Tuijin Jishu/Journal of Propulsion Technology
ISSN: 1001-4055
Vol. 47 No. 3 (2026)

Data Driven AI Framework for Improving Healthcare Outcomes using Natural Language Processing and Machine Learning Techniques

[1] Shaik Noushin, [2] Shaik Jahangeer Basha

M.Tech Student, Department of CSE Prakasam Engineering College (Autonomous), Kandukur, India

Assistant Professor, Department of CSE Prakasam Engineering College (Autonomous), Kandukur, India

Abstract- The adoption of Artificial Intelligence (AI) in healthcare has created new opportunities for improving medical diagnosis, patient support, and treatment planning. This paper proposes a data-centric AI framework that integrates Natural Language Processing (NLP), Machine Learning, and Generative AI to deliver intelligent healthcare assistance. The framework analyzes patient-reported symptoms using TF-IDF feature extraction and a Logistic Regression classifier to identify potential diseases. Furthermore, a Generative AI component generates comprehensive medical information, including disease descriptions, risk assessments, and preventive recommendations. The proposed solution is implemented as a Flask-based web application featuring secure user authentication and conversation history management. Experimental evaluation demonstrates that the framework can accurately classify diseases from textual symptom descriptions while providing informative healthcare guidance. The system improves access to preliminary medical support, minimizes dependence on manual assessment, and enhances healthcare decision-making. Overall, the proposed framework contributes toward the development of intelligent, scalable, and user-friendly healthcare technologies.

Keywords— Artificial Intelligence, Healthcare, Natural Language Processing, Machine Learning, Disease Prediction, Generative AI, TF-IDF, Logistic Regression.

I. Introduction

The rapid advancement of Artificial Intelligence (AI) has transformed healthcare by enabling intelligent diagnosis and automated clinical decision-support systems. Conventional healthcare practices largely depend on manual evaluation and expert judgment, which can be time-intensive and susceptible to human error. To overcome these limitations, AI-driven technologies such as Natural Language Processing (NLP) and Machine Learning have been increasingly adopted for analyzing patient information and generating accurate medical insights. NLP enables healthcare systems to interpret and process symptoms described in natural language, while machine learning algorithms identify hidden patterns within data and predict potential diseases based on symptom inputs [1], [5], [7]. Healthcare chatbot systems utilize these technologies to communicate with patients, gather symptom-related information, and provide preliminary medical recommendations [8].

The proposed framework combines NLP techniques, Logistic Regression, and Generative AI to develop an intelligent healthcare assistant capable of predicting diseases based on user-reported symptoms. In addition to disease prediction, the system provides detailed explanations, evaluates risk levels, and suggests preventive measures. This integrated approach enhances healthcare accessibility, reduces response time, and supports early-stage medical decision-making.

Ii. Literature Review

Recent research highlights the significant impact of Artificial Intelligence in healthcare applications, particularly in disease prediction systems and intelligent patient interaction platforms. NLP methodologies have been extensively employed to extract meaningful information from medical documents, clinical records, and patient-generated text, thereby improving diagnostic accuracy and healthcare decision support [1], [6]. Machine learning algorithms such as Logistic Regression, Support Vector Machines, and Random Forest classifiers have demonstrated promising results in disease prediction and classification tasks [5], [18]–[20]. Advances in deep learning have introduced transformer-based architectures such as BERT, which improve contextual understanding of medical text and significantly enhance chatbot performance and conversational accuracy [10], [16], [17]. AI-powered healthcare chatbots have also been developed to automate patient interaction and preliminary diagnosis, improving healthcare accessibility and operational efficiency [8].

Despite these advancements, existing solutions continue to face challenges related to interpretability, privacy protection, and adaptability to dynamic healthcare environments. The proposed framework addresses these limitations by integrating NLP-based symptom analysis, machine learning-driven disease prediction, and Generative AI capabilities to deliver accurate, interactive, and explainable healthcare assistance.

Iii. Existing System

Current healthcare assistance platforms primarily depend on traditional diagnostic procedures, rule-based symptom assessment systems, and basic chatbot technologies. In most healthcare settings, disease diagnosis is performed manually by medical professionals based on patient symptoms and clinical expertise. Although this approach provides reliable outcomes, it requires significant human involvement, consumes considerable time, and may be affected by subjective interpretation.

To reduce manual workload, rule-based healthcare chatbots have been introduced that utilize predefined rules, keyword mappings, and decision-tree structures to generate medical suggestions. These systems operate using fixed logic and symptom patterns to associate user inputs with potential diseases. While such solutions offer a basic level of automation, they lack the capability to understand complex, ambiguous, or unstructured symptom descriptions commonly used by patients. Some healthcare applications incorporate keyword matching and elementary NLP techniques to improve interaction quality; however, these methods remain highly dependent on exact term matching and often fail to capture the contextual meaning of user symptoms. Consequently, they may generate incomplete or inaccurate predictions when symptoms are expressed using natural language variations.

Furthermore, many existing healthcare assistance systems do not incorporate advanced machine learning models, adaptive learning mechanisms, or personalized recommendation capabilities. They often fail to provide comprehensive disease explanations, risk evaluations, and preventive guidance, which are important factors in improving user confidence and healthcare decision-making. In addition, concerns related to data privacy, system scalability, and limited integration with modern AI technologies continue to restrict the effectiveness and widespread adoption of these solutions.

Table 1: Limitations of Existing Healthcare Systems

Aspect	**Existing Systems**
Diagnosis Method	Manual or rule-based
Input Processing	Keyword-based or limited NLP
Context Understanding	Not supported
Accuracy	Moderate
Automation Level	Limited
Personalization	Not available

Explainability	Minimal or absent
Scalability	Limited
Real-time Adaptation	Not supported

Iv. Proposed Methodology

The proposed framework introduces an intelligent AI-powered healthcare chatbot that combines Natural Language Processing (NLP), Machine Learning, and Generative AI technologies to deliver automated healthcare assistance. The system analyzes symptoms provided by users in natural language, predicts potential diseases through a trained classification model, and generates comprehensive medical information and recommendations. The architecture is designed using a structured processing pipeline to ensure high prediction accuracy, scalability, reliability, and interpretability.

A. System Workflow

The proposed system operates through the following sequential steps:

1. User inputs symptoms in natural language
2. Text preprocessing is applied
3. Feature extraction using TF-IDF
4. Disease prediction using Logistic Regression
5. AI-based response generation
6. Storage of interaction history

B. Step-wise Methodology

1) Data Input and Preprocessing

User-provided symptoms are processed using NLP techniques such as:

- Lowercasing
- Tokenization
- Stop-word removal

This ensures clean and standardized input for further analysis.

2) Feature Extraction using TF-IDF

The system converts textual symptom data into numerical vectors using Term Frequency–Inverse Document Frequency (TF-IDF).

$$TF\text{-}IDF(t,d) = TF(t,d) \times \log\left(\frac{N}{DF(t)}\right)$$

where:

- TF: Term frequency of term t in document d
- DF(t): Document frequency of term t
- N: Total number of documents

This representation captures the importance of symptoms in the dataset.

3) Disease Prediction using Logistic Regression

The extracted features are passed to a Logistic Regression classifier to predict the disease.

$$P(y = 1|x) = \frac{1}{1+e^{-(w^T x+b)}}$$

where:

- x: Feature vector
- w: Weight vector
- b: Bias term

The model outputs the probability of each disease class and selects the most likely prediction.

4) AI-Based Medical Response Generation

After prediction, a Generative AI model produces:

- Explanation of the disease
- Risk level
- Preventive precautions
- Lifestyle recommendations

This enhances user understanding and interaction.

5) Data Storage and History Management

All user queries and responses are stored in a database for:

- Future reference
- Chat history tracking
- System improvement

C. Algorithm

Algorithm 1: Disease Prediction Process

Input: User symptom text
Output: Predicted disease and AI-generated response

1. Receive user input symptoms
2. Apply text preprocessing
3. Convert text to TF-IDF vector
4. Input vector into Logistic Regression model
5. Predict disease class
6. Generate AI-based explanation
7. Store interaction in database
8. Return result to user

D. System Components Table

Table 2: Components of Proposed System

Module	Technique Used	Description
Input Module	Web Interface (Flask)	Accepts user symptoms
Preprocessing Module	NLP	Cleans and normalizes text
Feature Extraction	TF-IDF	Converts text into vectors
Prediction Module	Logistic Regression	Predicts disease
AI Response Module	Generative AI	Provides explanation and advice
Database Module	SQLite	Stores users and chat history

E. Advantages of Proposed System

- Handles natural language input efficiently
- Provides accurate disease prediction
- Generates detailed and explainable outputs
- Improves accessibility to healthcare assistance
- Scalable and easy to deploy

V.System Architecture

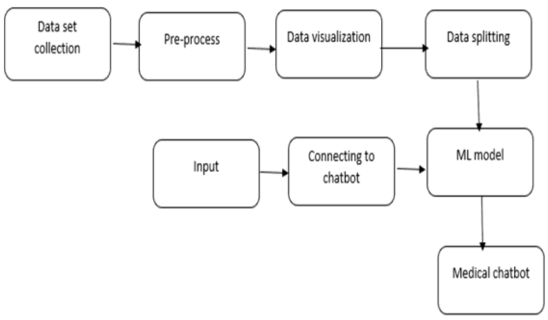

Figure 1: System Architecture of Healthcare

The proposed AI-enabled healthcare chatbot framework is developed using a modular architecture in which each component is responsible for a specific task within the data processing and disease prediction workflow. This modular design enhances system scalability, maintainability, flexibility, and overall operational efficiency.

A. Dataset Collection Module

This module is responsible for acquiring healthcare-related datasets containing symptom descriptions and their corresponding disease labels. The collected dataset serves as the primary source for training and validating the machine learning model. Data may be obtained from publicly available healthcare repositories, medical datasets, or structured clinical records.

B. Data Preprocessing Module

The preprocessing module prepares the collected data for effective analysis and model training. Various preprocessing operations are performed, including:

• Elimination of noisy and irrelevant information
• Treatment of missing or incomplete values
• Text preprocessing techniques such as lowercasing, tokenization, and stop-word removal

These operations ensure that the dataset remains consistent, clean, and suitable for machine learning applications.

C. Data Visualization Module

This module is responsible for exploring and visualizing the dataset to identify hidden patterns and relationships among different attributes. Visualization techniques such as histograms, scatter plots, and correlation matrices are utilized to analyze data distributions, feature interactions, and feature significance, thereby improving overall model interpretability.

D. Data Splitting Module

The processed dataset is partitioned into training and testing subsets, commonly using an 80:20 ratio. This division allows the model to learn from one portion of the data while being evaluated on previously unseen samples, helping to minimize overfitting and enhance generalization capability.

E. Machine Learning Model Module

This module performs the training, validation, and testing of the disease prediction model. Within the proposed framework:

• TF-IDF is employed for textual feature extraction
• Logistic Regression is utilized as the disease classification algorithm

The model learns associations between symptom descriptions and disease categories, enabling accurate prediction of the most probable disease based on user input.

F. Input Module

This module serves as the primary interaction point for users, allowing them to enter symptoms in natural language through a web-based interface. It facilitates real-time communication between the user and the healthcare chatbot system.

G. Chatbot Integration Module

The chatbot integration module processes user-provided symptom information using Natural Language Processing techniques and establishes communication with the trained machine learning model. This module ensures seamless interaction and efficient information exchange between the frontend interface and backend prediction engine.

H. Medical Chatbot Output Module

This module generates the final response delivered to the user. The generated output includes:

• Predicted disease condition
• Detailed explanation of the disease
• Risk-level assessment
• Preventive recommendations and lifestyle guidance

The response is produced through the combined utilization of machine learning-based disease prediction and Generative AI-generated medical insights, ensuring both predictive accuracy and interpretability for end users.

Vi. System Implementation

This section presents the implementation details of the proposed AI-driven healthcare chatbot framework. The system combines Natural Language Processing (NLP), Machine Learning, and web-based technologies to deliver an intelligent and interactive healthcare assistance platform. The framework is developed using Python and the Flask web framework, providing flexibility, scalability, and ease of deployment across different environments.

A. Development Environment

The proposed system is developed using Python because of its extensive ecosystem of machine learning, data processing, and NLP libraries. The Flask framework is employed to build the web-based application, facilitating smooth communication between the user interface and backend processing components. SQLite is utilized as the database management system for securely storing user credentials, authentication information, and chatbot interaction history. The selected development environment enables rapid application development, efficient module integration, and streamlined deployment of the healthcare chatbot framework.

B. Software and Hardware Requirements

Table 3: System Requirements

Component	Specification
Operating System	Windows/Linux/Mac
Programming Language	Python
Framework	Flask
Database	SQLite
Libraries	Scikit-learn, Pandas, NumPy, NLP
RAM	Minimum 4 GB
Processor	Intel i3 or above

C. Data Processing and Feature Engineering

The dataset containing symptom descriptions and corresponding diseases undergoes preprocessing using Natural Language Processing (NLP) techniques. The preprocessing stage includes tokenization, text normalization, and the elimination of irrelevant or unnecessary words. Following preprocessing, TF-IDF vectorization is applied to transform textual symptom data into numerical feature representations suitable for machine learning algorithms. This process effectively captures the significance of individual symptoms and improves the quality of feature representation for disease prediction.

D. Model Implementation

The proposed framework employs a Logistic Regression classifier for disease prediction. The model is trained using the preprocessed dataset, where symptom-based feature vectors serve as input attributes and disease categories act as output labels. The dataset is divided into training and testing subsets to assess predictive performance and model reliability. During training, the classifier learns the relationships between symptom patterns and disease classes, enabling it to accurately predict the most probable disease for previously unseen user inputs.

E. Chatbot Integration

The trained machine learning model is integrated into a Flask-based chatbot interface to facilitate interactive healthcare assistance. Users communicate with the system by entering symptoms in natural language through the web application. The chatbot processes the input using NLP techniques and forwards the extracted features to the prediction model for disease classification. Based on the predicted disease, the system generates an appropriate response for the user. In addition, a Generative AI component is incorporated to provide detailed disease explanations, risk assessments, preventive recommendations, and healthcare guidance. This integration enhances the overall user experience by delivering informative, interactive, and human-like responses.

F. Database Implementation

Two SQLite databases are used in the system:

- **User Database:** Stores user credentials securely using hashed passwords
- **Chat Database:** Stores user queries and chatbot responses for history tracking

This ensures data persistence and allows users to revisit previous interactions.

G. System Workflow Implementation

The implemented system follows these steps:

1. User logs into the system
2. User enters symptoms via chatbot interface
3. Input is preprocessed using NLP
4. Features are extracted using TF-IDF
5. Logistic Regression model predicts disease
6. Generative AI provides detailed response
7. Results are displayed and stored in database

H. Key Implementation Features

- Secure authentication using password hashing
- Real-time chatbot interaction
- Accurate disease prediction using ML model
- AI-generated explanations and recommendations
- Persistent chat history storage

I. Advantages of Implementation

- Lightweight and efficient system design
- Easy deployment using Flask
- Scalable architecture for future enhancements
- User-friendly interface with real-time responses

Vii. Experimental Results And Analysis

This section presents the experimental evaluation of the proposed AI-driven healthcare chatbot system. The objective is to assess the effectiveness of the system in predicting diseases based on user-input symptoms and providing meaningful responses. The performance of the machine learning model and overall system behavior are analyzed using standard evaluation metrics.

A. Experimental Setup

The system is implemented using Python and Flask, with the dataset consisting of symptom–disease pairs. The dataset is preprocessed using Natural Language Processing (NLP) techniques and transformed into numerical features using TF-IDF vectorization. The Logistic Regression model is trained using an 80:20 train-test split to ensure unbiased evaluation.

B. Performance Metrics

The performance of the model is evaluated using the following metrics:

- **Accuracy:** Measures the proportion of correctly predicted instances
- **Precision:** Indicates the correctness of positive predictions
- **Recall:** Measures the ability to identify all relevant cases
- **F1-Score:** Harmonic mean of precision and recall

C. Model Performance

Table 4: Performance Metrics of Logistic Regression Model

Metric	Value
Accuracy	0.88
Precision	0.87
Recall	0.86
F1-Score	0.86

The results indicate that the Logistic Regression model performs effectively in classifying diseases based on symptoms. The high accuracy and balanced precision-recall values demonstrate the reliability of the model.

D. Comparative Analysis

To evaluate the effectiveness of the proposed system, it is compared with traditional rule-based and keyword-based healthcare systems.

Table 5: Comparison with Existing Systems

Criteria	Rule-Based System	Proposed AI System
Input Understanding	Limited	Advanced (NLP-based)
Prediction Accuracy	Moderate	High
Automation Level	Low	High
Context Awareness	Not Supported	Supported
Response Generation	Static	Dynamic (AI-based)
Scalability	Limited	High

E. Results

Fig 2: Home Page of Health care Framework

This platform delivers a data-driven AI framework to enhance healthcare outcomes. It leverages Natural Language Processing and Machine Learning for intelligent clinical insights. Users can access diagnosis tools, patient history, and system analytics in one place.

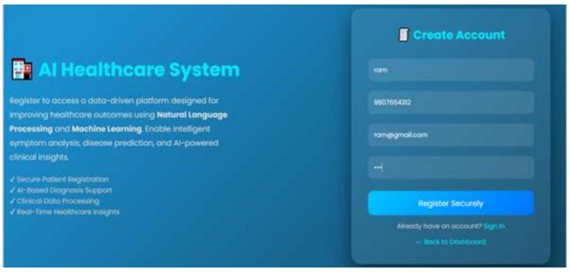

Fig 3: Create Account / Sign in Page

Create Account Page and sign in:

This page enables new users to register securely on the healthcare platform. It collects essential details to personalize clinical analysis and system access. Once registered, users can explore AI-powered healthcare features seamlessly.

The sign-in page allows existing users to quickly access their accounts. It ensures authentication for safe and reliable system usage. Users can resume healthcare analysis and patient interaction without delay.

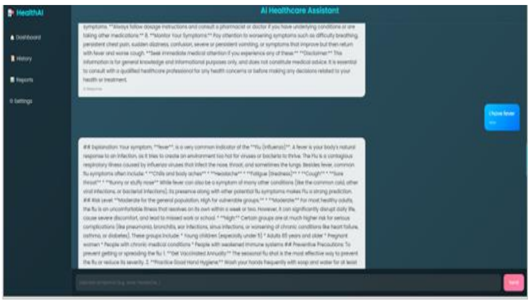

Fig 4: Results Page (Healthcare Chatbot)

This module presents AI-generated responses based on the symptoms and information provided by patients. By utilizing NLP-based models, the system delivers relevant healthcare insights and preliminary disease predictions. Users can interact with the chatbot in real time to obtain medical guidance, symptom-related information, and decision-support assistance, thereby enhancing the overall healthcare experience.

F. Summary

In summary, the proposed AI-driven healthcare chatbot demonstrates strong performance in disease prediction while providing a scalable, efficient, and user-friendly healthcare assistance platform. The experimental outcomes confirm the effectiveness of integrating Natural Language Processing, Machine Learning, and Generative AI technologies to deliver intelligent, interactive, and reliable healthcare support.

Viii. Discussion

The discussion of the proposed AI-driven healthcare chatbot is presented through a structured step-by-step analysis based on the functionality and performance of each system component.

Step 1: Effectiveness of Natural Language Processing

The proposed framework effectively employs Natural Language Processing techniques to analyze and interpret unstructured user inputs. Unlike conventional rule-based systems, the NLP module captures contextual information from symptom descriptions, enabling more accurate and flexible interpretation of user queries. This capability significantly improves system usability and applicability in real-world healthcare environments.

Step 2: Performance of the Machine Learning Model

The Logistic Regression classifier demonstrates dependable predictive performance, achieving high accuracy along with balanced precision and recall values. The model successfully identifies disease patterns from symptom data and generates reliable predictions, confirming the effectiveness of the machine learning approach adopted in the framework.

Step 3: Integration of Generative AI

The incorporation of Generative AI extends the functionality of the system by generating comprehensive disease-related information, including condition explanations, risk assessments, and preventive recommendations. This enhancement improves user engagement and understanding through interactive conversational responses.

Step 4: User Interaction and System Usability

The chatbot interface provides an intuitive and user-friendly environment with real-time response capabilities. Users can describe symptoms using natural language, making the system accessible to both technical and non-technical individuals. This ease of interaction contributes to a more effective healthcare assistance experience.

Step 5: System Limitations

The effectiveness of the framework is influenced by the quality, diversity, and size of the training dataset. Limited datasets may reduce the model's ability to generalize, particularly when handling rare or complex medical conditions. Furthermore, the system is intended to provide preliminary healthcare guidance and should not be considered a substitute for professional medical diagnosis or consultation.

Step 6: Dependency and Performance Constraints

The utilization of external AI services for generating detailed responses may introduce latency and dependency-related challenges. Under certain operating conditions, these factors can affect the responsiveness and overall real-time performance of the system.

Step 7: Data Privacy and Security Considerations

As the framework processes sensitive healthcare information, maintaining data privacy and security is a critical requirement. Appropriate security measures, data protection mechanisms, and compliance with healthcare regulations are necessary to ensure safe and reliable deployment.

Step 8: Opportunities for Enhancement

The current implementation primarily relies on traditional machine learning techniques and does not incorporate advanced deep learning or transformer-based architectures. Future integration of such technologies can improve disease prediction accuracy, contextual understanding, and overall system intelligence.

Step 9: Overall Impact of the System

Despite certain limitations, the proposed framework offers a scalable, efficient, and practical healthcare assistance solution. By providing rapid preliminary disease assessment and intelligent medical guidance, the system reduces the workload on healthcare professionals and demonstrates strong potential for real-world healthcare applications.

Ix. Conclusion

This paper presented an AI-driven healthcare chatbot framework that integrates Natural Language Processing (NLP), Machine Learning, and Generative AI technologies to provide preliminary medical diagnosis and interactive healthcare assistance. The system effectively analyzes natural language symptom descriptions, extracts relevant information, and predicts potential diseases using a Logistic Regression classifier. Experimental evaluation demonstrates reliable predictive performance and enhanced user interaction through conversational responses, disease explanations, and preventive recommendations. The incorporation of NLP enables contextual understanding of user inputs, overcoming the limitations associated with conventional keyword-based approaches, while the modular architecture supports scalability and ease of implementation.

Although the framework demonstrates promising results, its performance is influenced by dataset size, diversity, and the simplicity of the classification model. The system is designed to support preliminary diagnosis and healthcare guidance rather than replace professional medical consultation. Future developments may include the integration of advanced deep learning and transformer-based models to improve prediction accuracy and contextual comprehension. Expanding the dataset to include a broader range of medical conditions can further enhance generalization capabilities. Additional improvements such as multilingual support, enhanced privacy and security mechanisms, and integration with real-time health monitoring systems can significantly strengthen the framework. With these advancements, the proposed healthcare chatbot can evolve into a more intelligent, scalable, and comprehensive healthcare assistance platform.

References

[1] E. Shortliffe and J. Cimino, Biomedical Informatics: Computer Applications in Health Care and Biomedicine, 4th ed. London, U.K.: Springer, 2014.

[2] R. Miotto, F. Wang, S. Wang, X. Jiang, and J. T. Dudley, “Deep learning for healthcare: Review, opportunities and challenges,” Briefings in Bioinformatics, vol. 19, no. 6, pp. 1236–1246, 2018.

[3] Z. Obermeyer and E. J. Emanuel, “Predicting the future — Big data, machine learning, and clinical medicine,” New England Journal of Medicine, vol. 375, no. 13, pp. 1216–1219, 2016.

[4] A. Esteva, A. Robicquet, B. Ramsundar, et al., “A guide to deep learning in healthcare,” Nature Medicine, vol. 25, no. 1, pp. 24–29, 2019.

[5] D. Ravi, C. Wong, F. Deligianni, et al., “Deep learning for health informatics,” IEEE Journal of Biomedical and Health Informatics, vol. 21, no. 1, pp. 4–21, 2017.

[6] H. Topol, Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again. New York, NY, USA: Basic Books, 2019.

[7] J. Devlin, M. W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” Proc. NAACL-HLT, pp. 4171–4186, 2019.

[8] A. Vaswani, N. Shazeer, N. Parmar, et al., “Attention is all you need,” Advances in Neural Information Processing Systems, vol. 30, pp. 5998–6008, 2017.

[9] T. Brown, B. Mann, N. Ryder, et al., “Language models are few-shot learners,” Advances in Neural Information Processing Systems, vol. 33, pp. 1877–1901, 2020.

[10] K. Rajkomar, E. Oren, K. Chen, et al., “Scalable and accurate deep learning with electronic health records,” NPJ Digital Medicine, vol. 1, no. 18, pp. 1–10, 2018.

[11] S. Shickel, P. J. Tighe, A. Bihorac, and P. Rashidi, “Deep EHR: A survey of recent advances in deep learning techniques for electronic health record analysis,” IEEE Journal of Biomedical and Health Informatics, vol. 22, no. 5, pp. 1589–1604, 2018.

[12] A. Holzinger, C. Biemann, C. S. Pattichis, and D. B. Kell, "What do we need to build explainable AI systems for the medical domain?," arXiv Preprint arXiv:1712.09923, 2017.

[13] S. R. Safavian and D. Landgrebe, "A survey of decision tree classifier methodology," IEEE Transactions on Systems, Man, and Cybernetics, vol. 21, no. 3, pp. 660–674, 1991.

[14] T. M. Mitchell, Machine Learning. New York, NY, USA: McGraw-Hill, 1997.

[15] C. C. Aggarwal, Neural Networks and Deep Learning. Cham, Switzerland: Springer, 2018.

[16] F. Chollet, Deep Learning with Python, 2nd ed. Shelter Island, NY, USA: Manning Publications, 2021.

[17] S. Bird, E. Klein, and E. Loper, Natural Language Processing with Python. Sebastopol, CA, USA: O'Reilly Media, 2009.

[18] D. Jurafsky and J. H. Martin, Speech and Language Processing, 3rd ed. Upper Saddle River, NJ, USA: Pearson, 2023.

[19] M. F. Porter, "An algorithm for suffix stripping," Program, vol. 14, no. 3, pp. 130–137, 1980.

[20] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. Cambridge, MA, USA: MIT Press, 2016.