

AI-Driven Diabetes Risk Assessment and Personalized Health Communication: A Review of Machine Learning, Explainable AI, and Generative AI

Anurag Gupta, Dr. Gurvinder Singh & Dr. Mahesh Kumar Singh

Research Scholar, Department of CSE, Arni University Indora, H.P.

Dean Engineering, Arni University, Indora, H.P.

Associate Professor, Dronacharya Group of Institutions, Greater Noida, India.

ARTICLE INFO

Received: 05 Jan 2024
Revised: 17 March 2024
Accepted: 27 March 2024

ABSTRACT

Generative Artificial Intelligence (GenAI), Large Language Models (LLMs), and Retrieval-Augmented Generation (RAG). However, these technologies have largely evolved as separate research directions, creating a gap between accurate risk prediction and understandable, trustworthy health communication. This critical review examines the progression of AI-based approaches for diabetes risk assessment, with particular emphasis on ML prediction, XAI, GenAI/LLMs, RAG, and intelligent healthcare assistants. The review synthesizes existing evidence to examine how ML models support risk estimation, how SHAP and LIME improve model interpretability, and how GenAI and LLMs enable natural-language and personalized healthcare communication. It further considers RAG as a mechanism for grounding generated responses in external knowledge sources and reducing dependence on internally generated information. The critical synthesis identifies persistent gaps in the integration of prediction, explanation, evidence grounding, personalization, and safety. Existing systems frequently provide either predictive outputs, technical explanations, or general conversational support rather than combining these capabilities within a unified diabetes risk-awareness framework. Based on these findings, an integrated conceptual framework is proposed linking ML prediction, XAI-based risk interpretation, RAG, GenAI/LLM-based communication, personalization, and safety validation. The review highlights the need for future empirical validation, standardized evaluation, privacy-aware personalization, and patient-centred assessment of integrated AI healthcare assistants.

Keywords: Artificial Intelligent, GenAI, XAI, LLMs, RAG.

Introduction

Diabetes mellitus is one of the major chronic health conditions worldwide and represents a significant challenge for modern healthcare systems. Persistent abnormalities in blood glucose regulation can progressively contribute to complications affecting the cardiovascular system, kidneys, eyes, nerves, and other organs. The increasing prevalence of diabetes places a substantial clinical, economic, and social burden on individuals, healthcare providers, and public health systems, particularly because of long-term treatment, monitoring, productivity loss, and disease-related complications. The burden is particularly important in low- and middle-income countries, where healthcare resources, awareness, and access to preventive screening may vary across populations. India represents an important context because of its large population, changing lifestyle patterns, urbanization, dietary transitions, reduced physical activity, and increasing exposure to metabolic risk factors. The interaction of genetic, demographic, behavioural, and environmental factors makes diabetes prevention and early risk identification an important public-health priority [1], [2]. Diabetes-related metabolic abnormalities and risk factors may develop before formal diagnosis. Early identification of individuals who may be at

increased risk can therefore support timely health awareness, screening, preventive lifestyle measures, and appropriate consultation with healthcare professionals [2], [4]. Artificial Intelligence (AI) and Machine Learning (ML) have increasingly been investigated for analysing demographic, clinical, and behavioural characteristics associated with diabetes risk [5]. However, the value of AI-based diabetes risk assessment extends beyond predictive accuracy. A risk probability or classification alone may not explain why an individual has been identified as being at risk or how the result should be understood. This has encouraged research toward explainable, personalized, and conversational AI-based healthcare systems that can connect prediction with understandable and responsible health communication [3], [8], [15].

Need for Early Diabetes Risk Assessment

Early identification of individuals at elevated risk of diabetes is an important component of preventive healthcare. Conventional healthcare pathways rely on clinical assessment, laboratory investigations, and professional interpretation, which remain essential for diagnosis and clinical decision-making. Computational risk-assessment approaches can complement these processes by analysing multiple risk-associated characteristics simultaneously and providing preliminary risk awareness [4], [16]. Machine learning-based risk assessment has attracted considerable attention because ML algorithms can identify relationships among multiple input variables and outcome patterns. Demographic and clinical characteristics such as age, glucose-related measurements, body mass index (BMI), blood pressure, family-history-related factors, and other risk indicators can be analysed to estimate the likelihood of an adverse outcome [2], [16], [19]–[22]. Such systems may assist in identifying individuals who could benefit from further screening or professional evaluation. Importantly, risk assessment should be distinguished from clinical diagnosis. An AI model estimates the likelihood of an outcome based on patterns learned from data; it does not establish the presence of diabetes. Therefore, AI-based diabetes risk assessment should be positioned primarily as a health-awareness and decision-support mechanism rather than as a replacement for laboratory diagnosis or professional healthcare [1], [3], [6], [15]. Another important requirement is the effective communication of risk information. Numerical probabilities or binary classifications may be difficult for non-technical users to interpret and may not indicate which characteristics contributed to the estimated risk. Consequently, effective diabetes risk assessment requires not only prediction but also meaningful interpretation and understandable communication [8], [9], [25], [28]. This requirement has led to broader research combining predictive AI with Explainable Artificial Intelligence (XAI), Generative AI, Large Language Models (LLMs), and evidence-grounded information retrieval. XAI can help identify factors contributing to predictions, while GenAI/LLMs can transform structured outputs into understandable natural-language communication. Retrieval-Augmented Generation (RAG) can further provide access to relevant external knowledge sources to support evidence-grounded responses [5]–[8], [11], [12]. Such integration can shift diabetes risk assessment from a prediction-centred approach toward a more user-oriented framework focused on explanation, evidence, personalization, and responsible communication.

Evolution from Conventional Prediction to AI-Based Risk Assessment

The development of diabetes risk assessment has gradually progressed from conventional statistical methods toward machine learning and Artificial Intelligence (AI). Traditional statistical models have been used to examine relationships between clinical variables and diabetes outcomes and to estimate disease-related risk. Their relatively interpretable mathematical structure continues to make them useful for understanding statistical associations [22]. With the increasing availability of healthcare datasets and computational advances, Machine Learning (ML) methods such as Logistic Regression, Support Vector Machines (SVM), Decision Trees, Random Forest, and ensemble techniques have been widely applied to diabetes prediction and classification [2], [18], [20], [40]. More recently, boosting and deep learning approaches have been explored for modelling nonlinear relationships and complex patterns in healthcare data [2], [4], [16], [19], [21]. This progression represents a shift from predefined

statistical relationships toward data-driven pattern recognition. However, increasing model complexity has also highlighted the need for transparency, particularly when predictions are used in healthcare settings. A predictive result alone may not explain the factors contributing to an individual's estimated risk [25], [28]. The development has therefore expanded beyond prediction. XAI approaches aim to improve understanding of model decisions [8], [9], while Generative AI and LLMs can transform structured model outputs into natural-language and conversational communication [11], [12], [29]–[32]. RAG further supports this process by retrieving information from external knowledge sources before generating responses [7]. Collectively, these developments indicate a movement toward AI systems that combine prediction, interpretability, evidence grounding, personalization, and responsible communication [5]–[8], [25], [33], [34], [36].

Limitations of Conventional ML-Based Prediction

Despite substantial progress, conventional ML-based diabetes prediction has several limitations when considered from a patient-centred healthcare perspective.

First, interpretability remains a major concern. Many ML systems provide a probability, class label, or risk category without clearly explaining the factors contributing to the prediction. This can limit understanding and transparency in healthcare applications [25], [28]. Second, predictive performance does not directly measure communication quality. Metrics such as accuracy, precision, recall, F1-score, and ROC-AUC are important for evaluating predictive models, but they do not indicate whether the resulting information is understandable or useful to patients [29], [31], [33]. Third, individual-level explanations are required in addition to population-level feature importance. Global importance describes general model behaviour, whereas individual predictions may be influenced by different combinations of characteristics. Consequently, healthcare-oriented systems benefit from both global and local explanations [8], [9], [26]. Personalization is another limitation. Conventional ML models generally convert structured patient variables into a prediction but do not inherently provide individualized educational or contextual information. Users with different risk profiles may therefore receive similar forms of output [29], [31], [34], [36].

Generalizability is also an important concern. Diabetes datasets can vary in population characteristics, sample size, data quality, feature distributions, and healthcare settings. Performance observed on one dataset may therefore not automatically transfer to other populations, highlighting the importance of external validation using diverse and representative datasets [2], [4], [16], [21].

Finally, extending ML prediction into conversational healthcare introduces additional concerns, including factual accuracy, hallucination, bias, privacy, inappropriate recommendations, and safety [5], [6], [29], [33], [34], [36], [37]. These issues indicate that predictive performance alone is insufficient for developing a trustworthy AI-enabled diabetes health-awareness system.

Accordingly, recent research increasingly considers ML prediction, XAI, Generative AI/LLMs, and RAG as complementary components. This progression provides the basis for examining how predictive results can be explained, evidence-grounded, personalized, and communicated responsibly [5]–[8], [11], [15], [25].

Emergence of XAI, GenAI and LLMs

The limitations of conventional machine learning prediction have contributed to the emergence of Explainable Artificial Intelligence (XAI) as an important research direction in healthcare. XAI aims to provide meaningful information about the factors contributing to an AI model's decision. In diabetes risk assessment, explainability can help identify the variables that have the greatest influence on predictions and can provide individual-level explanations for specific risk estimates [8], [9], [13], [14]. Among the widely used model-agnostic XAI approaches are SHapley Additive exPlanations (SHAP) and Local Interpretable Model-Agnostic Explanations (LIME). SHAP can be used to analyse the contribution of individual features to model predictions at both global and local levels, whereas LIME focuses on constructing interpretable local approximations of complex models. Such techniques can make machine learning outputs more transparent and can potentially improve communication between

AI systems, healthcare professionals, and users [8], [9], [25], [28]. However, XAI alone does not completely solve the problem of patient-centred communication. Feature contribution values, importance rankings, and local explanation plots may be understandable to researchers or technical users but can remain difficult for individuals without a machine learning background. Consequently, an additional mechanism is required to transform structured explanations into clear and understandable health information [25], [28], [34]. The emergence of Generative Artificial Intelligence (GenAI) and Large Language Models (LLMs) has created new possibilities in this area. Unlike conventional predictive models that primarily produce structured outputs, generative models can generate natural-language responses, summarize information, answer questions, and support conversational interaction. These capabilities have increased interest in healthcare applications including patient education, health information support, documentation, decision support, and personalized communication [5], [11], [12], [29], [33], [34]. For diabetes risk awareness, GenAI can potentially act as a communication layer over an underlying predictive model. Instead of directly presenting a probability score or technical XAI output, a generative system can receive structured information such as estimated risk, important contributing factors, and relevant patient characteristics and transform these inputs into a more understandable explanation. Such an architecture addresses two different requirements. ML provides the computational risk estimate, while XAI provides interpretability regarding the prediction. GenAI can then translate these structured outputs into patient-oriented language. This separation is important because the generative component should not independently determine a patient's risk when a validated predictive model is intended to provide that assessment [5], [8], [15], [25]. At the same time, the introduction of GenAI and LLMs into healthcare creates additional concerns. Generated responses may contain factual errors, hallucinated information, inappropriate recommendations, or statements that could be interpreted as clinical diagnoses. Issues related to transparency, privacy, bias, and safety therefore become particularly important in health-related applications. The literature consequently emphasizes the need for evidence-grounded generation, controlled response mechanisms, and appropriate human or professional oversight [6], [11], [12], [29], [33], [36]. The development of XAI and GenAI/LLM technologies therefore represents an important transition from “Can AI predict diabetes risk?” toward broader questions such as “Why was this risk predicted?”, “How can the result be explained to the individual?”, and “How can relevant health information be communicated accurately and safely?” These questions provide the foundation for examining Retrieval-Augmented Generation as a complementary approach for improving the reliability and evidence grounding of AI-generated healthcare communication [5], [7], [11], [12].

Motivation and Research Gap

The preceding discussion demonstrates that artificial intelligence has progressively expanded from conventional disease prediction toward explainable and conversational healthcare applications. Machine learning models can identify patterns associated with diabetes risk, while XAI techniques can provide information regarding the factors contributing to individual predictions. GenAI and LLMs can further transform structured information into natural-language communication, and RAG can provide access to external knowledge sources to support evidence-grounded responses.

Despite these advances, the existing research landscape remains fragmented. A substantial proportion of diabetes-related AI research continues to focus primarily on predictive performance, with emphasis on classification accuracy, sensitivity, specificity, F1-score, or ROC-AUC. Although these measures are essential for evaluating predictive models, they do not adequately describe whether the resulting information is interpretable, personalized, understandable, and safe for end users.

Similarly, the incorporation of XAI improves transparency but does not necessarily solve the communication problem. SHAP and LIME can identify important features and explain individual predictions, but their outputs are generally designed for analytical interpretation rather than direct patient communication. Consequently, there remains a gap between technical model explanation and patient-oriented health information.

Generative AI and LLM-based healthcare systems address part of this gap by providing natural-language interaction. However, many existing conversational healthcare systems function primarily as general-purpose information or question-answering tools. The thesis literature analysis indicates that such systems often do not combine validated diabetes-specific predictive analytics with explainability and personalized risk information [4], [13], [14], [25].

However, relatively limited research has systematically examined how these components can be brought together into a unified diabetes risk-awareness framework. The literature comparison in the thesis similarly identifies that existing healthcare assistants rarely integrate diabetes risk prediction, XAI, Generative AI, and personalization within a single framework. A further gap concerns patient-level personalization. An individual's risk profile may be influenced by a particular combination of clinical and demographic characteristics. Therefore, providing the same generic diabetes information to every user may not adequately address individual needs. A more useful healthcare assistant should be capable of using the individual's predicted risk and explanatory factors as contextual inputs while retrieving appropriate evidence and generating understandable health information.

Another important gap concerns safety and responsible communication. Healthcare-oriented GenAI systems must avoid presenting probabilistic predictions as confirmed diagnoses and should not generate inappropriate medication or treatment recommendations. Therefore, an integrated framework requires explicit safety mechanisms in addition to prediction, explanation, and generation. The thesis positions the proposed system as a health-awareness and information-support framework rather than a replacement for professional diagnosis or clinical decision-making.

Based on these observations, the principal motivation for this review is to critically examine the evolution and convergence of AI technologies used for diabetes risk assessment and personalized healthcare communication. Rather than reviewing machine learning, XAI, and GenAI as completely independent research areas, this review investigates their relationships and identifies opportunities for integration.

Literature Review

AI-Based Approaches for Diabetes Risk Assessment

Early diabetes risk assessment approaches were largely based on conventional statistical methods that examined relationships between clinical characteristics and diabetes outcomes. Logistic Regression has remained an important baseline because of its relatively simple mathematical structure and interpretability [19]–[22]. However, predefined statistical assumptions may be less suitable for representing complex nonlinear relationships and interactions among the multiple demographic, physiological, metabolic, and behavioural factors associated with diabetes risk [4], [16], [19]–[21]. The increasing availability of healthcare datasets subsequently encouraged the use of Machine Learning (ML) approaches for diabetes prediction. Studies have investigated algorithms including Logistic Regression, Decision Trees, Support Vector Machines (SVM), k-Nearest Neighbours (KNN), Naïve Bayes, Random Forest, and ensemble methods [18]–[22], [40]. These approaches can analyse multiple variables simultaneously and learn patterns from data, making them suitable for modelling complex relationships associated with diabetes risk [2], [4], [16]. Model performance, however, depends not only on the algorithm but also on dataset characteristics, predictor selection, preprocessing, class distribution, and validation strategy [2], [21], [40]. Reviews and meta-analyses have reported substantial variation in predictive performance across diabetes risk studies, indicating that the performance of a particular algorithm on one dataset should not be treated as universally applicable to other populations or settings [2], [16], [21]. Logistic Regression continues to provide an interpretable baseline, while tree-based and ensemble methods can capture more complex relationships. SVM can model complex classification boundaries, whereas approaches such as Naïve Bayes offer computational efficiency under their underlying assumptions [18], [19], [21], [40]. However, the literature does not establish a universally superior algorithm across diabetes datasets. Differences in population, feature distributions, sample size, and methodological design can substantially influence reported results [2], [21], [40]. More recent research has therefore moved beyond prediction alone toward interpretable and

application-oriented AI systems. Although ML can provide useful risk estimates, healthcare applications also require transparency, robustness, fairness, privacy, and reliability [6], [25], [28]. This has encouraged the integration of ML prediction with XAI and other supporting technologies. Overall, ML provides an important computational foundation for diabetes risk assessment by enabling multiple risk-related characteristics to be analysed simultaneously. However, predictive performance alone does not explain the basis of an individual prediction or communicate its meaning in patient-oriented language. These limitations motivate the subsequent examination of XAI for model interpretation and GenAI/LLMs for personalized health communication.

Ensemble and Advanced ML Approaches

Ensemble learning has become an important approach for improving the predictive performance and robustness of diabetes risk models. Methods such as Random Forest, AdaBoost, Gradient Boosting, and XGBoost combine multiple learners to capture nonlinear relationships and interactions among diabetes-related variables [20], [24], [27], [40]. Random Forest aggregates predictions from multiple decision trees, whereas boosting methods sequentially improve model performance by addressing errors from previous learners. XGBoost has been specifically investigated for type 2 diabetes risk prediction, demonstrating the potential of boosting-based approaches for modelling complex diabetes-related risk factors [20]. Comparative studies have also reported the use of ensemble and multi-classifier approaches across diabetes prediction tasks [24], [27], [40]. However, reported performance varies with dataset characteristics, feature distributions, preprocessing, and validation strategy; therefore, no single ensemble method can be considered universally optimal across diabetes datasets. An important limitation of increasingly complex ensemble models is reduced direct interpretability. Although such models may provide strong predictive discrimination, their combined decision processes can be difficult to communicate to healthcare users. Consequently, XAI techniques such as SHAP and LIME are increasingly used alongside complex ML models to identify influential features at both global and individual levels [8], [9], [25]–[28]. Thus, the evaluation of ensemble and advanced ML approaches should extend beyond predictive metrics to include interpretability, robustness, generalizability, and relevance to healthcare applications. This provides an important basis for integrating predictive modelling with explainability and patient-oriented communication.

Deep Learning Approaches

The increasing availability of large and high-dimensional healthcare datasets has encouraged the application of Deep Learning (DL) techniques to diabetes prediction and management. Deep Neural Networks and recurrent neural architectures have been investigated using structured clinical data, electronic health records, and sequential physiological measurements [16], [38]. Recurrent approaches, for example, have been explored for blood-glucose prediction using temporal physiological information [38]. A major advantage of DL is its ability to learn complex data representations with comparatively limited manual feature construction, particularly when large and high-dimensional datasets are available. However, DL models generally require greater computational resources and larger datasets than many conventional ML approaches. Their complex internal representations can also make individual predictions difficult to interpret, creating an important transparency challenge in healthcare [25], [28]. XAI methods are therefore increasingly combined with complex predictive models to identify influential features and improve the interpretability of model decisions [25], [26]. Furthermore, the suitability of a DL architecture depends strongly on the data modality and prediction task. Structured clinical prediction, medical imaging, and sequential glucose forecasting involve different data characteristics and therefore require different modelling strategies [16], [38]. Overall, deep learning expands the methodological possibilities for diabetes-related prediction and management, but greater model complexity does not necessarily imply greater practical suitability for patient-centred risk assessment. Data requirements, computational cost, interpretability, generalizability, and validation should therefore be considered alongside predictive performance.

Feature engineering and feature selection

Feature engineering and feature selection are important components of diabetes risk prediction because healthcare datasets may contain redundant, weakly informative, or highly correlated variables. Selecting relevant predictors can reduce model complexity, improve computational efficiency, and potentially enhance model generalization [40]. The literature indicates that predictive performance can vary substantially with different feature configurations and that the effect of feature selection is model-dependent [2], [16], [21], [40]. Therefore, removing variables based solely on their apparent importance in one model may not necessarily improve another. Feature selection can also contribute to interpretability by identifying variables that are influential for prediction [8], [9], [25], [26]. In patient-centred applications, carefully selected predictors may make model outputs easier to interpret and communicate. However, predictive feature importance should not be treated as evidence of causality, and statistical contribution should be distinguished from clinical relevance [25], [28]. Thus, feature engineering and selection should be considered not only preprocessing steps but also components that influence prediction, interpretability, and subsequent risk communication.

Class Imbalance and Model Evaluation

Class imbalance is an important challenge in diabetes prediction because some datasets contain substantially fewer positive or high-risk cases than non-diabetes cases. Under such conditions, a model may achieve high overall accuracy while performing poorly in identifying minority-class observations [2], [21], [40]. Common strategies include Synthetic Minority Over-sampling Technique (SMOTE), random over-sampling, under-sampling, class weighting, and ensemble-based approaches [17]. However, balancing procedures must be incorporated carefully within the validation framework. Applying oversampling before data partitioning can introduce information leakage and produce overly optimistic performance estimates. Consequently, balancing should be performed using only the training data within an appropriate validation procedure. Model evaluation should extend beyond accuracy and include complementary measures such as precision, recall, F1-score, ROC-AUC, and PR-AUC [2], [21], [40]. Recall is particularly relevant when identifying potentially at-risk individuals is important, while precision reflects the reliability of positive predictions. F1-score provides a balance between precision and recall, and ROC-AUC evaluates discrimination across classification thresholds. PR-AUC can provide additional insight when the positive class is relatively less frequent. Reported performance varies considerably across diabetes prediction studies because of differences in datasets, populations, preprocessing, and validation methods [2], [4], [16], [21]. Therefore, a single accuracy value provides an incomplete assessment of model usefulness. Where appropriate, evaluation should also consider calibration and threshold analysis. The operating threshold is another application-dependent consideration. A default probability threshold may not provide the desired balance between false-positive and false-negative classifications for every use case. Threshold selection should therefore be treated as a modelling decision based on the intended application rather than as a universal clinical cut-off. Overall, reliable diabetes risk-model evaluation requires leakage-aware validation, appropriate handling of class imbalance, multiple complementary performance measures, and careful consideration of the intended operating threshold.

Critical Analysis of Existing Prediction Approaches

The reviewed literature demonstrates substantial progress in AI-based diabetes risk assessment. Conventional statistical models provide interpretable baselines, while ML, ensemble, boosting, and deep learning approaches offer greater flexibility for modelling complex and nonlinear relationships across different healthcare datasets [16], [19]–[22], [38], [40]. However, increasing algorithmic complexity does not consistently result in superior or more clinically useful prediction. Reported performance varies with population characteristics, dataset size and distribution, feature selection, preprocessing, class imbalance, validation methodology, and evaluation metrics [2], [16], [21], [40]. Therefore, algorithm selection and reported performance need to be interpreted within the context of the dataset and prediction task rather than through direct cross-study comparison alone. A further

observation is that much of the existing literature remains prediction-centred, with emphasis on accuracy and discrimination. Comparatively less attention is often given to explaining individual predictions and communicating their meaning to non-technical users. This has encouraged the integration of XAI techniques such as SHAP and LIME to improve transparency and interpretation [8], [9], [13], [14], [25], [26]. Importantly, predictive feature importance should not be interpreted as evidence of clinical causation. A variable may contribute substantially to a model's prediction because of statistical patterns in the data without establishing a causal relationship with diabetes. This distinction is particularly important when technical model outputs are converted into patient-oriented information [25], [28]. Personalization also remains a limitation of conventional prediction systems. Risk models generally produce a probability, class label, or risk estimate but do not inherently explain how the result relates to an individual's specific characteristics. Even when XAI is incorporated, technical explanations may require further transformation before they can be effectively understood by non-technical users [25], [28], [34], [36]. Overall, the literature indicates a broader shift from focusing solely on predictive performance toward understanding why a particular risk is predicted, how that prediction can be interpreted, and how it can be communicated meaningfully to the individual. This provides the rationale for examining XAI as a key component for improving transparency and interpretability in diabetes risk assessment.

Explainable AI in Diabetes Prediction

The increasing use of ML for diabetes risk prediction has improved the identification of individuals who may be at elevated risk, but predictive performance alone may be insufficient for healthcare applications. Conventional models often provide a probability or class label without clearly indicating which characteristics contributed to the prediction, limiting transparency for healthcare professionals and non-technical users. Explainable Artificial Intelligence (XAI) has therefore emerged as an important approach for improving the interpretability of healthcare AI systems [25], [28]. XAI methods provide information about the contribution of individual features to model predictions. In diabetes risk assessment, such explanations can indicate which demographic, physiological, or clinical characteristics influence an estimated risk and can therefore complement predictive modelling [13], [14], [27]. Among widely used approaches, SHapley Additive exPlanations (SHAP) estimates feature contributions at both global and individual levels. Aggregated SHAP values can indicate features that generally influence model behaviour, while individual SHAP explanations can show how specific features contribute to a particular prediction [8], [26]. Local Interpretable Model-Agnostic Explanations (LIME) provides local interpretations by approximating the behaviour of a complex model around a specific observation using an interpretable surrogate model [9]. Both approaches are therefore relevant to individualized diabetes risk interpretation. A distinction between global and local explanations is particularly important. Global explanations describe overall model behaviour and commonly influential features, whereas local explanations focus on the factors associated with a specific individual's prediction [8], [9], [26]. Combining these perspectives can support assessment of general model behaviour while also providing patient-specific information [25], [28]. However, the presence of XAI does not automatically make an AI system understandable to patients. Feature rankings, numerical contribution values, and technical visualizations may remain difficult for non-technical users to interpret [25], [28]. This creates an important distinction between model interpretability and patient-centred explanation: an explanation may be mathematically meaningful while still requiring transformation into understandable health language [34], [36]. Another limitation is that feature importance should not be interpreted as evidence of causality. XAI primarily explains the behaviour of a predictive model and does not establish that a highly influential feature causes diabetes [25], [26], [28]. Explanations may also vary with the underlying model, dataset, feature correlations, and explanation method. Consequently, XAI outputs require careful interpretation and validation before being translated into patient-oriented health information.

Generative AI and LLMs in Healthcare

The emergence of Generative Artificial Intelligence (GenAI) and Large Language Models (LLMs) has expanded the role of AI in healthcare from predictive modelling toward natural-language communication, information support, question answering, education, and interactive healthcare applications [5], [11], [12], [29], [33], [34]. Unlike conventional predictive models that primarily generate numerical or categorical outputs, LLMs can transform complex information into natural-language responses and support conversational interaction. Medical LLMs such as Med-PaLM and evaluations of systems such as ChatGPT and GPT-4 have demonstrated the potential of LLMs for medical question answering and healthcare information tasks [11], [12], [30]–[32]. However, performance on question-answering or examination-style tasks does not by itself establish reliability for real-world healthcare use. Existing research highlights concerns related to factual accuracy, hallucination, bias, contextualization, safety, and clinical oversight [6], [11], [12], [29], [33], [36], [37]. Thus, language-generation capability and healthcare reliability should be considered distinct properties. For diabetes risk awareness, GenAI provides an opportunity to extend conventional prediction systems toward personalized health education and information support. A static risk score or classification may provide limited value if users cannot understand its meaning or relevance to their own characteristics. GenAI can potentially convert structured information into accessible natural-language communication while incorporating relevant contextual factors such as age, BMI, glucose-related characteristics, lifestyle information, family history, and other available risk indicators. XAI can further support this personalization by providing structured information about the factors contributing to an individual's prediction [8], [9], [25], [26], [28]. Accordingly, the role of GenAI in the proposed research direction is primarily to transform predictive and explanatory outputs into understandable, personalized health information rather than independently determine diabetes risk. In an integrated diabetes risk-awareness architecture, the ML model estimates risk, XAI identifies relevant contributing factors, and the generative layer communicates these results in natural language. RAG can additionally retrieve relevant information from defined knowledge sources to support evidence-grounded responses and reduce dependence on unsupported model-generated knowledge [7]. This separation maintains a clear distinction between risk estimation and language generation while supporting more accessible communication. Despite these opportunities, the use of GenAI in healthcare requires appropriate safeguards. Generated responses may contain inaccurate, incomplete, outdated, or unsupported information, even when expressed fluently. Therefore, GenAI-based healthcare applications should incorporate evidence grounding, response validation, safety controls, and appropriate professional oversight [6], [11], [12], [29], [33], [36], [37]. Within diabetes risk awareness, GenAI should consequently function as a communication and health-education component rather than a replacement for professional diagnosis or clinical decision-making.

These concerns indicate that the value of GenAI in healthcare cannot be judged solely by the fluency or naturalness of generated responses. Healthcare communication requires factual consistency, contextual relevance, personalization, and safety. Therefore, mechanisms for grounding generated responses in trusted external information and validating their content are necessary. This requirement provides the motivation for Retrieval-Augmented Generation (RAG), which combines external knowledge retrieval with generative language models.

Overall, the literature indicates that XAI represents an important step toward transparent diabetes risk prediction by addressing the question of why a model produces a particular prediction. However, explainability alone does not fully address the challenge of transforming technical model outputs into understandable, personalized health information. This limitation provides a strong rationale for examining Generative AI and Large Language Models as a communication layer capable of converting structured predictions and explanations into natural-language healthcare information.

RAG-Based Healthcare Communication

Retrieval-Augmented Generation (RAG) has emerged as an approach for improving the reliability of Large Language Model (LLM) responses by retrieving relevant information from external knowledge

sources before generation [7]. Unlike conventional LLMs that rely primarily on pretrained knowledge, RAG provides retrieved evidence as contextual information, allowing responses to be grounded in defined sources. In healthcare, the retrieval layer can incorporate trusted medical literature, clinical guidelines, institutional resources, or curated healthcare knowledge bases. This evidence-grounding capability can potentially improve factual consistency, but RAG does not guarantee correctness. The quality of the final response depends on the relevance and reliability of retrieved information, the generation model, and the validation and safety mechanisms used [6], [11], [12], [29], [33], [36], [37]. For diabetes-related health communication, RAG can retrieve relevant information on risk factors, lifestyle practices, screening, prevention, and disease-management education. When combined with ML-based risk assessment and XAI, retrieved information can also provide context for an individual's risk profile. However, RAG does not perform risk prediction or explain a model's decision. Its primary function is evidence retrieval. An integrated system can therefore assign separate roles to ML for risk estimation, XAI for interpretation, RAG for evidence retrieval, and GenAI/LLM for natural-language communication. Evaluation of RAG-based healthcare systems should extend beyond linguistic quality to include factual consistency, evidence relevance, personalization, usability, and safety. In particular, retrieval quality is critical because irrelevant or outdated evidence can lead to inappropriate generated responses. Current research therefore indicates a need for broader evaluation frameworks covering retrieval, factuality, personalization, relevance, and safety [6], [11], [12], [29], [33], [36], [37]. Thus, RAG should be considered a grounding mechanism rather than a complete solution to GenAI reliability. Its value in diabetes healthcare is likely to depend on its integration with reliable knowledge sources, controlled generation, response validation, predictive modelling, and explainability.

Intelligent Healthcare Assistants


The convergence of Machine Learning, Explainable AI, Generative AI, Large Language Models, and knowledge-grounding approaches has led to the development of Intelligent Healthcare Assistants (IHAs). These systems support healthcare information, patient education, conversational guidance, and self-management while enabling more flexible interaction than traditional rule-based systems [6], [29], [33]. Healthcare conversational systems have evolved from predefined rule-based interactions toward NLP- and AI-enabled systems capable of more adaptive communication. Existing studies demonstrate potential for healthcare information and patient support, but also highlight concerns regarding evaluation, safety, reliability, and clinical applicability [34], [35]–[37]. In diabetes applications, ML-based systems have primarily focused on prediction and classification, whereas conversational assistants generally emphasize health information and education [25], [28], [34], [36], [40]. These functions have therefore often developed separately. The integration of predictive modelling with explainability and natural-language communication provides an opportunity for a more comprehensive diabetes risk-awareness assistant. LLMs such as ChatGPT and medical language models have expanded conversational flexibility, but conversational capability alone does not provide individualized risk prediction or explain the factors underlying a prediction. Concerns regarding hallucination, factual reliability, bias, clinical validation, and safe deployment further indicate the need for additional mechanisms [5], [6], [11], [12], [29], [33]. RAG-based assistants address part of this limitation by connecting LLMs with external knowledge sources [7]. However, RAG primarily provides evidence grounding and does not independently perform risk prediction or explainability. A diabetes-focused assistant therefore requires coordinated components for risk estimation, explanation, evidence retrieval, personalized communication, and safety validation. The literature indicates that existing systems commonly emphasize individual functions such as prediction, education, conversational interaction, or knowledge retrieval. Comparatively fewer systems combine these capabilities within a single diabetes risk-awareness architecture. Such an assistant can be conceptualized as a system that estimates potential risk, identifies contributing factors, retrieves relevant evidence, and communicates the resulting information in understandable language. Evaluation should consequently extend beyond conversational quality to include predictive performance, explainability, factual accuracy, personalization, relevance, usability, and safety. The assistant should remain positioned as a health-

awareness and information-support system rather than an autonomous diagnostic system. Overall, the literature supports a transition from isolated AI capabilities toward integrated healthcare assistance. The combination of diabetes risk prediction, XAI, RAG, GenAI-based communication, personalization, and safety-aware interaction provides the basis for the subsequent research-gap and conceptual-framework analysis.

Critical Synthesis and Research Gaps

The reviewed literature shows a gradual transition in diabetes-related AI research from conventional statistical and machine learning-based prediction toward explainable, generative, conversational, and evidence-grounded systems. However, these capabilities are often investigated independently. As a result, predictive performance, interpretability, personalization, evidence grounding, and safety are not consistently addressed together within diabetes risk-awareness applications. Machine learning provides an important foundation for diabetes risk prediction using demographic, clinical, physiological, and lifestyle variables. Approaches including Logistic Regression, Decision Trees, Support Vector Machines, Random Forest, boosting, and deep learning have been widely investigated [2], [16], [38], [40]. However, variations in datasets, features, preprocessing, validation strategies, and evaluation metrics lead to substantial differences in reported performance [2], [16], [21]. Consequently, model performance on one dataset cannot necessarily be generalized to other populations or healthcare settings. A further limitation is the emphasis on prediction rather than interpretation. Although metrics such as accuracy, F1-score, and ROC-AUC measure predictive performance, they do not explain the factors contributing to an individual prediction [2], [16], [21], [40]. XAI techniques such as SHAP and LIME address this limitation by providing global and local feature-level explanations [8], [9], [13], [14], [25]–[27]. However, technical explanations may still be difficult for non-technical users to understand. Thus, model explainability does not automatically result in patient-centred communication [25], [28]. Generative AI and LLMs provide an additional communication layer by transforming complex information into natural-language and conversational responses [5], [11], [29], [33], [34]. However, general-purpose LLMs do not inherently provide validated patient-specific diabetes risk prediction or explicit model explanations. Therefore, personalized diabetes risk awareness requires coordination between predictive models, XAI, and generative communication rather than relying on the LLM alone. Personalization remains another important gap. Diabetes prediction systems generally focus on estimating risk, whereas conversational healthcare systems commonly provide general information and support [16], [19]–[22], [34], [36]. Integrating an individual's predicted risk and contributing factors with understandable, personalized health information could provide a more connected risk-awareness process. RAG further supports this architecture by retrieving relevant information from external knowledge sources before response generation [7]. Although this can support evidence-grounded communication, RAG primarily addresses knowledge retrieval rather than risk prediction or model explainability. Its effectiveness also depends on retrieval quality, source reliability, factual consistency, and response validation [6], [11], [12]. Therefore, evidence grounding requires additional safety and validation mechanisms.

The literature also indicates that evaluation remains fragmented. Predictive systems are generally evaluated using statistical metrics, XAI using interpretability measures, and generative systems using response quality or factuality. An integrated healthcare assistant requires a broader evaluation perspective covering predictive performance, explainability, factuality, personalization, relevance, usability, and safety. The proposed research therefore positions these dimensions as complementary evaluation requirements. Finally, the system should remain clearly positioned as a diabetes risk-awareness and personalized health-information tool, rather than a replacement for laboratory diagnosis or professional clinical judgement. This distinction is important when translating probabilistic predictions and model explanations into patient-oriented communication. The literature therefore reveals a functional separation among AI technologies:



Based on the synthesis, five interconnected research gaps are identified:

Integration Gap: ML prediction, XAI, GenAI/LLMs, and RAG are commonly studied as separate components, with limited integration into a unified diabetes risk-awareness framework.

Personalization Gap: Existing systems often provide either risk prediction or general health information, with limited integration of patient-specific risk, contributing factors, and personalized communication.


Explainability-to-Communication Gap: XAI can explain model behaviour, but technical feature contributions are not automatically transformed into understandable patient-oriented information.

Evidence-Grounding and Reliability Gap: RAG supports access to external evidence, but retrieval quality, source reliability, factual consistency, and response safety require systematic evaluation.

Multidimensional Evaluation Gap: Existing studies frequently evaluate individual components independently. Integrated systems require evaluation across predictive performance, explainability, factuality, personalization, relevance, usability, and safety.

Conceptual Framework for Integrated Diabetes Risk Awareness

The critical synthesis indicates a gradual shift toward integrated AI systems for diabetes risk assessment and personalized health communication, in which Machine Learning (ML) provides risk estimation, Explainable AI (XAI) identifies the factors contributing to predictions, Retrieval-Augmented Generation (RAG) provides evidence grounding, and Generative AI/ Large Language Models (GenAI/LLMs) transform structured information into understandable natural-language communication.



The framework begins with ML-based prediction, where relevant demographic, clinical, physiological, and lifestyle characteristics are processed to estimate an individual's potential diabetes risk, followed by XAI, using techniques such as SHAP and LIME to identify contributing features at global and individual levels [8], [9], [25], [26]. The predicted risk and explanatory outputs are then converted into contextual information that can be understood by non-technical users, while maintaining the

distinction between risk estimation and diagnosis [25], [28]. RAG retrieves relevant information from trusted external knowledge sources to support evidence-grounded responses, although the quality and reliability of retrieved information remain important considerations [6], [7], [11], [12]. The retrieved evidence and structured risk information are subsequently processed by the GenAI/LLM layer to generate natural-language responses, with the LLM serving primarily as a communication component rather than independently determining diabetes risk [5], [11], [29], [33]. The system further supports personalized communication by combining an individual's risk estimate, contributing factors, and relevant evidence to provide context-sensitive health information [29], [34], [36]. Finally, safety validation is applied to identify factual inconsistencies, unsupported claims, inappropriate recommendations, and other potential safety concerns, since neither RAG nor LLM generation independently guarantees accurate or clinically appropriate responses [6], [7], [11], [29], [33]. Overall, the framework connects ML Prediction → XAI → Risk Interpretation → RAG → GenAI/LLM → Personalized Communication → Safety Validation and is proposed as a conceptual research direction for diabetes risk awareness and personalized health information rather than as a replacement for laboratory diagnosis or professional clinical judgement.

S The conceptual framework consequently represents a layered progression from data-driven prediction to evidence-grounded personalized communication. Each component performs a distinct function:

Component	Primary Function
ML Prediction	Estimates diabetes risk
XAI	Explains contributing factors
Risk Interpretation	Converts model outputs into understandable risk information
RAG	Retrieves relevant external evidence
GenAI/LLM	Generates natural-language communication
Personalization	Adapts information to individual context
Safety Validation	Check's reliability and appropriate communication

The framework should be interpreted as a conceptual synthesis of the reviewed literature rather than as evidence that all components have already been validated together in a single clinical setting. Its primary value is to provide a research direction for integrating complementary AI capabilities while maintaining appropriate boundaries between prediction, explanation, information retrieval, and communication.

Future Research Directions

The reviewed literature identifies several directions for future research in AI-enabled diabetes risk assessment and personalized healthcare communication. First, future studies should investigate the integrated implementation and validation of ML, XAI, RAG, and GenAI, with attention to interoperability, component-level contributions, and end-to-end system performance.

Second, external validation and generalizability should receive greater attention. Variations in datasets, populations, predictors, preprocessing, and validation strategies can substantially affect reported performance [2], [16], [21], [40]. Future research should therefore evaluate integrated systems using independent datasets and diverse demographic and healthcare settings.

Third, future work should strengthen patient-centred explainability and personalization. Although SHAP and LIME provide useful model-level and individual-level explanations [8], [9], [26], further research is required to convert technical explanations into understandable, context-sensitive health information. User-centred studies can help assess whether explanations are understandable and relevant to different users [25], [28].

Fourth, greater attention is required to RAG quality and knowledge governance. Future systems should examine source selection, updating, retrieval relevance, conflicting or outdated information, and

retrieval failures [7]. Evidence grounding should be combined with response validation because retrieval alone cannot guarantee factual or clinically appropriate outputs [6], [11], [12], [29], [33].

Fifth, standardized multidimensional evaluation frameworks should be developed for integrated healthcare assistants. Evaluation should extend beyond predictive performance or language quality to include explainability, factual consistency, evidence relevance, personalization, usability, user satisfaction, and safety.

Finally, future research should address privacy, human-AI interaction, and clinical boundaries. Personalized systems may require sensitive health information and therefore need appropriate privacy-preserving mechanisms, consent, and secure data processing. User studies involving relevant stakeholders can further assess understanding and acceptance. At the same time, AI-generated risk information should remain positioned as health awareness and information support rather than a replacement for established screening, diagnosis, or professional clinical judgement.

Overall, future research should move toward validated, explainable, evidence-grounded, personalized, privacy-aware, and safety-conscious AI healthcare systems.

Conclusion

Artificial Intelligence has expanded diabetes risk assessment from conventional statistical approaches toward machine learning prediction, explainable modelling, generative communication, and evidence-grounded conversational systems. ML provides risk estimation, XAI improves understanding of contributing factors, GenAI/LLMs support natural-language communication, and RAG provides access to relevant external knowledge [5], [7], [11], [29], [33], [34]. The critical synthesis shows that these capabilities are often investigated separately. Diabetes prediction studies commonly emphasize predictive performance, XAI focuses on model interpretation, while GenAI and RAG research primarily addresses communication and evidence grounding. This fragmentation highlights the need to consider prediction, explanation, evidence, personalization, and safety as interconnected components of diabetes risk-awareness systems. The conceptual framework presented in this review therefore connects ML Prediction → XAI → Risk Interpretation → RAG → GenAI/LLM → Personalized Communication → Safety Validation. The framework is intended as a research direction for connecting these complementary capabilities rather than as evidence of established clinical effectiveness. Future empirical research is required to validate the integrated approach across diverse populations and independent datasets, evaluate generated responses for factuality and safety, assess user understanding and acceptance, and address privacy and responsible healthcare communication.

In conclusion, integrating predictive AI, explainability, evidence retrieval, and generative communication provides a structured direction for developing personalized diabetes risk-awareness systems. Such systems should remain clearly distinguished from autonomous diagnosis and professional clinical decision-making.

References:

- [1] Anjana, R. M., et al. (2023). Metabolic non-communicable disease health report of India: The ICMR-INDIAB national cross-sectional study (ICMR-INDIAB-17). *The Lancet Diabetes & Endocrinology*, 11(7), 474–489. [https://doi.org/10.1016/S2213-8587\(23\)00119-5](https://doi.org/10.1016/S2213-8587(23)00119-5)
- [2] Olusanya, M. O., Ogunsakin, R. E., Ghai, M., & Adeleke, M. A. (2022). Accuracy of Machine Learning Classification Models for the Prediction of Type 2 Diabetes Mellitus: A Systematic Survey and Meta-Analysis Approach. *International Journal of Environmental Research and Public Health*, 19(21), 14280. <https://doi.org/10.3390/ijerph192114280>
- [3] Mohsen, F., et al. (2023). A scoping review of artificial intelligence-based methods for diabetes risk prediction. *npj Digital Medicine*, 6, 197. <https://doi.org/10.1038/s41746-023-00933-5>
- [4] Thirunavukarasu, A. J., Ting, D. S. J., Elangovan, K., et al. (2023). Large language models in medicine. *Nature Medicine*, 29, 1930–1940. <https://doi.org/10.1038/s41591-023-02448-8>
- [5] World Health Organization. (2021). *Ethics and Governance of Artificial Intelligence for Health*. WHO.

- [6] Lewis, P., et al. (2020). Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. *Advances in Neural Information Processing Systems*, 33, 9459–9474.
- [7] Lundberg, S. M., & Lee, S.-I. (2017). A Unified Approach to Interpreting Model Predictions. *Advances in Neural Information Processing Systems*, 30.
- [8] Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). “Why Should I Trust You?” Explaining the Predictions of Any Classifier. *Proceedings of the 22nd ACM SIGKDD*, 1135–1144. <https://doi.org/10.1145/2939672.2939778>
- [9] Vaswani, A., et al. (2017). Attention Is All You Need. *Advances in Neural Information Processing Systems*, 30.
- [10] Singhal, K., et al. (2023). Large Language Models Encode Clinical Knowledge. *Nature*, 620, 172–180. <https://doi.org/10.1038/s41586-023-06291-2>
- [11] Nori, H., et al. (2023). Capabilities of GPT-4 on Medical Challenge Problems. *arXiv:2303.13375*.
- [12] Tasin, I., Nabil, T. U., Islam, S., & Khan, R. (2023). Diabetes prediction using machine learning and explainable AI techniques. *Healthcare Technology Letters*, 10(1–2), 1–10. <https://doi.org/10.1049/htl2.12039>
- [13] Joseph, L. P., Joseph, E. A., & Prasad, R. (2022). Explainable diabetes classification using hybrid Bayesian-optimized TabNet architecture. *Computers in Biology and Medicine*, 151, 106178. <https://doi.org/10.1016/j.compbimed.2022.106178>
- [14] Contreras, I., & Vehi, J. (2018). Artificial intelligence for diabetes management and decision support: Literature review. *Journal of Medical Internet Research*, 20(5), e10775.
- [15] Fregoso-Aparicio, L., Noguez, J., Montesinos, L., & García-García, J. A. (2021). Machine learning and deep learning predictive models for type 2 diabetes: a systematic review. *Diabetology & Metabolic Syndrome*, 13, 148. <https://doi.org/10.1186/s13098-021-00767-9>
- [16] Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: Synthetic Minority Over-sampling Technique. *Journal of Artificial Intelligence Research*, 16, 321–357. <https://doi.org/10.1613/jair.953>
- [17] Sisodia, D., & Sisodia, D. S. (2018). Prediction of diabetes using classification algorithms. *Procedia Computer Science*, 132, 1578–1585. <https://doi.org/10.1016/j.procs.2018.05.122>
- [18] Zou, Q., Qu, K., Luo, Y., Yin, D., Ju, Y., & Tang, H. (2018). Predicting diabetes mellitus with machine learning techniques. *Frontiers in Genetics*, 9, 515. <https://doi.org/10.3389/fgene.2018.00515>
- [19] Wang, L., Wang, X., Chen, A., Jin, X., & Che, H. (2020). Prediction of type 2 diabetes risk and its effect evaluation based on the XGBoost model. *Healthcare*, 8(3), 247. <https://doi.org/10.3390/healthcare8030247>
- [20] De Silva, K., Lee, W. K., Forbes, A., Demmer, R. T., Barton, C., & Enticott, J. (2020). Use and performance of machine learning models for type 2 diabetes prediction in community settings: A systematic review and meta-analysis. *International Journal of Medical Informatics*, 143, 104268. <https://doi.org/10.1016/j.ijmedinf.2020.104268>
- [21] Talaei-Khoei, A., & Wilson, J. M. (2018). Identifying people at risk of developing type 2 diabetes: A comparison of predictive analytics techniques and predictor variables. *International Journal of Medical Informatics*, 119, 22–38. <https://doi.org/10.1016/j.ijmedinf.2018.08.008>
- [22] Alabi, et al. (2023). Machine learning explainability in nasopharyngeal cancer survival using LIME and SHAP. *Scientific Reports*, 13, 8984. <https://doi.org/10.1038/s41598-023-35795-0>
- [23] Abnoosian, K., Farnoosh, R., & Behzadi, M. H. (2023). Prediction of diabetes disease using an ensemble of machine learning multi-classifier models. *BMC Bioinformatics*, 24, 337. <https://doi.org/10.1186/s12859-023-05465-z>
- [24] Holzinger, A., Biemann, C., Pattichis, C. S., & Kell, D. B. (2017). What do we need to build explainable AI systems for the medical domain? *arXiv:1712.09923*.
- [25] Lundberg, S. M., et al. (2020). From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence*, 2, 56–67. <https://doi.org/10.1038/s42256-019-0138-9>

- [26] Kibria, H. B., Nahiduzzaman, M., Goni, M. O. F., Ahsan, M., & Haider, J. (2022). An ensemble approach for the prediction of diabetes mellitus using a soft voting classifier with an explainable AI. *Sensors*, 22(19), 7268. <https://doi.org/10.3390/s22197268>
- [27] Guidotti, R., et al. (2018). A survey of methods for explaining black box models. *ACM Computing Surveys*, 51(5), Article 93. <https://doi.org/10.1145/3236009>
- [28] Sallam, M. (2023). The utility of ChatGPT as an example of large language models in healthcare education, research and practice. *Healthcare*.
- [29] Kung, K. M., et al. (2023). Performance of ChatGPT on USMLE: Potential for AI-assisted medical education using large language models. *PLOS Digital Health*, 2(2), e0000198. <https://doi.org/10.1371/journal.pdig.0000198>
- [30] Ayers, J. B., et al. (2023). Comparing Physician and Artificial Intelligence Chatbot Responses to Patient Questions Posted to a Public Social Media Forum. *JAMA Internal Medicine*, 183(6), 589–596. <https://doi.org/10.1001/jamainternmed.2023.1838>
- [31] Gilson, A., et al. (2023). How Does ChatGPT Perform on the United States Medical Licensing Examination (USMLE)? The Implications of Large Language Models for Medical Education and Knowledge Assessment. *JMIR Medical Education*, 9, e45312. <https://doi.org/10.2196/45312>
- [32] Moor, M., et al. (2023). Foundation models for generalist medical artificial intelligence. *Nature*, 616, 259–265. <https://doi.org/10.1038/s41586-023-05881-4>
- [33] [34] Laranjo, L., et al. (2018). Conversational agents in healthcare: a systematic review. *Journal of the American Medical Informatics Association*, 25(9), 1248–1258. <https://doi.org/10.1093/jamia/ocy072>
- [34] Fitzpatrick, K., Darcy, A., & Vierhile, M. (2017). Delivering Cognitive Behavior Therapy to Young Adults With Symptoms of Depression and Anxiety Using a Fully Automated Conversational Agent (Woebot): A Randomized Controlled Trial. *JMIR Mental Health*, 4(2), e19. <https://doi.org/10.2196/mental.7785>
- [35] Abd-Alrazaq, A. B., et al. (2021). An overview of the features of chatbots in healthcare: A systematic literature review. *Journal of Medical Internet Research*, 23(12), e42088. <https://doi.org/10.2196/42088>
- [36] Abd-Alrazaq, A. B., et al. (2020). The effectiveness and safety of using chatbots to improve mental health: Systematic review and meta-analysis. *Journal of Medical Internet Research*.
- [37] [38] Martinsson, J., Schliep, A., Eliasson, B., & Mogren, O. (2020). Blood Glucose Prediction with Variance Estimation Using Recurrent Neural Networks. *Journal of Healthcare Informatics Research*, 4(1), 1–18. <https://doi.org/10.1007/s41666-019-00059-y>
- [38] Valderrábano, R. J., & Linares, M. I. (2018). Diabetes mellitus and bone health: epidemiology, etiology and implications for fracture risk stratification. *Clinical Diabetes and Endocrinology*, 4, 9. <https://doi.org/10.1186/s40842-018-0060-9>
- [39] Ahamed, B. S., Arya, M. S., & Nancy, A. O. V. (2022). Diabetes Mellitus Disease Prediction Using Machine Learning Classifiers with Oversampling and Feature Augmentation. *Advances in Human-Computer Interaction*, 2022, 1–14. <https://doi.org/10.1155/2022/9220560>
- [40] Saxena, R., et al. (2023). Machine learning algorithms for diabetes detection: a comparative evaluation of performance of algorithms. *Evolutionary Intelligence*, 16, 587–603. <https://doi.org/10.1007/s12065-021-00685-9>
- [41] Reza, M. M., et al. (2023). Improving SVM performance for type II diabetes prediction with an improved non-linear kernel: Insights from the PIMA dataset. *Computer Methods and Programs in Biomedicine Update*, 4, 100118. <https://doi.org/10.1016/j.cmpbup.2023.100118>
- [42] Ganie, S. M., et al. (2023). An ensemble learning approach for diabetes prediction using boosting techniques. *Frontiers in Genetics*, 14, 1252159. <https://doi.org/10.3389/fgene.2023.1252159>
- [43] Tripathi, et al. (2023). Soft voting ensemble approach for diabetes prediction using the PIMA dataset.