Clinical Report Interpretation Using Explainable Artificial Intelligence Techniques

Anitha R[1*], Sindhuja S[2]

[1,2] *Department of Computer Science and Engineering, Sri Venkateswara College of Engineering, Sriperumbudur, Tamilnadu, India*

***Abstract* - Electronic health records, diagnostic summaries, and other clinical documents are being generated at an increasingly large scale due to the rapid digital transformation of healthcare systems. However, these records are often highly technical and difficult for patients or non-expert users to understand. As a result, many individuals face challenges in interpreting their medical information, which can negatively impact overall health awareness and decision-making. To overcome this limitation, the paper introduces an Explainable Artificial Intelligence (XAI) framework aimed at automatically processing and simplifying medical reports. The primary goal of the system is to improve accessibility by transforming complex clinical data into clear, understandable language while also supporting knowledge-based explanations. The proposed framework follows a structured multi-stage pipeline that combines multiple AI techniques. It begins with Optical Character Recognition (OCR) to extract text from scanned or image-based medical documents. This is followed by biomedical Named Entity Recognition (NER), which identifies and categorizes key medical terms such as diseases, symptoms, medications, and procedures. Finally, transformer-based models are used to generate simplified and coherent summaries of the extracted clinical information. The system also includes a knowledge retrieval module that links identified medical entities with relevant explanations, definitions, and contextual insights. This helps users better understand medical terminology and improves the interpretability of the generated summaries. By combining extraction, summarization, and explanation, the framework ensures that medical information is not only processed but also made meaningful for end users. The effectiveness of the proposed system is validated through experimental evaluation, focusing on accuracy, readability, and user comprehension.**

***Keywords:* Explainable Artificial Intelligence, Medical Report Processing, Biomedical Named Entity Recognition, Transformer-Based Summarization, Clinical Knowledge Retrieval.**

I. INTRODUCTION

The ongoing digitalization of healthcare services has resulted in a rapid increase in the amount of medical data being generated in clinical environments. Hospitals and healthcare institutions now produce large volumes of documents such as patient reports, diagnostic summaries, laboratory results, and treatment records. These documents contain crucial information about a patient's medical condition, history, and treatment process, which makes them essential for effective healthcare decision-making. However, the highly technical language and unstructured nature of these records often make them difficult to interpret for both healthcare professionals and patients.

In recent years, advancements in artificial intelligence and natural language processing have significantly improved the ability to analyze and understand medical text automatically. Modern deep learning approaches, particularly transformer-based models, have demonstrated strong capabilities in handling complex biomedical language tasks, including text understanding and summarization. Alongside this, named entity recognition techniques are widely used to extract important clinical elements such as diseases, drugs, symptoms, and medical procedures from unstructured documents. Retrieval-based approaches further enhance understanding by providing additional context and explanations for extracted medical information. Additionally, summarization techniques help condense lengthy medical reports into simpler and more structured formats that are easier to read and interpret. Even with these technological improvements, most existing solutions still focus on individual tasks rather than offering a fully integrated system. Many current models also function as black-box systems, which makes their decision-making process difficult to interpret. This lack of transparency reduces user trust, especially in critical fields like healthcare where clarity and reliability are essential. Therefore, there is a strong need for systems that not only process medical information effectively but also provide understandable and transparent outputs. To overcome these limitations, this work presents an integrated framework for automated analysis of medical reports. The proposed system combines multiple stages, including extraction of text from scanned documents, identification of relevant medical entities, generation of simplified summaries, and retrieval of supporting explanations. These components work together to form a complete pipeline that processes medical data in a structured and meaningful way. The key contribution of this study lies in combining multiple natural language processing techniques into a single unified and explainable framework. Unlike traditional approaches that treat each task separately, this system enhances interpretability by providing contextual explanations for extracted medical information. This improves understanding for end users and helps bridge the gap between complex clinical data and practical usability.

II. LITERATURE SURVEY

The rapid digital transformation of healthcare systems has resulted in a substantial rise in the amount of medical data being generated and stored. This includes a wide range of clinical documents such as physician notes, laboratory results, discharge summaries, and diagnostic reports. These records contain essential information related to a patient's diagnosis, treatment history, and ongoing care, making them crucial for effective healthcare delivery and decision-making [1]. However, the unstructured nature and technical complexity of these documents often make them difficult to interpret for both medical professionals and patients. Recent progress in artificial intelligence and natural language processing has significantly improved the ability to automatically process and analyze clinical text. In particular, transformer-based architectures have demonstrated strong performance in biomedical text mining and summarization tasks [2], [3]. Biomedical Named Entity Recognition (NER) techniques are widely used to extract important clinical entities such as diseases, medications, symptoms, and medical procedures from unstructured text [4]–[6]. In addition, advanced transformer-based biomedical models have enhanced the accuracy of medical language understanding tasks [10], while retrieval-based methods have improved the ability to provide clearer explanations for complex clinical information [11]. Summarization techniques further assist in reducing the length and complexity of medical documents, making them more accessible and easier to comprehend. Despite these technological advancements, most existing systems are limited to isolated tasks such as entity extraction or summarization, rather than offering a fully integrated solution. Furthermore, many current AI models function as black-box systems, lacking transparency and interpretability in their decision-making processes. These limitations reduce their reliability and hinder their adoption in real-world clinical environments. To overcome these challenges, Explainable Artificial Intelligence (XAI) has emerged as a promising approach that enhances model transparency and allows users to better understand system outputs, thereby improving trust and usability [13], [14]. To address these issues, this work proposes a unified framework for automated medical report analysis called MediScan AI. The proposed system integrates multiple NLP and AI components, including Optical Character Recognition (OCR) for text extraction, biomedical Named Entity Recognition for identifying clinical entities, transformer-based models for summarization, and a knowledge retrieval module for generating contextual explanations. Together, these components enable complete end-to-end processing of medical documents in a structured and meaningful manner. The main contribution of this work lies in the integration of multiple natural language processing techniques into a single explainable AI framework. Unlike traditional approaches that focus on individual tasks, the proposed system incorporates a knowledge retrieval mechanism that enhances interpretability by explaining extracted medical entities in context. This ensures that both healthcare professionals and patients can better understand complex medical information, ultimately improving communication and decision-making in healthcare settings.

III.PROPOSED FRAMEWORK

The proposed framework, MediScan AI, is designed to provide an automated and explainable solution for analyzing medical reports using natural language processing techniques. Medical documents are often complex, unstructured, and difficult to interpret, which creates challenges in extracting meaningful information efficiently. Existing methods generally focus on isolated tasks such as entity extraction or summarization, without offering clear explanations of the results. This limitation highlights the need for a unified and interpretable system for medical report analysis. To address these challenges, the proposed framework integrates multiple components, including Optical Character Recognition (OCR), preprocessing modules, biomedical Named Entity Recognition (NER), transformer-based summarization, and knowledge retrieval mechanisms. This combination enables the system to process different types of medical documents such as PDFs and scanned images, converting them into structured and usable information. Unlike traditional approaches that handle individual tasks separately, the proposed system follows an end-to-end pipeline for medical report understanding. It not only focuses on efficient information extraction but also emphasizes explainability, making the outputs more useful for both healthcare professionals and patients. The architecture of the MediScan AI system is shown in Figure 1 and follows a sequential pipeline design. In this structure, the output of each module becomes the input for the next, ensuring smooth and organized data flow. The system consists of four main components: a user interface, backend processing unit, AI processing engine, and data storage layer. The process begins when a user uploads medical documents through a web-based interface. Once uploaded, the backend system receives the files and performs initial handling, including basic validation and privacy-preserving preprocessing steps to ensure secure data handling. After preprocessing, the data is forwarded to the AI processing engine. The first stage within this engine involves Optical Character Recognition (OCR), which converts scanned or image-based documents into machine-readable text. The extracted text then undergoes preprocessing steps such as normalization, noise removal, tokenization, and sentence segmentation to prepare it for further analysis. Following this, biomedical Named Entity Recognition (NER) is applied to identify key clinical entities such as diseases, medications, and procedures. The processed information is then passed to the transformer-based summarization module. This module generates two types of outputs: a clinical summary, which provides detailed and accurate medical information for healthcare professionals, and a simplified explanation mode, which presents the content in an easy-to-understand format for patients and non-expert users. Finally, the system produces structured outputs that include summarized reports, simplified explanations, and contextual insights generated through the knowledge retrieval module. These outputs are stored and displayed in a user-friendly format, enabling easy access and interpretation. The modular design of the proposed framework ensures scalability and flexibility, allowing individual components to be updated or replaced without affecting the overall system performance. The sequential pipeline reduces information loss and maintains consistency across processing stages. Additionally, the inclusion of explainability through dual-mode summarization improves usability in real-world healthcare scenarios by catering to both experts and non-experts, thereby enhancing accessibility and understanding of medical information.

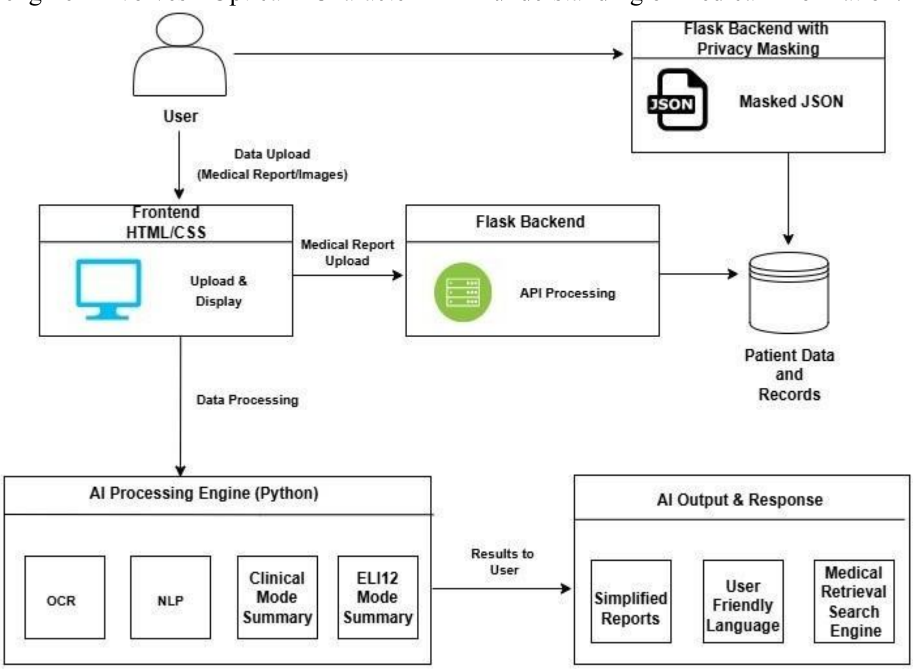

Figure 1. Architecture of the proposed MediScan AI framework

IV.METHOD

The proposed MediScan AI framework follows a structured and modular methodology for the automated analysis of medical reports. The complete process is organized into a sequence of stages that include data acquisition, text extraction, preprocessing, biomedical entity recognition, transformer-based summarization, knowledge retrieval, and explainability. Together, these steps form a unified pipeline that enables efficient processing of medical documents and supports accurate interpretation of clinical information. This study adopts an experimental research design in which a full pipeline system is developed to automatically process medical reports. The proposed framework transforms unstructured medical documents into structured outputs through multiple stages such as Optical Character Recognition (OCR), preprocessing, Named Entity Recognition (NER), summarization using transformer models, and knowledge retrieval. The main aim of this design is to ensure systematic conversion of raw clinical data into meaningful and interpretable results. The system is also evaluated using appropriate performance measures to assess its effectiveness in real-world applications. The implementation of the system uses a combination of modern tools and technologies to achieve efficient processing. Text extraction from scanned documents is performed using Tesseract OCR, while natural language processing tasks such as preprocessing and biomedical entity recognition are handled using spaCy and SciSpaCy. Transformer-based models are used for advanced language tasks, where DistilBART is applied for summarization, FLAN-T5 is used for generating simplified explanations, and DistilBERT supports semantic understanding and knowledge retrieval. These models are integrated through the HuggingFace Transformers library. A hybrid retrieval approach combining TF-IDF and BM25 is used to improve the relevance of extracted contextual information. The backend is developed using Python Flask, while the frontend interface is built using HTML, CSS, and JavaScript. The system also generates structured outputs using ReportLab and is tested using biomedical datasets sourced from PubMed and related repositories. The dataset used in this work consists of publicly available medical reports, including discharge summaries, diagnostic reports, and clinical documents. Additionally, the MedQuAD dataset containing approximately 47,000 medical question-answer pairs is used to support the knowledge retrieval module. This diverse dataset ensures that the system can handle different types of medical terminology and document structures effectively, making the evaluation more reliable and comprehensive. The overall research procedure begins with uploading medical reports in formats such as PDF or images. These inputs are first processed using OCR to convert them into machine-readable text. The extracted text then undergoes preprocessing steps including normalization, noise removal, tokenization, and case conversion to improve data quality. After preprocessing, SciSpaCy is applied to identify key biomedical entities such as diseases, drugs, and procedures. The processed text is then passed to transformer-based models to generate concise summaries while retaining essential medical information. Finally, a knowledge retrieval mechanism using MedQuAD is applied to provide explanations for the extracted entities, enabling a more interpretable and informative analysis of medical reports.

V.RESULTS AND DISCUSSION

This section will present the results of the proposed MediScan AI framework, and a detailed discussion of the performance of the framework will be conducted. The effectiveness of the proposed system is evaluated using commonly used natural language processing metrics, including precision, recall, F1-score, and response time. These metrics help measure the accuracy and efficiency of the system in processing medical text and generating outputs.

Precision is defined as the ratio of correctly identified positive predictions to the total predicted positives:

Precision = TP / (TP + FP) (1)

Recall measures the proportion of correctly identified positives out of all actual positive cases:

Recall = TP / (TP + FN) (2)

The F1-score represents the harmonic mean of precision and recall, providing a balanced measure of model performance:

F1 = (2 × Precision × Recall) / (Precision + Recall) (3)

Response time is calculated as the difference between output generation time and input processing time:

Response Time = T_output − T_input (4)

Here, TP refers to true positives, FP refers to false positives, and FN refers to false negatives. In addition, T_input represents the time when the medical report is submitted to the system, while

T_output represents the time when the final analysis results are generated.

5.2 Performance Analysis

The performance results of the proposed MediScan AI framework are summarized in Table 1. The system achieves a precision of 0.96, recall of 0.90, and an F1-score of 0.93, indicating strong overall performance in medical report analysis tasks. When compared with baseline models such as SciSpaCy and BioBERT, the proposed approach demonstrates improved accuracy and better balance between precision and recall. The high precision value indicates that the system produces fewer false positive results, meaning it accurately identifies relevant medical entities. The improved recall shows that fewer relevant entities are missed during extraction. The strong F1-score reflects a balanced performance between precision and recall, confirming the reliability of the proposed framework in handling clinical text data effectively.

5.3 Comparative Analysis with Existing Models

A comparative evaluation between the proposed system and existing baseline models is illustrated in Figure 2. The results show that MediScan AI outperforms traditional approaches across multiple performance indicators, including entity recognition accuracy, summarization quality, OCR effectiveness, and text compression capability. Overall, the proposed framework demonstrates superior performance due to its integrated architecture that combines OCR, NLP, and retrieval-based techniques in a unified pipeline. The proposed MediScan AI system has better performance metrics, such as entity recognition, summarization, OCR accuracy, and compression ratio.

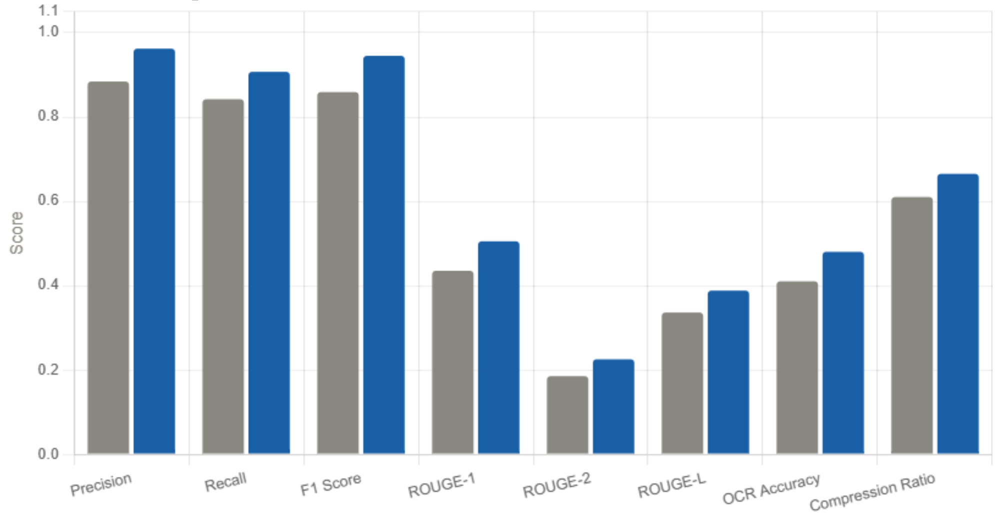

Figure 2. Performance comparison of MediScan AI (blue) with existing models (gray) across multiple evaluation metrics.

Performance Evaluation

The performance of the proposed system is evaluated using standard evaluation metrics for NLP. For biomedical named entity recognition, precision, recall, and F1-score are used as evaluation metrics. For summarization, ROUGE-1, ROUGE-2, and ROUGE-L are used as evaluation metrics. A comparative analysis of the proposed model with baseline models has been presented in Table 1.

Table 1. Performance comparison of models

Model	Precision	Recall	F1-Score
SciSpacy (Baseline)	0.89	0.85	0.87
BioBERT	0.92	0.88	0.90
MediScan AI (Proposed)	0.96	0.90	0.93

The summarization module scored a ROUGE-1 score of 0.50, a ROUGE-2 score of 0.22, and a ROUGE-L score of 0.38. This shows that the system is efficient in extracting key information from medical reports. The OCR module scored a moderate accuracy of 0.48. The relatively low OCR accuracy is due to different document formats and image quality but the pre-processing and NLP phases reduce its effect. However, this can be attributed mainly to variations in document structure and image quality.

VI. DISCUSSION

This paper presented MediScan AI, an explainable artificial intelligence framework designed for the automated analysis of medical reports using natural language processing techniques. The proposed system integrates OCR, biomedical Named Entity Recognition, transformer-based summarization, and knowledge retrieval into a unified pipeline to address the challenges associated with interpreting complex and unstructured medical data. The experimental evaluation shows that the framework achieves strong performance, with a precision of 0.96, recall of 0.90, and an F1-score of 0.93 for biomedical entity recognition tasks. The summarization component also delivers effective results in generating concise and meaningful clinical summaries. Furthermore, the inclusion of a knowledge retrieval module significantly improves system explainability by providing contextual explanations for extracted medical entities, thereby enhancing user understanding and trust. Overall, the proposed system offers a scalable and efficient solution that benefits both healthcare professionals and patients by simplifying access to complex medical information and supporting better decision-making. Despite its strengths, limitations such as OCR dependency and knowledge base coverage still exist. Future work will focus on improving OCR accuracy, integrating more advanced domain-specific models, enhancing multilingual capabilities, and refining the knowledge retrieval mechanism for better contextual understanding. In summary, MediScan AI demonstrates strong potential for real-world healthcare applications and establishes a solid foundation for further research in explainable AI for medical text analysis.

VII. CONCLUSION

The results demonstrate that the proposed MediScan AI framework successfully integrates multiple components—Optical Character Recognition (OCR), biomedical Named Entity Recognition, transformer-based summarization, and knowledge retrieval—into a unified system for medical report analysis. Unlike traditional approaches that handle these tasks independently, the proposed method provides an end-to-end solution that supports complete understanding of medical documents. The inclusion of explainable AI further improves transparency, making the system more trustworthy and suitable for healthcare applications. Experimental outcomes indicate that the proposed framework consistently outperforms baseline models across all evaluation metrics. The hybrid design enhances the accuracy of entity recognition while also improving the quality and relevance of generated summaries. This combined improvement highlights the advantage of integrating multiple NLP techniques within a single architecture rather than using isolated models. However, the system has certain limitations. The performance of the OCR module is strongly dependent on the quality of the input documents, meaning low-quality scans may reduce accuracy. In addition, the effectiveness of knowledge retrieval is influenced by the completeness and coverage of the underlying medical knowledge base. These factors can impact overall system reliability in certain real-world conditions. Future improvements may focus on enhancing OCR robustness, adopting more advanced biomedical transformer models, and expanding support for multilingual medical documents.

REFERENCES

[1] Y. Liu, F. Wang, P. Wang, Z. Zhou, H. Wang, J. Li, AI-MDT: An automatic and intelligent multidisciplinary team consultations platform for lung cancer diagnosis, J. Cancer Res. Clin. Oncol. 152 (2026) 32.

[2] Y. Vasilev, I. Raznitsyna, A. Pamova, T. Burtsev, T. Bobrovskaya, P. Kosov, Evaluating medical text summaries using automatic evaluation metrics and LLM-as-a-judge approach: A pilot study, Diagnostics 16(3) (2026).

[3] Q. Xie, J.A. Bishop, P. Tiwari, S. Ananiadou, Pre-trained language models with domain knowledge for biomedical extractive summarization, Knowl. Based Syst. 252 (2022) 109460.

[4] M. Kumar, P. Singh, P. Kashtriya, Enhanced biomedical named entity recognition using SpaCy and BERT models, Procedia Comput. Sci. 258 (2025) 1954–1961.

[5] A. Dash, S. Darshana, D.K. Yadav, V. Gupta, A clinical named entity recognition model using pretrained word embedding and deep neural networks, Decis. Anal. J. 10 (2024) 100426.

[6] V. Kocaman, D. Talby, Accurate clinical and biomedical named entity recognition at scale, Softw. Impacts 13 (2022) 100373.

[7] M. Asghari, D. Sierra-Sosa, A.S. Elmaghraby, BINER: A low-cost biomedical named entity recognition, Inf. Sci. 602 (2022) 184–200.

[8] S. Liu, A.B. McCoy, Q. Chen, A. Wright, Integrating rule-based NLP and large language models for statin information extraction from clinical notes, Int. J. Med. Inform. 205 (2026) 106104.

[9] R. El Shawi, L. Jamel, Leveraging ChatGPT and explainable AI for enhancing clinical decision support, Sci. Rep. 15(1) (2025).

[10] S. Madan, M. Lentzen, J. Brandt, D. Rueckert, M. Hofmann-Apitius, H. Fröhlich, Transformer models in biomedicine, BMC Med. Inform. Decis. Mak. 24 (2024) 214.

[11] Y. Guo, W. Qiu, G. Leroy, S. Wang, T. Cohen, Retrieval augmentation of large language models for lay language generation, J. Biomed. Inform. 149 (2024) 104580.

[12] M.M. Rahman, M.S. Irbaz, K. North, M.S. Williams, M. Zampieri, K. Lybarger, Health text simplification: An annotated corpus for digestive cancer education, J. Biomed. Inform. 158 (2024) 104727.

[13] Z. Sadeghi, R. Alizadehsani, M.A. Cifci, S. Kausar, R. Rehman, P.M. Pardalos, A review of explainable artificial intelligence in healthcare, Comput. Electr. Eng. 118 (2024) 109370.

[14] I.D. Mienye, G. Obaido, N. Jere, E. Mienye, K. Aruleba, I.D. Emmanuel, A survey of explainable artificial intelligence in healthcare, Inform. Med. Unlocked 51 (2024) 101587.

[15] B. Picton, S. Andalib, A. Spina, B. Camp, S.S. Solomon, J. Liang, Assessing AI simplification of medical texts, Int. J. Med. Inform. 195 (2025) 105743.

[16] G. Mariammal, K. Swetha, S. Jimrys, Medical report simplification system: Enhancing healthcare accessibility using NLP-based extraction and AI-driven explanation, in: Proc. ICICV, 2025, pp. 541–546.

[17] J.D. Oliveira et al., Development and evaluation of a clinical note summarization system using large language models, Commun. Med. 5 (2025).

[18] A.N.A. Abeo, S. Armstrong, M. Scriney, H. Goss, Artificial intelligence techniques and health literacy: A systematic review, Mayo Clin. Proc. Digit. Health 3(4) (2025).

[19] D. Gifu, AI-backed OCR in healthcare, Procedia Comput. Sci. 207 (2022) 1134–1143.

[20] M. Ponnuru et al., Image-based extraction of prescription information using OCR-Tesseract, Procedia Comput. Sci. 235 (2024) 1077–1086.

[21] D. Verma, D. Darlan, R. Mallipeddi, Multi-objective optimization of large language model summaries for clinical text narratives, IEEE Access 14 (2026) 13597–13607.

[22] D. Hazime, L. Raffi, E. Burner, Impact of mHealth-augmented social support on health care use among patients with diabetes: Secondary analysis of the TExT-MED+FANS trial, Interact. J. Med. Res. 15 (2026) e65113.

[23] S. Bandare, Natural language processing with AI for unstructured claims data, Am. J. Bioinform. 7(1) (2026).

[24] M.H. Nguyen et al., Assessing genetic counseling efficiency with natural language processing, J. Am. Med. Inform. Assoc. 33(2) (2026) 295–303.

[25] P. El-Kafrawy, O. Guedhami, Text analytics techniques: Sentiment analysis of healthcare narratives, in: M.F. El-Amin, P. El-Kafrawy (Eds.), Mathematical Modeling for Big Data Analytics, Morgan Kaufmann, 2026, pp. 193–204.

[26] A.T. Tran et al., Explainable natural language processing (NLP) models to predict 90-day mortality of different stroke types from clinical note, Stroke 57(Suppl. 1) (2026).