publication of this paper.

### References

- [1] Sevilla-Gonzalez, M.D.R., Bourguet-Ramirez, B., Lazaro-Carrera, L.S., Martagon-Rosado, A.J., Gomez-Velasco, D.V. and Viveros-Ruiz, T.L. (2022) Evaluation of a Web Platform to Record Lifestyle Habits in Subjects at Risk of Developing Type 2 Diabetes in a Middle-Income Population: Prospective Interventional Study. *JMIR Diabetes*, **7**, e25105. <https://doi.org/10.2196/25105>
- [2] Alam, T.M., Iqbal, M.A., Ali, Y., Wahab, A., Ijaz, S., Baig, T.I., Hussain, A., Malik, M.A., Raza, M.M., Ibrar, S., et al. (2019) A Model for Early Prediction of Diabetes. *Informatics in Medicine Unlocked*, **16**, Article ID: 100204. <https://doi.org/10.1016/j.imu.2019.100204>
- [3] Bhat, S.S., Selvam, V., Ansari, G.A., Ansari, M.D., Rahman, M.H., et al. (2022) Prevalence and Early Prediction of Diabetes Using Machine Learning in North Kashmir: A Case Study of District Bandipora. *Computational Intelligence and Neuroscience*, **2022**, Article ID: 2789760. <https://doi.org/10.1155/2022/2789760>
- [4] American Diabetes Association (2010) Diagnosis and Classification of Diabetes Mellitus. *Diabetes Care*, **33**, S62-S69. <https://doi.org/10.2337/dc10-S062>
- [5] Bhat, S.S. and Ansari, G.A. (2021) Predictions of Diabetes and Diet Recommendation System for Diabetic Patients Using Machine Learning Techniques. 2021 *2nd International Conference for Emerging Technology (INCET)*, Belagavi, 21-23 May 2021, 1-5.
- [6] Chen, T.Q. and Guestrin, C. (2016) Xgboost: A Scalable Tree Boosting System. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, San Francisco, 13-17 August 2016, 785-794. <https://doi.org/10.1145/2939672.2939785>
- [7] Ahamed, B.S., Arya, M.S. and Nancy, A.O. (2022) Diabetes Mellitus Disease Prediction Using Machine Learning Classifiers and Techniques Using the Concept of Data Augmentation and Sampling. In: Tuba, M., Akashe, S. and Joshi, A., Eds., *ICT Systems and Sustainability: Proceedings of ICT4SD 2022*, Springer, Berlin, 401-413. [https://doi.org/10.1007/978-981-19-5221-0\\_40](https://doi.org/10.1007/978-981-19-5221-0_40)
- [8] Zhang, X.J. and Zhang, Q.R. (2020) Short-Term Traffic Flow Prediction Based on

- LSTM-XGBoost Combination Model. *CMES-Computer Modeling in Engineering & Sciences*, **125**, 95-109. <https://doi.org/10.32604/cmes.2020.011013>
- [9] Zhu, X., Chu, J., Wang, K.D., Wu, S.F., Yan, W. and Chiam, K. (2021) Prediction of Rockhead Using a Hybrid N-XGboost Machine Learning Framework. *Journal of Rock Mechanics and Geotechnical Engineering*, **13**, 1231-1245. <https://doi.org/10.1016/j.jrmge.2021.06.012>
- [10] Bai, L. and Pinson, P. (2019) Distributed Reconciliation in Day-Ahead Wind Power Forecasting. *Energies*, **12**, Article No. 1112. <https://doi.org/10.3390/en12061112>
- [11] Ganie, S.M. and Malik, M.B. (2022) An Ensemble Machine Learning Approach for Predicting Type-II Diabetes Mellitus Based on Lifestyle Indicators. *Healthcare Analytics*, **2**, Article ID: 100092. <https://doi.org/10.1016/j.health.2022.100092>
- [12] Kopitar, L., Kocbek, P., Cilar, L., Sheikh, A. and Stiglic, G. (2022) Early Detection of Type 2 Diabetes Mellitus Using Machine Learning-Based Prediction Models. *Scientific Reports*, **10**, Article No. 11981. <https://doi.org/10.1038/s41598-020-68771-z>
- [13] Balci, F. (2022) A Hybrid Attention-Based LSTM-XGboost Model for Detection of ECG-Based Atrial Fibrillation. *Gazi University Journal of Science Part A: Engineering and Innovation*, **9**, 199-210. <https://doi.org/10.54287/gujisa.1128006>
- [14] Miao, Y.J., Gowayed, M. and Metze, F. (2015) Eesen: End-to-End Speech Recognition Using Deep RNN Models and WFST-Based Decoding. 2015 *IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, Scottsdale, 13-17 December 2015, 167-174. <https://doi.org/10.1109/ASRU.2015.7404790>
- [15] Sak, H., Senior, A.W. and Beaufays, F. (2014) Long Short-Term Memory Recurrent Neural Network Architectures for Large Scale Acoustic Modeling. *Proceedings Interspeech 2014*, Singapore, 14-18 September 2014, 338-342. <https://doi.org/10.21437/Interspeech.2014-80>
- [16] Chen, T.Q., He, T., Benesty, M., Khotilovich, V., Tang, Y., Cho, H., Chen, K., Mitchell, R., Cano, I., Zhou, T.Y., et al. (2015) Xgboost: Extreme Gradient Boosting. *R Package Version 0.4-2*, 1, 1-4.
- [17] Deng, L., Yu, D., et al. (2014) Deep Learning: Methods and Applications. *Foundations and Trends® in Signal Processing*, **7**, 197-387. <https://doi.org/10.1561/20000000039>
- [18] Ciregan, D., Meier, U. and Schmidhuber, J. (2012) Multi-Column Deep Neural Networks for Image Classification. 2012 *IEEE Conference on Computer Vision and Pattern Recognition*, Providence, 16-21 June 2012, 3642-3649. <https://doi.org/10.1109/CVPR.2012.6248110>
- [19] Shwartz-Ziv, R. and Armon, A. (2022) Tabular Data: Deep Learning Is Not All You Need. *Information Fusion*, **81**, 84-90. <https://doi.org/10.1016/j.inffus.2021.11.011>
- [20] Jin, Y.R., Qin, C.J., Huang, Y.X., Zhao, W.Y. and Liu, C.L. (2020) Multi-Domain Modeling of Atrial Fibrillation Detection with Twin Attentional Convolutional Long Short-Term Memory Neural Networks. *Knowledge-Based Systems*, **193**, Article ID: 105460. <https://doi.org/10.1016/j.knsys.2019.105460>
- [21] Mitchell, R. and Frank, E. (2017) Accelerating the XGboost Algorithm Using GPU Computing. *PeerJ Computer Science*, **3**, e127. <https://doi.org/10.7717/peerj-cs.127>