Bilkent University

Department of Computer Engineering

Senior Design Project

DiabetAid

Project Specifications Report

Çağatay Akpınar (22003508), Sude Topaloğlu (22102677), Osman Berkay Artaç (22103071), Yamaç Yiğit Ozan (22003595), Begüm Kunaç (22103838)

Supervisor: Halil Altay Güvenir

Innovation Expert: Deniz Katırcıoğlu Öztürk

Nov 21, 2024

This report is submitted to the Department of Computer Engineering of Bilkent University in partial fulfillment of
the requirements of the Senior Design Project course CS491/1.

1. Introduction

Diabetes mellitus is a chronic disease that occurs when the pancreas produces insufficient insulin or when the body cannot effectively use the insulin it produces. Insulin is a vital hormone responsible for regulating blood sugar levels. Uncontrolled diabetes leads to elevated blood sugar, or hyperglycemia, which can cause long-term damage to various organs, including the nerves, blood vessels, and kidneys.

Diabetes is a growing health concern globally. In 2022, 14% of adults aged 18 and older were living with diabetes, and more than 59% of adults aged 30 and older with diabetes were not receiving medication. The disease remains a leading cause of death, directly causing 1.6 million deaths in 2021 alone, with almost half occurring before the age of 70 years. Additionally, diabetes is responsible for 11% of deaths due to cardiovascular disease and 530000 deaths due to kidney diseases [1].

Our project, DiabetAid, is being developed as an innovative solution to improve diabetes care through machine learning (ML) techniques. This project, in collaboration with Ankara Bilkent City Hospital, aims to provide personalized recommendations for medication and dosage to support doctors in making accurate treatment decisions. By analyzing patient data, such as age, gender, BMI, laboratory results, current medical conditions, and current treatments, DiabetAid uses ML algorithms to deliver precise, real-time suggestions during patient visits. DiabetAid is a user-friendly application that integrates seamlessly with hospital systems and supports doctors in making accurate treatment decisions more efficiently. This accelerates the treatment stage, reduces the workload on healthcare professionals, and enhances the overall quality of diabetes management.

1.1 Description

DiabetAid represents a very current undertaking in the field of AI in which machine learning methods are used to provide medication and dosage recommendations in a personalized and data-driven manner to transform how diabetes is treated. In development in conjunction with Ankara Bilkent City Hospital, the goal is to embed ML-driven decisions into clinical applications that will better enable health professionals to make appropriate, evidence-based therapeutic decisions. The approval by the ethics committee of the hospital implies satisfying all ethical and regulatory considerations.

The DiabetAid program involves collecting and processing, in a secure manner, patients' data of 6,500 subjects accessed through the hospital's VPN. This is then further compiled into Excel sheets so it can be subsequently processed and analyzed. It includes all the relevant information related to a patient, such as laboratory test reports, age, sex, BMI, other coexisting chronic illnesses, and current therapeutic regimes. Comorbidities, the details of treatments, among other information, are then obtained from the physician's notes within the hospital's information system to ensure completeness and weight in the dataset for model training. These elements of data are very essential in training and evaluating various machine learning models that will ensure the system provides accurate, practical, and clinically meaningful recommendations.

It applies various machine learning algorithms in order to identify the best model that provides recommendations of treatments. The cyclic approach ensures that the model identified will be robust and reliable for the model to make personalized recommendations on how diabetes can be managed. The team will document the entire process from data preparation, model training, evaluation, and results into a technical paper once the project is complete.

DiabetAid was adapted for use at Ankara Bilkent City Hospital and will be installed on computers used by the medical staff. It works with an interactive interface through which health professionals can insert patient data and receive, in real time, suggestions about medication and dosages. Of course, these suggestions are made with a rationale behind each recommendation. Thus, with this approach, the system is supportive, helping the physician through the treatment, without taking his place.

Prospectively, this is because DiabetAid can be implemented in other hospitals, associations, and institutions, subject to obtaining necessary permissions and approvals. Therefore, this scalability of the initiative means that the project is able to expand into households other than those of Ankara Bilkent City Hospital, allowing the target population to be expanded with increased cooperation on diabetes management. DiabetAid represents a huge step into providing personalized and optimal treatment of diabetes using a patient's data, such as laboratory test results, comorbid diseases, treatments, age, gender, and body mass.

1.2 High Level System Architecture & Components of Proposed Solution

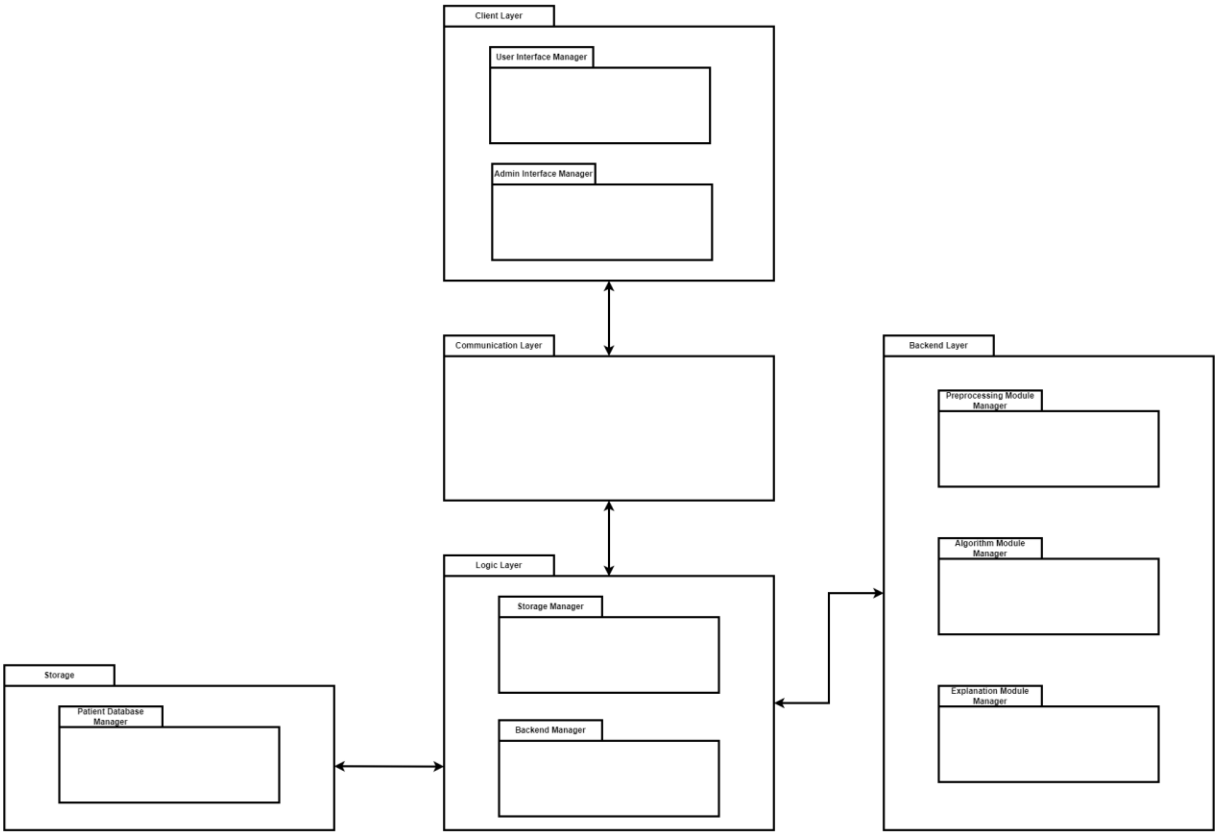

Explanation of Each Layer

- **Client Layer**: The Client Layer is the interface used by the people who interact with the system, such as administrators and doctors. Inside it, the User Interface Manager stores aids for doctors in inputting information for patients, visualizing treatment recommendations and their respective explanations, and the Admin Interface Manager controls access to data in order to limit exposure to sensitive patient information. This layer represents the interface through which users will interact with the system, transmitting user inputs to the backend and displaying results to the users.
- **Communication Layer**: The Communication Layer is an abstraction layer between the Client Layer and the Backend Layer, enabling the secure and efficient transfer of data by handling requests created within the Client Layer for processing at the back-end, returning responses to the Client Layer with the rationale and suggested treatments.

- **Backend Layer**: The backend layer controls the central processing of the system. The Preprocessing Module Manager prepares and cleans the patient data, whereas the Algorithm Module Manager generates the treatment recommendation through either machine learning or rule-based algorithms. In addition, it provides understandable explanations for the suggested treatments given by the Explanation Module Manager. It acts as the main processing center for the system.
- **Logic Layer**: The Logic Layer regulates system resources to support the back-end operations. The Backend Manager controls the flow of data across the Backend and Storage Layers, while the Storage Manager executes the operations of retrieval and update of data at the storage. It ensures that all the data necessary to carry out the processing and storage is managed effectively.
- **Storage Layer**: Data durability is the crucial role of storing and retrieving the treatment history of the patient, which the Storage Layer performs. The Patient Database Manager enables secured and efficient access to all data related to the patients. The interaction between this layer and the Logic Layer allows the back-end operation and maintenance of the system logs.

Explanation of Connections:

- **Client Layer - Communication Layer**: The Client Layer communicates with the Communication Layer to send user inputs, such as information related to patients, and it also receives the output, like clarifications and therapeutic recommendations. This, therefore, allows for the correct flow of information from the user interface to the back-end system.
- **Communication Layer - Logic Layer**: The Communication Layer also establishes the channel where the Client Layer can send the user requests for processing to the Logic Layer. The latter, after making the necessary calculations, utilizes the Communication Layer in order to send the results back to the Client Layer.
- **Logic Layer - Storage**: The Logic Layer interacts with the Storage Layer for any decision on when to persist the data, fetches patient records and treatment history from the storage, securely saves changed or new data in storage.
- **Logic Layer - Backend Layer**: It carries out various computations using preprocessed data for certain operations; for example, the execution of algorithms and/or fetching of required data in support of the Backend Layer.

Without this layer, the Backend Layer would not be able to access any system resources to generate results.

1.3 Constraints

1.3.1 Implementation Constraints

- The project will be implemented as a web solution for the doctors in personal computers running on web at the Ankara Bilkent City Hospital and then adapted for use in other hospitals.
- GitHub and Git will be used for version control and collaboration during development.
- Jira will be used for project management and tracking development tasks to ensure smooth progress.
- Python will be the main programming language due to its strong libraries for machine learning and data processing.
- TensorFlow, Scikit-learn and LIME will be used for training, testing and explaining machine learning models.
- Patient data will be retrieved via the hospital's VPN and temporarily stored in structured formats like CSV or Excel for processing.
- The interface will be built using React and Flask to ensure ease of use for hospital staff.

1.3.2 Economic Constraints

- This will be costlier if the system were to be scaled up to other hospitals or institutions, considering infrastructural adaptation, permission that will have to be taken, probably retraining the model on new data.
- The system uses open-source libraries and tools, which eliminates licensing costs.
- Data storage and VPN access are provided by the hospital at no additional cost, reducing the overall expenses.
- There are no fees associated with using the hospital's existing IT infrastructure, as it is already maintained by the hospital's team.

1.3.3. Ethical Constraints

- All the information related to patients is anonymized. No personal data, including names or IDs, is either kept or processed in any form that would lead to an individual being identified. The project is in conformance with KVKK law to protect the security of information.
- The project has approval from the hospital's ethics committee to use patient data. If the system is expanded to other hospitals, new permissions will be required.
- Patient data is only used for training the system and providing treatment recommendations. Data is deleted immediately after processing to avoid unnecessary storage.
- The system will clearly explain how it generates its recommendations so that doctors can understand the logic and trust the outputs.
- The system will be a tool to assist doctors, not replace their expertise. Doctors are always responsible for the final treatment decisions.

1.4 Professional and Ethical Issues

1.4.1 Data Privacy and Security

All patient data is anonymized and treated securely for confidentiality and in accordance with KVKK. Data is used only for the training and testing of systems and no other purposes, in accordance with regulatory rules, which minimizes the risk of data misuse. Permission to access the patients' data on this project was also provided by the ethics committee of the hospital, ensuring that all procedures meet ethical guidelines.

1.4.2 Transparency and Accountability

Doctors actively participate in validating system recommendations. A collaborative approach ensures clinical safety, reliability, and trust in AI-driven decisions. Regular audits help find out and address biases and inaccuracies within the system through feedback.

1.4.3 Collaboration with Third Parties

Expanding to other institutions requires new ethical approvals. Clear agreements will outline data usage, ownership, and responsibilities to avoid conflicts.

1.4.4 Licensing and Use of Open-Source Software

The project uses open-source tools in line with licensing requirements, ensuring compliance and sustainability.

1.4.5 Patient Consent and Autonomy

The system respects patient autonomy by only providing recommendations that are non-binding, thereby allowing doctors to retain full control over the clinical decisions. Patients are informed about the use of AI in their treatment process and the measures taken to protect their data.

1.5 Standards

In the development of DiabetAid, several international and local standards are considered to ensure safety, reliability, and compliance, particularly given the sensitive nature of medical data in Turkey. Below are the key standards considered:

KVKK (Personal Data Protection Law)

- KVKK is the Turkish version of the EU GDPR for the protection of personal data. All the patient data utilized in DiabetAid is anonymized and undergoes secure processing to fully comply with KVKK recommendations. This includes taking prior permission from the hospital's ethics committee for the use of patient data with the view of protecting their privacy and maintaining confidentiality [2].

ISO/IEC 27001 - Information Security Management System

- The ISO/IEC 27001 standard ensures that all patient information handled by DiabetAid is well kept and processed with security in mind. The security management of health data shall always have this standard to minimize the risks associated with data breach and unauthorized access [3].

ISO 9241 - Usability and Human-Centered Design

- It should be user-friendly for any and all medical practitioners irrespective of their technical background. ISO 9241 presents a standard of usability that gives rules for designing user interfaces with the intention of making DiabetAid as intuitive and effective as possible for care providers [4].

ISO 14971 - Risk Management for Medical Devices

- Since DiabetAid is a treatment recommender, there will need to be some regulation of risk and its mitigation regarding its use. ISO 14971 has been applied for the identification, assessment, and control of risks associated with patient safety so that the system can advance healthcare with safety for the patients [5].

ISO 62304 - Software Life Cycle Processes for Medical Devices

- Development of software for DiabetAid is based on the accepted standard ISO 62304 for guidance of design, development, and maintenance of software for medical devices. This standard provides assurance that our approach is well-organized and managed with substantial emphasis on software quality, safety, and reliability [6].

IEEE 11073 - Health Informatics

- DiabetAid uses the IEEE 11073 standard for the integration of the application with medical devices and health informatics systems. Adherence to this standard will ensure the smooth functionality of the application with the available hospital systems hence ensuring successful capture and sharing of information [7].

HL7 (Health Level Seven)

- This will work with most EHR systems available today since it uses HL7 standards, which are common in hospitals and facilitate the interaction between different applications through the interface by providing smooth connections and sharing of data. This enables DiabetAid to communicate efficiently with other hospital information systems for better output in patient care [8].

ISO 13485 - Medical Device Quality Management Systems

- ○ DiabetAid has been developing and continuously maintaining following the norms of ISO 13485 for quality management in the development of medical software solutions. This helps the product maintain high-quality standards throughout the life cycle of the product [9].

Open-Source Software Compliance

- ○ DiabetAid uses different open-source tools like TensorFlow and Scikit-learn. The project follows all the licensing conditions and hence is with the policy of open-source software, which avoids all the possible legal issues concerning intellectual property rights accordingly [10].

SSL/TLS Encryption

- ○ Communication between DiabetAid, hospital servers, and doctor devices is secured using **SSL/TLS** encryption protocols. This ensures data integrity and confidentiality during transmission [11].

Together, these standards ensure that DiabetAid is built to be a reliable, secure, and user-friendly healthcare tool. By following these guidelines, we aim to enhance diabetes treatment while keeping patient safety and data privacy at the highest level.

2. Design Requirements

2.1. Functional Requirements

- The doctor must be able to securely log in to the system using hospital-provided credentials.
- The doctor should be able to input patient data such as demographics, lab results, and medical history into the application.
- The doctor should be able to receive personalized treatment recommendations, including medication and dosage, based on the patient’s data.
- The doctor should be able to view and update treatment history for each patient through the system.
- The doctor should be able to access support or help documentation through the application interface.

- The doctor should be able to generate and download a detailed report summarizing patient data, recommendations, and treatment history.
- The doctor should be able to provide feedback on the system's recommendations to improve machine learning model performance over time.
- The system should analyze patient data using machine learning models to generate personalized treatment recommendations, including medication and dosage.
- The system must store treatment logs and recommendations securely to allow for future reference and follow-up.

2.2. Non-Functional Requirements

2.2.1 Usability

- The application should be easy to use, with a simple and clean design to help doctors of all technical levels.
- Clear instructions, tooltips, and error messages should guide users and make the system easy to navigate.
- Data entry should be quick and efficient so that it doesn't slow down doctors during patient care..

2.2.2 Reliability

- The system must work well even if some patient data is missing, by providing alerts or making reasonable assumptions.
- It should be available and ready to use at all times since it supports important medical decisions.
- If the VPN connection fails, the system should save any entered or processed data and try to reconnect automatically.
- The recommendations must be accurate, using well-tested machine learning models to ensure reliability.

2.2.3 Performance

- The system should provide recommendations quickly after the doctor enters patient data, ensuring it does not slow down the workflow.
- The system must handle large patient datasets quickly and efficiently.

2.2.4 Supportability

- Updates should be simple to apply without interrupting hospital operations.
- The system should connect smoothly with hospital systems like electronic health records (EHRs) or lab reports.

2.2.5 Scalability

- The system should handle more users, patient data, and hospitals as it grows without performance issues.
- The system should support multiple doctors using it simultaneously without performance issues or slowing down.

3. Feasibility Discussions

3.1. Market & Competitive Analysis

There are various different kinds of tools for managing diabetes in the healthcare market today, ranging from glucose monitoring devices to decision-support systems and EHR integrations. Most of these tools have been designed either just to monitor patient data or to provide general guidelines for treatment, without offering detailed personalized treatment options.

DiabetAid stands out in this competitive landscape because of its novel coupling of explainability and efficiency. Many AI healthcare tools function as "black boxes" offering treatment suggestions without explaining how they were derived, however DiabetAid ensures transparency by incorporating Explainable AI. This allows doctors to receive personalized medication and dosage suggestions, and understand the reasoning behind these suggestions. This explainability is critical for building trust, supporting informed decision-making in clinical practice, and ensuring regulatory compliance. DiabetAid, through automation and optimization of the processes of medication and dosage selection, substantially reduces the time doctors spend performing manual calculations and making complex decisions. In this way, it can allow health professionals more time for patient care and other critical tasks. Traditional diabetes management often involves a trial-and-error approach to finding the optimal treatment plan. DiabetAid eliminates much of this uncertainty by providing data-driven, evidence-based recommendations tailored to individual patients, improving patient outcomes while reducing inefficiencies.

The application is being developed in collaboration with Bilkent City Hospital. The scalable design enables the application to be adaptable for various healthcare institutions, allowing for potential adoption across a broader market. The system will be designed to integrate into hospital systems and EHRs, offering doctors a real-time and user-friendly tool to manage diabetes accurately and efficiently. Early user feedback from this collaboration underlines the tool's potential for enhancing workflow efficiency and improving outcomes of treatments.

3.2. Academic Analysis

DiabetAid utilizes machine learning methodologies by training several models and testing them through empirical metrics such as accuracy, precision, recall, and F1 score. This whole process of iteration is well documented and considered a great contribution to the academic research of AI applications in medicine, especially personalized treatments.

One of the great points with respect to this project is that it deals with Explainable AI-XAI. That would mean the system provides understandable and clear explanations with regard to the treatment proposals provided. It could, therefore, enable doctors to review and validate model suggestions with a view to ensuring trust and confidence in AI-driven decisions. By embedding XAI into DiabetAid, embedding transparency into DiabetAid strengthens clinically informed decision-making-easier translation at the interface between machine output and human expertise.

Various academic papers are being examined as part of the research for DiabetAid. For instance, the paper "Why Should I Trust You? Explaining the Predictions of Any Classifier" [12] introduces LIME (Local Interpretable Model-Agnostic Explanations), a method for explaining machine learning predictions in a transparent and interpretable manner. The principles in this paper are being incorporated to ensure that the treatment recommendations provided by DiabetAid are reliable and easy for doctors to understand.

The role of explainability in medical AI is explored in the paper "Explainability for Artificial Intelligence in Healthcare: A Multidisciplinary Perspective" [13] further. This paper highlights the ethical, legal, medical, and patient-centered implications of explainable AI. It also highlights how explainability is necessary for strengthening trust in clinical decision support systems (CDSS) and avoiding potential harms caused by black-box AI models. The insights from this research emphasize the importance of developing explainable features in DiabetAid to ensure alignment with ethical and professional standards in medicine.

Another relevant contribution comes from the paper "Artificial Intelligence and Black-Box Medical Decisions: Accuracy versus Explainability" [14], which addresses the trade-off between accuracy and the interpretability of ML models in medicine. It states that black-box models may cause ethical concerns because of their lack of transparency but empirical verification of results can sometimes outweigh the need for complete interpretability, particularly in cases where causal knowledge is incomplete. This perspective adds depth to the development of DiabetAid, reinforcing the importance of balancing accuracy with explainability in clinical applications.

This project, therefore, focuses on the research of different machine learning approaches, with particular emphasis on decision trees to optimize treatment strategies. The paper "Decision Trees: An Overview and Their Use in Medicine" [15] provides the necessary insights into using them in medical treatment decisions. Decision trees are mainly known for their simplicity and reliability; through them, suggestions on certain lines of treatment can be presented in a comprehensively simplified manner. As clarity and ease are important for the diffusion in clinical practice, the characteristic of decision trees fits with the objectives of DiabetAid.

The application of machine learning to other diseases is being analyzed to improve the understanding of optimizing treatment. For example, the paper "Learning Stroke Treatment Progression Models for an MDP Clinical Decision Support System" [16] explores a data-driven clinical decision support system (CDSS) for stroke treatment using Markov Decision Processes (MDPs).

Moreover, this collaboration gives the project more clinical relevance with hospital physicians and ensures that recommendations align with real practice, thus contributing to better decision-making. Therefore, it will be thoroughly described in an academic publication, presenting the methodology applied, the performance of the model, explainability features, and integration approaches. Hence, the paper will yield important insights for both researchers and practitioners and connect theoretical developments in artificial intelligence with practical applications in healthcare. DiabetAid constitutes a significant contribution to the existing literature on the impact of artificial intelligence within contemporary medicine and clinical practice.

4. Glossary

Accuracy: A metric used to measure how often the machine learning model's predictions are correct.

BMI (Body Mass Index): A numerical value derived from an individual's weight and height to assess whether they are underweight, normal weight, overweight, or obese.

Clinical Decision Support System (CDSS): A software system designed to assist healthcare professionals in making clinical decisions by analyzing patient data and providing evidence-based recommendations.

CSV (Comma-Separated Values): A file format that stores tabular data in plain text, with each row separated by a newline and each column separated by a comma.

EHR (Electronic Health Record): A digital version of a patient's paper chart, containing medical history, diagnoses, medications, and treatment plans.

Empirical Metrics: Quantitative measures, such as accuracy, precision, recall, and F1 score, used to evaluate the performance of machine learning models.

Explainable AI (XAI): A type of AI that provides clear explanations of how and why it arrives at specific outputs or decisions.

F1 Score: A measure of a model's accuracy, combining precision and recall into a single metric.

Hyperglycemia: A condition characterized by abnormally high blood sugar levels.

KVKK (Personal Data Protection Law): Turkish legislation that governs the protection and processing of personal data.

LIME (Local Interpretable Model-Agnostic Explanations): A tool used to explain predictions made by machine learning models.

Markov Decision Processes (MDPs): A mathematical framework used for modeling decision-making where outcomes are partly random and partly under the control of a decision-maker.

Precision: A metric that measures the accuracy of positive predictions made by a machine learning model.

React: React is a JavaScript library for building user interfaces, focusing on creating reusable components.

Flask: Flask is a lightweight Python web framework used to build web applications quickly and easily.

Recall: A metric that measures the proportion of actual positives correctly identified by the model.

Scikit-learn: A Python library for building and evaluating machine learning models.

TensorFlow: An open-source machine learning library used for building and training ML models.

VPN (Virtual Private Network): A secure connection used to access remote systems and protect data during transmission.

5. References

[1] "Diabetes," World Health Organization, [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/diabetes.

[2] "6698 Sayılı Kişisel Verilerin Korunması Kanunu (KVKK)," www.mevzuat.gov.tr/mevzuat?MevzuatNo=6698&MevzuatTur=1&MevzuatTertip=5. Accessed 21 Nov. 2024.

[3] "ISO/IEC 27001:2022," ISO, https://www.iso.org/standard/27001 (accessed Nov. 21, 2024).

[4] "ISO 9241-210:2019," ISO, https://www.iso.org/standard/77520.html (accessed Nov. 21, 2024).

[5] "ISO 14971:2019," ISO, https://www.iso.org/standard/72704.html (accessed Nov. 21, 2024).

[6] "IEC 62304:2006," ISO, https://www.iso.org/standard/38421.html (accessed Nov. 21, 2024).

[7] "IEEE SA - IEEE Standard for Health Informatics - device interoperability - part 10701: Point-of-care medical device communication - metric provisioning by participants in a service-oriented device connectivity (SDC) system," IEEE Standards Association, https://standards.ieee.org/ieee/11073-10701/7538/ (accessed Nov. 21, 2024).

[8] Health Level Seven international - homepage | HL7 international, https://www.hl7.org/ (accessed Nov. 21, 2024).

[9] "ISO 13485:2016," ISO, https://www.iso.org/standard/59752.html (accessed Nov. 21, 2024).

[10] J. Perez, "Unpacking open source compliance," OpenLogic by Perforce, https://www.openlogic.com/blog/open-source-compliance-overview (accessed Nov. 21, 2024).

[11] "What is SSL/TLS encryption?," F5, Inc., https://www.f5.com/glossary/ssl-tls-encryption (accessed Nov. 21, 2024).

[12] M. T. Ribeiro, S. Singh, and C. Guestrin, "Why Should I Trust You? Explaining the Predictions of Any Classifier," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, San Francisco, CA, USA, 2016, pp. 1135-1144. Available: http://dx.doi.org/10.1145/2939672.2939778.

[13] J. Amann, P. Blasimme, G. Vayena, J. Frey, and A. Madai, "Explainability for artificial intelligence in healthcare: a multidisciplinary perspective," *BMC Med Inform Decis Mak* , Jun. 2020. Available: https://doi.org/10.1186/s12911-020-01332-6.

[14] A. J. London, "Artificial Intelligence and Black-Box Medical Decisions: Accuracy versus Explainability," *The Hastings Center Report*, vol. 49, no. 1, pp. 15–21, Jan./Feb. 2019. Available: https://doi.org/10.1002/hast.973.

[15] V. Podgorelec, P. Kokol, B. Stiglic, and I. Rozman, "Decision Trees: An Overview and Their Use in Medicine," *Journal of Medical Systems*, vol. 26, no. 5, pp. 445–463, Oct. 2002.

[16] D. C. Coroian and K. Hauser, "Learning Stroke Treatment Progression Models for an MDP Clinical Decision Support System," *SIAM International Conference on Data Mining (SDM)*, pp. 676–684.