ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25apr596

Human-Centric Approach to Diabetes Prediction Using Machine Learning Models

Mohammad Fahad[1]; Majid Hussain[2]; Mohd Arif Khan[3]; Ehteshaam Hussain[4]

[1]Department of Computer Science and Engineering Integral University
[2]Department of Computer Science and Engineering Integral University
[3]Department of Computer Science and Engineering Integral University
[4]Department of Computer Science and Engineering Integral University

Publication Date: 2025/04/26

Abstract: Diabetes mellitus is one of the most pressing global health issues, affecting millions worldwide. Early prediction and timely management can significantly reduce the disease's impact and improve the quality of life for individuals at risk. This research presents a detailed and human- centric approach to building a diabetes prediction model using machine learning algorithms. By leveraging real-world patient data, we explore various supervised learning techniques, assess their accuracy, and highlight the importance of interpretability in predictive healthcare. This paper emphasizes the ethical implications, real-world applications, and the need to bridge the gap between technology and patient-centered care.

How to Cite: Mohammad Fahad; Majid Hussain; Mohd Arif Khan; Ehteshaam Hussain (2025). Human-Centric Approach to Diabetes Prediction Using Machine Learning Models. *International Journal of Innovative Science and Research Technology*, 10(4), 1410-1415. https://doi.org/10.38124/ijisrt/25apr596

I. INTRODUCTION

In today's world, where technology continues to transform every facet of our lives, healthcare stands at the forefront of this digital revolution. One of the diseases that demands constant vigilance is diabetes—a chronic condition that, if left unchecked, can lead to severe complications including heart disease, kidney failure, and even death.

According to the World Health Organization (WHO), over 422 million people worldwide have diabetes, with most living in low- and middle-income countries. The silent nature of this disease makes early diagnosis a crucial step in combating its effects. Fortunately, with the advancement of machine learning and data analytics, we now have powerful tools that can aid in predicting the onset of diabetes, especially Type 2 diabetes, which is largely preventable.

The goal of this research is not just to build a model that predicts diabetes, but to do so in a way that is understandable, ethical, and aligned with the need of real people—patient, doctor, and healthcare providers alike.

II. LITERATURE REVIEW

Many researchers have explored diabetes prediction using different methodologies. Early studies focused on statistical models such as logistic regression, which provided reasonable accuracy but lacked sophistication.

Studies using datasets like the Pima Indians Diabetes Database (PIDD) have experimented with models like Decision Trees, Support Vector Machines (SVM), Random Forests, and Neural Networks. While these models show promise, one common issue is the lack of explainability. For medical practitioners, understanding why a model predicts a certain outcome is just as important as the prediction itself.

Additionally, most studies tend to focus solely on accuracy metrics, often overlooking human- centric considerations like usability, interpretability, and ethical transparency.

III. METHODOLOGY

To build a diabetes prediction model that is both practical and clinically relevant, we began by selecting a dataset that reflects real-world medical conditions. We chose the Pima Indians Diabetes Database (PIDD), which is widely used in health-related machine learning research. This dataset includes 768 medical records of female patients over the age of 21, each with eight health-related attributes such as glucose level, BMI, insulin level, number of pregnancies, age, and a diabetes pedigree function.

The outcome column indicates whether or not the individual has diabetes. Before applying any machine learning techniques, we thoroughly cleaned and preprocessed the data to address common challenges found in clinical records, such as missing or inconsistent values. Specifically, we found several instances where features like glucose and BMI had zero entries, which are not physiologically realistic. To correct for this, we treated those

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25apr596

entries as missing and replaced them using median imputation, a method particularly effective for skewed datasets. After cleaning, we normalized the dataset using Min-Max Scaling to bring all the features onto a uniform scale, ensuring that no single attribute had undue influence on model training. With the data prepped, we divided it into two parts—80% for training and 20% for testing—to evaluate how well the models would perform on unseen data. We then applied a variety of supervised machine learning algorithms, including Logistic Regression, Decision Tree, Random Forest, Support Vector Machine (SVM), K-Nearest Neighbors (KNN), Gradient Boosting, and a simple neural network model known as Multilayer Perceptron. This diverse mix allowed us to compare traditional and advanced models in terms of both prediction accuracy and ease of interpretation. To measure performance, we used key metrics such as accuracy, precision, recall, F1-score, and ROC-AUC. These metrics gave us a well-rounded view of each model's strengths and limitations. In particular, we emphasized recall and precision, which are especially important in healthcare settings where failing to identify a diabetic patient could lead to severe consequences. This step-by-step approach helped ensure our final model was not only technically robust but also aligned with the needs and priorities of real-world clinical practice.

IV. RESULTS AND EVALUATION

Once the models were trained using the prepared dataset, their performance was rigorously evaluated using a variety of statistical metrics. However, in the context of healthcare, numbers are not enough—we must understand what these results mean in real terms for patients, doctors, and the broader healthcare ecosystem.

➢ *Model Comparison and Performance Metrics*

We assessed the performance of seven machine learning models: Logistic Regression, Decision Tree, Random Forest, Support Vector Machine (SVM), K-Nearest Neighbors (KNN), Gradient Boosting, and a basic Multilayer Perceptron (Neural Network).

- Accuracy: Overall correctness of the model.
- Precision: The ability of the model to avoid false positives (i.e., predicting a patient has diabetes when they don't).
- Recall (Sensitivity): The ability to catch true positives (i.e., identifying actual diabetic patients).
- F1-Score: Providing a balanced assessment, the harmonic mean of accuracy and memory.
- ROC-AUC: A metric that shows the trade-off between sensitivity and specificity, helpful for understanding model performance across thresholds.

- *Here's a Comparative Summary of model Performance:*

Table 1 Model Comparison and Performance Metrics

Model	Accuracy	Precision	Recall	F1 Score	ROC-AUC
Logistic Regression	78.5%	76.4%	74.2%	75.3%	0.81
Decision Tree	74.3%	70.5%	69.1%	69.8%	0.75
Random Forest	83.2%	80.1%	78.9%	79.5%	0.86
SVM	79.6%	77.8%	73.3%	75.5%	0.82
KNN	76.0%	72.6%	70.0%	71.2%	0.78
Gradient Boosting	84.5%	82.7%	80.1%	81.4%	0.88
Neural Network (MLP)	82.1%	79.4%	76.5%	77.9%	0.85

➢ *Observations and Insights*

From the table above, Gradient Boosting achieved the best results overall, with the highest accuracy (84.5%), strong precision (82.7%), and the highest recall (80.1%) among all tested models. This makes it a reliable choice for healthcare use, where both accuracy and sensitivity are crucial.

The Random Forest model also performed well, offering a balance between accuracy (83.2%) and explainability through feature importance plots. Logistic Regression, while slightly less accurate, remains a preferred model in many clinical setups due to its simplicity and interpretability. On the other hand, models like KNN and Decision Tree struggled slightly with recall, which may not be ideal in a medical context where missing a diabetic case is riskier than a false alarm.

Neural Networks, though promising, require more data to truly shine. In this case, with only 768 entries, their performance was limited. Nonetheless, they still managed to outperform simpler models like KNN and Decision Tree, highlighting their potential for future use with larger datasets.

➢ *Feature Importance and Interpretability*

One of the crucial aspects of building a healthcare model is understanding why it makes the predictions it does. Using models like Random Forest and Gradient Boosting, we extracted the top features influencing the prediction outcomes. These were:

- Glucose Levels: Unsurprisingly, this was the most influential predictor
- BMI (Body Mass Index): Obesity and excess body fat are well-known risk factors for diabetes, making this a highly relevant variable.
- Age: The risk of developing Type 2 diabetes increases with age.
- Diabetes Pedigree Function: This measures the genetic

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25apr596

predisposition to diabetes based on family history.
- Insulin Levels: Although sometimes missing in the dataset, insulin values help in determining insulin resistance.

Understanding feature importance helps doctors trust the model and relate its outcomes to known medical science. For example, if the model predicts high risk in a patient with high glucose, high BMI, and older age, it aligns with clinical reasoning.

➢ *Confusion Matrix Analysis*

To further evaluate performance, we examined the confusion matrix of the Gradient Boosting model:

- True Positives (TP): Patients correctly identified as diabetic.
- False Positives (FP): Patients incorrectly labeled diabetic.
- True Negatives (TN): Patients correctly labeled as non-diabetic.
- False Negatives (FN): Diabetic patients missed by the model.

The model showed a low rate of false negatives, which is vital. In real-world applications, this means fewer diabetic individuals will go undiagnosed—a significant strength of this model.

➢ *Real-World Implications*

In practice, these results mean that a model like Gradient Boosting could be deployed in clinics to screen high-risk individuals using basic health data. The tool could act as a second opinion for doctors or even a preliminary screening tool in rural areas. With interpretability features, clinicians can understand the reasoning behind a model's decision, enhancing trust and adoption.

V. DISCUSSION

Developing a predictive model for diabetes is not just a technical challenge—it's a human- centered mission. The models we trained and evaluated in this study aim to serve as more than mathematical tools; they are intended to support real-world healthcare decision-making. In this section, we reflect not only on the numbers and metrics but on the broader implications of what these results mean for healthcare delivery, policy, and the future of preventive medicine.

➢ *Interpreting the Results in Real-World Context*

While all the machine learning models showed promising performance, our focus extended beyond numerical accuracy. In medical contexts, recall (the ability to catch true cases) is often more important than sheer accuracy. A model that misses a diabetic patient—resulting in a false negative—could delay treatment and lead to serious complications such as neuropathy, kidney failure, or vision loss. On the other hand, false positives (predicting diabetes where it does not exist) may cause temporary anxiety or lead to unnecessary tests, but they are less dangerous. Hence, models like Gradient Boosting, which showed high recall alongside strong precision, are more clinically appropriate.

Another takeaway is that even simpler models like Logistic Regression, while not top performers in accuracy, hold value due to their transparency and ease of interpretation. In clinical environments, trust and explainability are critical—medical professionals must be able to justify decisions, especially when communicating risks to patients. Logistic Regression offers clear coefficients for each input variable, which can be directly linked to patient characteristics, thus facilitating clinical reasoning.

➢ *Human-Centered Design and Interpretability*

Technology should serve people—not the other way around. Predictive models must be designed with humans in the loop, especially when dealing with something as sensitive and impactful as medical diagnoses. Healthcare professionals often hesitate to adopt black-box models (e.g., deep neural networks) unless they can understand the rationale behind the prediction. This is where techniques such as SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-Agnostic Explanations) come into play. These tools offer a layer of interpretability, highlighting which features contributed most to an individual prediction.

For example, if a model flags a patient as high-risk for diabetes, SHAP can show that the decision was largely influenced by high glucose levels and BMI. This aligns with clinical intuition, strengthening the doctor's confidence in the system. Such explainable models are more likely to be accepted and used in real-world healthcare environments.

➢ *Practical Integration into Healthcare Systems*

A high-performing model is only useful if it can be seamlessly integrated into healthcare workflows. For instance, in a primary care setting, the model could be embedded into the patient intake system. As basic vitals and test results are entered into an electronic health record (EHR), the model could run in the background, flagging high-risk patients in real time. The clinician would then receive a simple notification, along with a breakdown of risk factors.

This system can be particularly valuable in resource-limited settings, such as rural clinics or underfunded hospitals, where access to specialists is scarce. A well-designed prediction tool can serve as a digital assistant, augmenting the clinical judgment of general practitioners and helping prioritize patients who need urgent attention.

➢ *Addressing Public Trust and Patient Empowerment*

Trust is essential in healthcare. If patients feel that AI-driven decisions are opaque or arbitrary, they may become skeptical or fearful. Transparency in how predictions are made, and the inclusion of patient-friendly explanations, can foster trust. For instance, if a patient is flagged at high risk, they should be shown a simple explanation like: ―Your risk

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25apr596

is elevated due to high glucose and BMI value.

Moreover, these models can be empowering tools for patients, especially when integrated into mobile apps or wellness platforms. A patient could monitor their risk score over time, adjusting their diet or exercise habits accordingly. This kind of proactive engagement shifts the paradigm from reactive to preventive healthcare.

- *Ethical Responsibility in Model Deployment*

Ethics cannot be an afterthought. From data privacy to decision accountability, we must embed ethical considerations at every stage of model development and deployment. This includes ensuring that patient data is anonymized, securely stored, and only used with informed consent. It also means being transparent about the limitations of the model, clearly stating that it is an aid—not a replacement—for professional diagnosis.

Furthermore, we must think critically about who has access to these tools. Will they be freely available in government hospitals? Can rural clinics afford to deploy them? These questions are vital to ensuring that AI in healthcare does not widen the gap between privileged and underserved populations

With great data comes great responsibility, especially in the world of healthcare. As artificial intelligence and machine learning begin to play a more central role in medical decision-making, it's essential that we go beyond metrics like accuracy and precision. The true value of a predictive model—especially in something as sensitive as diabetes diagnosis—lies in its ethical backbone. Every prediction made has the potential to impact someone's life, so these tools must be built and deployed with human dignity, fairness, and transparency in mind.

VI. ETHICAL CONSIDERATIONS

One of the most pressing ethical priorities is data privacy. Medical records carry deeply personal information—details about an individual's body, habits, and history. If this data is to be used in training or deploying machine learning models, it must be protected with the utmost care. That means securing data through encryption, anonymizing identities, limiting access, and clearly informing patients about how and why their data is being used. True informed consent goes beyond a checkbox—it involves real understanding and choice. Particularly in systems where data might be pulled automatically from electronic health records or fitness trackers, users should always have the option to opt in or out.

Another major issue is bias in model predictions. If the dataset used is skewed or lacks diversity, the model may end up favoring certain groups while failing others. In our own study, we used a dataset focused entirely on adult women from the Pima Indian community. While this gave us useful insights into that demographic, it also showed the risk of narrow representation. If such a model were rolled out across broader populations without being retrained on more diverse data, it could lead to misdiagnosis or unequal treatment. This isn't just a technical oversight—it's a matter of justice. Developers and healthcare organizations must actively audit their models for bias and continuously update them to reflect the people they are meant to serve.

In addition to fairness, explainability is a key part of ethical AI in healthcare. Patients deserve to know why a certain prediction or recommendation is being made. If a model flags someone as at risk for diabetes, they should be able to see which factors contributed to that result—was it their glucose level, BMI, or family history? Blindly following the advice of a "black box" algorithm without explanation can erode trust. By making models more transparent and understandable, we empower both patients and doctors to make informed, collaborative decisions, rather than turning healthcare into a guessing game guided by unseen logic.

Another critical issue is accountability. When an AI tool gets something wrong—perhaps missing a diagnosis or suggesting an incorrect treatment—who takes responsibility? Is it the developer who built the model, the institution that deployed it, or the healthcare professional who used it? Right now, these are grey areas in law and practice. We need clear frameworks that define roles and liabilities so that accountability doesn't get lost in the complexity.

AI-based healthcare tools must not become exclusive luxuries available only in advanced hospitals or to wealthy individuals. In fact, their greatest potential lies in reaching the underserved—rural areas, low-income communities, and public health centers. To truly democratize healthcare, these tools should be made affordable, scalable, and simple to use. Governments and healthcare providers have a shared responsibility to integrate AI into public health programs, making sure the benefits of innovation are shared by all—not just the privileged few.

Ultimately, ethical considerations are not add-ons—they are the foundation of responsible, human-centered healthcare. As we build smarter tools to predict and manage diseases like diabetes, we must also build stronger values into the systems that power them.

VII. FUTURE SCOPE

Integration with Wearables: Combining this model with data from smartwatches or fitness bands While the results of this research demonstrate the strong potential of machine learning models— especially Gradient Boosting—in predicting diabetes, this study is only a starting point. The future holds significant opportunities to improve, expand, and ethically implement these models in diverse real-world settings. With advancements in technology, broader data availability, and increasing integration of AI into healthcare systems, the scope for future research and application is vast and promising.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25apr596

- *Inclusion of More Diverse and Larger Datasets*

One of the most immediate steps forward involves the use of larger and more diverse datasets. The current model was trained on the Pima Indians Diabetes Dataset, which, while widely used in academic circles, is limited in terms of gender, ethnicity, geography, and age range. To make these models applicable across populations, future studies must gather and use data from various countries, ethnic groups, age demographics, and socioeconomic backgrounds.

Incorporating such diversity will not only improve the generalizability and fairness of models but will also help in identifying risk patterns unique to certain communities. This can be particularly valuable in understanding regional trends in diabetes prevalence and risk factors.

- *Integration with Wearable Technology and IoT*

With the rapid growth of wearable health devices, the future of diabetes prediction is moving beyond static health records. Devices like smartwatches, glucose monitors, and fitness trackers continuously collect data on heart rate, physical activity, blood glucose levels, sleep patterns, and more. By integrating machine learning models with real-time data from wearables, we can develop dynamic, continuously learning models that provide real-time risk assessments and early warnings to users and physicians.

This approach transforms diabetes prediction into continuous health monitoring, allowing patients to take proactive measures before symptoms worsen. Such integration also opens the door to personalized healthcare, where predictions and suggestions are tailored to each individual's health patterns and lifestyle.

- *Development of Personalized and Adaptive Models*

Future models could go beyond generalized predictions and evolve into personalized risk assessment tools. By taking into account a person's unique medical history, genetic background, diet, lifestyle, and even stress levels, models can offer more accurate and relevant predictions. With the help of adaptive algorithms, the model could update itself over time as new data is gathered, ensuring that its predictions remain current and personalized.

Such personalization would significantly improve patient engagement and compliance with lifestyle interventions, as the model's outputs would feel more relevant and actionable to the individual.

- *Interdisciplinary Collaboration for Holistic Solutions*

The future of AI in healthcare is not just a matter of data science—it involves interdisciplinary collaboration between computer scientists, medical professionals, public health experts, sociologists, and ethicists. Bringing these voices together ensures that AI tools are technically robust, clinically relevant, socially responsible, and culturally sensitive.

For instance, medical experts can help define the most meaningful features for prediction, while sociologists can guide the understanding of how different populations interact with healthcare systems. Ethicists can ensure fairness and accountability, and public health experts can explore how AI can be used to design community-level diabetes prevention programs.

- *Deployment in Rural and Low-Resource Areas*

In the coming years, one of the most impactful uses of diabetes prediction models could be their deployment in rural or underserved areas, where access to specialized care is limited. By integrating lightweight, efficient models into mobile health clinics or low-cost diagnostic tools, frontline health workers could screen patients with minimal equipment.

This would allow for early intervention and referral to higher-level care before complications arise, ultimately reducing the burden of diabetes-related illnesses in vulnerable populations. Governments and NGOs can work with developers to create user-friendly interfaces and multilingual platforms to facilitate such applications.

- *Integration into Electronic Health Records and Hospital Systems*

Another key area of growth lies in the seamless integration of prediction models into hospital information systems and EHRs. Instead of being standalone tools, future models can act as intelligent modules within hospital software, automatically analyzing incoming data and alerting doctors to potential diabetes risks in real time.

Such smart EHR systems can also maintain patient risk histories, compare them over time, and even suggest preventive measures or referrals. This level of automation can significantly improve the efficiency and effectiveness of routine checkups and chronic disease management.

- *Explainable AI and Decision Support Systems*

As AI tools gain wider adoption, the demand for explainable AI (XAI) will increase. Future work should focus on building models that not only predict but also explain their reasoning in simple, human-understandable terms. This is crucial for ensuring transparency and trust in clinical environments.

Further, diabetes prediction systems can be embedded into clinical decision support systems (CDSS) that guide healthcare professionals on the next steps—whether that be recommending lab tests, suggesting lifestyle interventions, or referring to endocrinologists. Such systems can be especially beneficial in busy outpatient departments or primary care settings where time and resources are limited.

- *Public Health Surveillance and Policy Planning*

Beyond individual diagnoses, future applications of predictive models could support population- level health monitoring. Governments and health organizations can use aggregated, anonymized predictions to track diabetes trends in real time, identify high-risk regions, and allocate resources accordingly.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25apr596

This data-driven approach to public health planning could result in more effective policies, targeted interventions, and awareness campaigns, ultimately helping reduce the overall burden of diabetes in society

VIII. CONCLUSION

Diabetes continues to be a major public health challenge, and early detection is key to managing its impact. In this study, we explored how machine learning can help identify individuals at risk of Type 2 Diabetes by analyzing simple clinical and physiological data. Among the various models tested, Gradient Boosting delivered the highest performance, though more interpretable models like Logistic Regression also showed practical strengths for real-world healthcare use.

Ethical responsibility is just as critical as technical success. Protecting patient privacy, ensuring model transparency, and eliminating bias are essential for earning and maintaining trust. These principles must guide how such tools are developed and deployed, especially when lives are at stake.

Looking ahead, these models have the potential to transform care—from enhancing clinical decision-making to supporting community health efforts in underserved regions. With careful implementation, collaboration across disciplines, and a commitment to fairness, predictive models can move healthcare toward a more proactive, personalized, and compassionate future.

REFERENCES

[1]. U.C. Irvine Machine Learning Repository, Pima Indians Diabetes Dataset, UCI Machine Learning Repository, https://archive.ics.uci.edu/ml/datasets/pima+indians+diabetes (last visited Apr. 8, 2025).
[2]. T. Chen & C. Guestrin, XGBoost: A Scalable Tree Boosting System, 22 Proc. ACM SIGKDD Int’l Conf. on Knowledge Discovery & Data Mining 785 (2016).
[3]. L. Breiman, Random Forests, 45 Machine Learning 5 (2001).
[4]. V.N. Vapnik, The Nature of Statistical Learning Theory (Springer 1995).
[5]. D. Witten, E. Frank, M. Hall & C. Pal, Data Mining: Practical Machine Learning Tools and Techniques (4th ed. 2016).
[6]. Z. Obermeyer & E.J. Emanuel, Predicting the Future — Big Data, Machine Learning, and Clinical Medicine, 375 New Eng. J. Med. 1216 (2016).
[7]. S. Panch, H. Mattie & L. Atun, Artificial Intelligence and Algorithmic Bias: Implications for Health Systems, 372 New Eng. J. Med. 2508 (2019).
[8]. World Health Org., Diabetes, World Health Organization, https://www.who.int/news-room/fact-sheets/detail/diabetes (last visited Apr. 8, 2025).
[9]. A. Rajkomar, J. Dean & I. Kohane, Machine Learning in Medicine, 380 New Eng. J. Med. 1347 (2019).
[10]. J. Ribeiro, F. Torgo & L. Neves, Ensemble Methods for Classification, in Data Mining and Analysis in the Engineering Field 33–45 (IGI Global 2014).
[11]. A. Holzinger, From Machine Learning to Explainable AI, 77 Int’l J. Interactive Multimedia & Artificial Intelligence 41 (2018).
[12]. U.S. Dep’t of Health & Human Servs., HIPAA Privacy Rule, https://www.hhs.gov/hipaa/for-professionals/privacy/index.html (last visited Apr. 8, 2025).
[13]. C. Dwork et al., Fairness Through Awareness, in Proc. of the 3rd Innovations in Theoretical Computer Science Conf. 214 (2012).
[14]. N. Mehrabi et al., A Survey on Bias and Fairness in Machine Learning, 54 ACM Computing Surveys 1 (2021).