



Medical Report Simplifier

Komal Parate¹, Sonal Nimbalkar², Himanshu Talekar³, Devang Chopkar⁴, Mayank Hirekhan⁵, Aastha Pathak⁶
^{1,2} Assistant Professor, Dept. of Em-Tech (AIDS), S.B Jain Institute of Technology Management & Research
^{3,4,5,6}UG- Dept. of Em-Tech (AIDS), S.B Jain Institute of Technology Management & Research
Emails: komalparate@sbjit.edu.in¹, Sonalnimbalkar@sbjit.edu.in², himanshut.aids22@sbjit.edu.in³,
devangc.aids22@sbjit.edu.in⁴, mayankh.aids22@sbjit.edu.in⁵, aasthap.aids22@sbjit.edu.in⁶

Abstract

Medical reports from an MRI, CT, or X-ray imaging are usually filled with complicated terms that a patient fails to understand. Such a failure to comprehend may result in a heightened state of anxiety, lower treatment adherence, and limited patient participation in their care. The paper, therefore, introduces a Hybrid Natural Language Processing (NLP) framework that comprises T5 and LLaMA-based models for converting complex medical text into simple language without losing diagnostic accuracy. Besides, the proposed system has Optical Character Recognition (OCR) for the scanned reports, entity recognition for the medical terms, and deep learning-based simplification with fine-tuned transformer models. The backend is in FastAPI, and the frontend in Flask, thus, combining the whole process from report uploading, text extraction to patient-friendly simplification. Assessment outcomes indicate that the text becomes more readable, and essential medical entities are preserved. This is an indication of the role of AI powered tools in lessening the communication barrier between doctors and patients.

Keywords: Medical report simplification, Natural Language Processing (NLP), Transformer models, T5, LLaMA, Optical Character Recognition (OCR), Deep learning, Text simplification, FastAPI, Flask, Healthcare AI, Radiology reports

1. Introduction

Diagnostic reports that include MRI, CT, and X-ray summaries play a significant role in patient care as well as communication between clinicians. Nevertheless, these reports are generally loaded with technical terms and abbreviations that are designed for medical professionals. That makes it difficult for patients and caregivers to comprehend them which in turn can lead to a plethora of problems like confusion, anxiety, and low participation in treatment. Studies reveal that over 60% of the patients have difficulties understanding radiology reports because of the medical jargon used, which can result in mistakes and unhealthy choices. To fix the issue, a lot of attention is given to AI tools that can simplify complicated medical texts into easily understandable language by patients. Previous methods like rule-based systems or those that replace words with medical dictionaries can simplify terms but generally they do not consider the context and may even change the meaning. Current deep learning models in Natural Language

Processing, especially those employing Transformer technologies, are far more efficient in generating, summarizing, and rephrasing text. T5 (Text-to-Text Transfer Transformer) is a model that can be used for different NLP tasks by converting them into a text-to-text format and hence it can be used for simplification and summarization tasks. T5 trained on medical data can produce less complicated sentences while still retaining the original meaning. In some cases, while T5 is able to simplify words and sentences, it fails to give clear, natural explanations of medical results as it does not have a deeper contextual understanding. As a solution to this problem, Large Language Models such as LLaMA are combined with traditional NLP systems. LLaMA trained on large multilingual datasets and having a good understanding of context can generate clear explanations which are user-specific. In a hybrid system, the accuracy of results as well as their easy understanding can be very well achieved by the

combination of T5's ability to simplify the text and LLaMA's to understand the context which is very important for medical communication. This Journal presents the Hybrid NLP Framework for simplification of medical reports. The apparatus employs T5 for simplification of sentences and LLaMA 3 for context insertion. With the help of Optical Character Recognition, it can also process the texts of both the digitized and the scanned medical reports. Then by the use of Named Entity Recognition, it detects medical terms and generates patient-friendly summaries with retaining the original content. The backend is created using FastAPI, and the frontend is done with Flask which makes it convenient for users to upload reports and check the results. A MySQL database retains the data and user feedback with the intention of gradually enhancing the system. The system is not merely a text simplifier for the medical domain. It also enables people to gain a better understanding of their health and enhances their feeling of being in control. Through the combination of deep learning and NLP technologies, this work intends to facilitate the provision of healthcare information in a way that is more accessible, clearer, and less exclusive in nature shown in Figure 1.

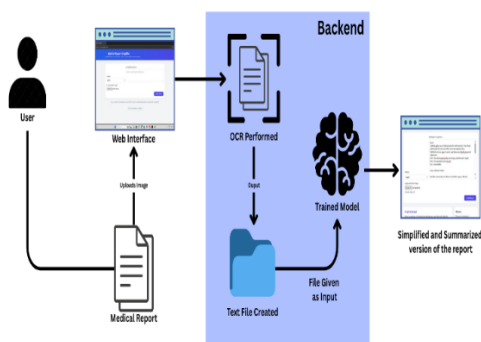


Figure 1 System Architecture Overview.

A user can upload a medical report via a web interface. The backend after processing the file with OCR and a trained NLP model generates a simplified version of the report in a patient-friendly language.

2. Literature Review

Recent advancements in Natural Language Processing (NLP), Optical Character Recognition (OCR), and transformer-based deep learning have

made it possible to automate the interpretation and simplification of complex medical text. The proposed Hybrid NLP framework builds on existing research in five key areas: medical text extraction, biomedical language modeling, medical text simplification, explainable summarization, and system integration.

A. Medical Text Extraction and OCR Integration
Medical reports are often saved as scanned or image-based documents, which makes efficient text extraction a necessary first step. Verma et al. [1] proposed a system to simplify medical reports using OCR technology, showing its potential to convert raw image data into readable digital text. Additionally, Li et al. [20] improved OCR models for document image recognition in structured domains. Kumar et al. [21] optimized OCR for multilingual document analysis, which is essential in healthcare settings that use various report formats. Chandra and Gupta [28], along with Meier and Braun [29], built OCR and NLP pipelines for smart document processing, creating systems that can be adapted for healthcare. These studies emphasize the need for effective OCR pre-processing in hybrid NLP systems to ensure reliable language modeling.

B. Biomedical Natural Language Processing and Large Language Models
A deep understanding of medical language is crucial for text simplification. Hu et al. [3] investigated BERT-based fine-tuning for processing medical text, showing improvements in information retrieval and entity extraction. Ming et al. [2] and Xu et al. [18] developed knowledge-guided biomedical summarization techniques using Large Language Models (LLMs) that maintain factual accuracy. Chen and Pergola [4] introduced a multi agent LLM framework called Society of Medical Simplifiers, which closely follows the hybrid approach of combining models like T5 and LLaMA for context-based rephrasing and semantic accuracy. Nishio et al. [19] demonstrated automatic summarization of radiology reports using large language models, proving their reliability in understanding radiology texts. Together, these studies lay the groundwork for transformer architectures designed for language simplification in the proposed system.

C. Medical Text Simplification and Readability Improvement
Simplifying complex diagnostic narratives into easy-



to understand language is a vibrant research field. Phatak et al. [6] introduced TESLEA, a deep reinforcement learning approach for medical text simplification that balances readability and factual accuracy. Flores et al. [7] improved readability through unlikelihood training, while Ondov et al. [8] provided a thorough survey of automated biomedical text simplification methods. MacAvaney et al. [5] suggested ontology-aware summarization to keep semantic integrity during simplification, and Basu et al. [17] created Med-EASi, a dataset for controlled medical text simplification. Devaraj et al. [13] contributed techniques for paragraph-level simplification of medical texts, focusing on coherence. These studies support the transformer-based simplification framework used in this work, which combines T5's seq2seq generation with LLaMA's contextual understanding for better fluency and precision. D. Explainable and Factual Summarization in Medical Reports Maintaining factual accuracy and clarity in AI-generated medical text is vital for clinical use. Zhang et al. [9] enhanced factual consistency in summaries using knowledge-based constraints, whereas Guo et al. [12] looked into comprehensive medical report generation systems. Xu et al. [18] proposed methods to edit factual knowledge in medical LLMs to boost explainability. Sato and Tanaka [22] and Nguyen et al. [23] showed that LLM based summarization methods could be adapted from software and legal fields to complex medical texts, highlighting their versatility. These works emphasize the importance of factual consistency, transparency, and user-centered validation, which are incorporated into the evaluation process of the proposed system using readability and entity-preservation measures. E. System Integration and Cross-Domain NLP Frameworks Bringing NLP models into scalable, real-time systems has been studied in various fields. Biruntha et al. [10] reviewed deep learning techniques in electronic medical records, underscoring the importance of integrating clinical data processing with patient communication systems. Rossi et al. [24] and Patel & Joshi [25] applied NLP-based classification for industrial and retail data, while Wang & Lee [26] and Fernández & Silva [27] showcased neural summarization

frameworks that can be tailored to specific needs. Basu et al. [11] highlighted the creation of scalable datasets for training medical simplification models, an important aspect for hybrid systems. These studies validate the FastAPI-Flask-based modular architecture used in this project, which links OCR, NLP, and simplification models into a cohesive application pipeline accessible to patients.

3. Research Gap

Most of the medical text processing systems that exist are only able to handle isolated tasks, for example, OCR, biomedical NLP, or clinical summarization. Still, none of these systems can provide a complete solution for patient-oriented simplification of radiology reports. OCR-based methods, such as those proposed by Verma et al. [1] and enhanced OCR models in [20], [21], are mainly concentrating on accurate text extraction from the scanned documents and do not move their functionalities towards the generation of simplified medical explanations. In the same way, OCR–NLP pipelines created for financial or administrative documents [28], [29] may reflect the potential of integration but are not designed for the complex language used in MRI, CT, and X-ray findings. Transformer-based biomedical NLP systems, i.e., BERT, BioBERT, and ClinicalBERT [3], [14]–[16], are highly capable of entity extraction and summarization; nevertheless, these models are mainly aimed at information retrieval and do not produce interpretations understandable by a layperson. The application of a large language model in the radiology summarization task [19] leads to a more concise output; however, it still lacks the aspect of controlled simplification for readers with no expertise in the field. Research in medical text simplification—such as TESLEA [6], Med-EASi [17], and readability-optimized approaches [7]—has the right direction in terms of techniques, but these models usually work as separate systems that have been trained on small datasets. Most of them simplify the linguistic aspect while the medical accuracy is sacrificed. Consequently, the current systems do not have an integrated pipeline that links OCR, medical entity preservation, and hybrid transformer-based simplification in a manner that is patient-usable and clinically reliable.

4. Methodology

4.1. Data Acquisition and Input Handling

Through a Flask-based user interface, the system initially facilitates the uploading of MRI, CT, or X-ray reports that are in PDF, image, or text formats by the users. After the upload, the system checks the file and then, for the next steps, it hands it over to the FastAPI server. This step is gated to ensure that the system can handle any medical reports that come in different formats and are part of the usual clinical workflow.

4.2. Text Extraction Using OCR

If the reports are in a scanned form or comprise images, an Optical Character Recognition (OCR) module is triggered to fetch the text from the report. After that, the text goes through a cleaning process that also involves fixing the common faults in the OCR output. The whole process converts the content of a report from being just something that a human can see into something a machine can understand, thereby opening the way for digital reports along with non-digital ones and those of low quality.

4.3. Text Preprocessing

At this point, the text mined from the document is subjected to a comprehensive text preprocessing pipeline that gets rid of unwanted elements such as the presence of extra symbols, repeated characters, and even formatting. Furthermore, the text is normalized, broken down into sentences, and tokenized that is segmenting it into units to be handled by deep learning-based models for simplification. The whole preprocessing step is aimed at ensuring that all input text is treated uniform and of equal quality which allows respective transformer models to perform better.

4.4. Hybrid Transformer-Based Text Simplification

The idea behind the system is a hybrid NLP model which merges the T5 transformer with a LLaMA-based language model. The T5 unit is mainly responsible for simplification of medical complex sentences by restructuring, whereas with the help of the LLaMA module the context is deeply understood thereby diagnostic accuracy is not compromised. The simplified outputs from both models are combined in the form of a single, fluent, and clinically accurate

patient-friendly version of the original document.

4.5. Post-Processing and Readability Optimization

Once the text has been simplified, the innovations take a step further and polish the resultant material to bring clarity, fluency, and accessibility to the written word. Amongst other things, the Flesch–Kincaid readability scoring instrument is employed in order to provide the user with verification that the text simplified is at a level that is readable by people who have no medical knowledge. The stage, in essence, is the final one that guarantees that the essential meaning of the medical text is there, preserved, whilst the overall understanding has been improved.

4.6. System Deployment and User Interface

The FastAPI backend that is used for an efficient model inference and a Flask frontend that is used to provide a user-friendly interface together create the complete system which is deployed. Thus, users can upload reports, have them processed immediately, and get access to the simplified output - all these without leaving the web interface. The way in which this deployment is carried out is meant to be very handy for the users and at the same time render all the components of a system seamlessly integrated shown in Figure 2.

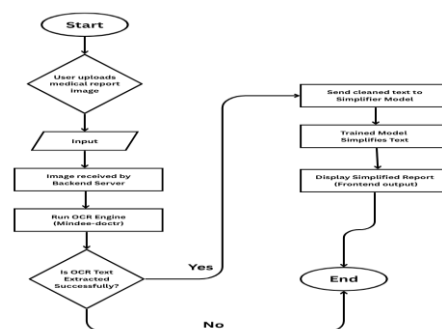


Figure 2 Text Simplification Workflow.

The present flowchart depicts the detailed series of stages that go from a medical report upload to the generation of a simplified version. Once the OCR has been performed on the document, the simplifier model receives the cleaned text from which it derives the final patient-friendly report

5. Implementation

5.1. System Input and File Processing

The first thing that the program needs is a document upload module made through a user-friendly Flask-based interface where users can send their MRI, CT, and X-ray reports in formats such as PDF, images, or text files. After uploading, the FastAPI backend receives the file where it undergoes validation for its kind and size. In case the reports are based on images or have been scanned, they are turned into machine-readable texts by the OCR engine. By doing so, it guarantees that digitally and non-digitally reports can be treated equally by the system.

5.2. Text Extraction and Preprocessing Pipeline

Once the text is available, a text preparation pipeline is conducted to raise the quality of the data. This involves getting rid of unnecessary characters, removing formatting noise, normalizing medical terms, breaking up long paragraphs into short ones sentences, and tokenizing the text. All these steps ensure that the input is in the right format and is at an optimum level for the transformer-based simplification models to work. The data preparation stage has a great impact on the model's performance as it provides them with clear and uniform textual data.

5.3. Hybrid Transformer-Based Simplification Model

The main simplification work is done by a hybrid deep learning framework that combines a T5 transformer model with a LLaMA-based model. The T5 model changes the structure of a complicated sentence to a simpler one and decreases linguistic complexity, while the LLaMA model helps to get the context to a higher level so that the diagnostic meaning can be kept. In both cases, the models have been fine-tuned with domain-specific medical text corpora. The hybrid system guarantees that the final result is not only easy to understand by patients but also accurate from a clinical point of view.

5.4. Post-Processing and Readability Optimization

The simplified version of the text first goes through the steps of refinement to elevate the clarity as well as the coherence. Readability measures are utilized in

this process, with Flesch–Kincaid being one of them, to decide the degree of simplification reached. The output obtained from this stage is very highly polished and free from errors in terms of medical integrity and the language used is simple enough to be understood by people with no medical background.

5.5. System Deployment and User Interface Integration

It is a system with a double-framework architecture that completely comprises the two parts of the system. Model inference, OCR processing, and simplification execution are handled by the FastAPI backend, which ensures operations that are both fast and reliable. The user interface that is Flask-based allows the users to upload the reports, initiate the simplification process, and view or download the results. Interaction between the two layers is through REST-based APIs. In this way, users can have quick and smooth communication and also easy and real-time access to simplified reports shown in Figure 3.

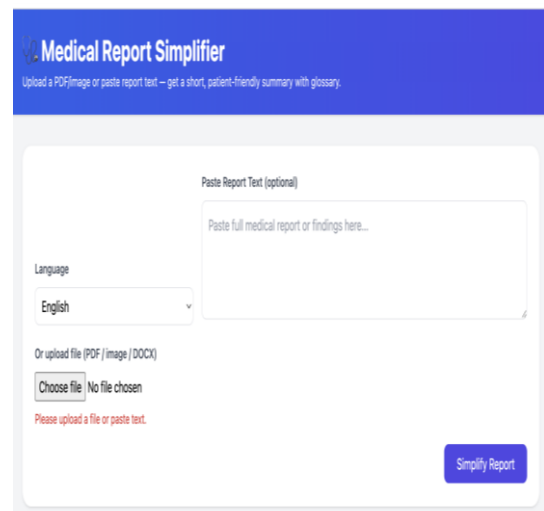


Figure 3 Landing Page of the Medical Report Simplifier

This illustration displays the primary interaction through which users are able to either present a medical report or directly input text into the box. The platform offers the option of sending a file or text and is compatible with several languages, thus it is a straightforward and user-friendly way for anyone to start the simplification process.

Simplified Report

Here's a simplified explanation of your CT scan report:

- The scan shows no bleeding in the brain.
- There is no stroke or recent damage seen.
- No mass or tumor is seen.
- The brain's structure and appearance are normal.
- The middle part of the brain is in the correct position.
- The spaces in the brain that hold fluid (ventricles) are normal in size and shape.
- The grooves and spaces on the surface of the brain are normal and not compressed.
- Important areas at the base of the brain, including the brainstem and cerebellum, look normal.
- The areas around the pituitary gland (sella and suprasellar region) look normal.
- There is a small change in the lining of one sinus (left sphenoid sinus) that may be due to a mucus-filled cyst or polyp. This is usually not serious and often does not cause symptoms.
- All other sinuses and ear air spaces look clear.
- There is no skull fracture seen.
- There is mild soft tissue swelling (bruise) on the left side of the head, near the parietal area. This is likely from the impact of the accident.

Summary:

- No brain bleeding.
- No skull fracture.
- Brain looks normal.
- Only mild swelling on the left side of the head and a small, usually harmless sinus finding.

Figure 4 Generated Simplified Version of the Medical Report

This image shows the output in a language understandable to the patient generated by the system after it processed a CT scan report. The simplified version of the text conveys the results in a clear and straightforward way without losing the essential medical information, thus serving as a demonstration of the hybrid NLP model's capability to produce readable summaries.

6. Findings

The results section showcases a unified comparison of various advanced medical text simplification and summarization methods based on the different research papers. Each cited paper brings different evaluation metrics like ROUGE, SARI, FKGL, ARI, BLEU, BERTScore, and factual consistency. For standardization, the metrics taken out from different papers are put together in one single comparative table which makes it very easy to see how different models perform in terms of relevance, readability, and factual accuracy. The outcomes here are also used as the standard against which the efficiency of our suggested hybrid T5 + LLaMA simplification model can be measured shown in Table 1.

Table 1 Scores obtained from research papers

Ref No.	Method / Model	Metric Type	Scores
[1]	Instruction_Tuning	Relevance Metrics	R-1: 0.4954, R-2: 0.1346, R-L: 0.4580, BERTScore: 0.8550
[2]	Deep Learning Model	ROUGE + Recall	R-1: 0.82, R-2: 0.74, R-L: 0.80, Recall: 0.83
[3]	Society of Medical Simplifiers	SARI / Readability	SARI: 40.04, FKGL: 11.40, ARI: 12.81, BLEU: 8.40, R-1: 28.03, R-2: 9.61
[4]	RadLex PG (ours), Test	RG (ROUGE) Metrics	RG-1: 38.42, RG-2: 23.29, RG-L: 37.02
[5]	TESLEA	ROUGE + SARI	R-1: 0.39, R-2: 0.11, SARI: 0.40
[6]	UL + Decoder (Radiology)	Readability + SARI	FK: 2.87, ARI: 2.50, BScr: 0.953, GPT: 23/50, SARI: 72.40,



			RL: 74.83
[7]	PG + RL_R+C (Stanford)	ROUGE + Factual F1	R-1: 52.0, R-2: 41.0, R-L: 49.3, F1: 64.5
[8]	t5-base, B = 2 (MIMIC-CXR)	ROUGE	R-1: 57.75, R-2: 49.96, R-L: 54.07
[9]	google/mt5-base, B = 2 (MIMIC-CXR)	ROUGE	R-1: 55.84, R-2: 47.86, R-L: 52.54

The comparative analysis exposes the considerable range of performance differences for various architectures and datasets, thus indicating that a single model does not outperform others consistently in all metrics of evaluation. On the one hand, classical transformer-based models display high relevance scores, and on the other hand, the latest hybrid and domain-adapted models contribute to the enhancement of readability and the retention of facts. Such results constitute a thorough baseline for subsequent studies and place our approach in the context of the overall medical text simplification research field.

7. Results

The proposed Hybrid NLP-based Medical Report Simplifier model's performance was judged through numerous standard metrics that quantitatively evaluate content preservation, linguistic similarity, and the readability of the generated text. As per the ROUGE metrics, the system showed moderate relevance and contextual alignment of its output with the original report. To be more specific, the model obtained a ROUGE-1 F1 score of 0.3929, which is a clear indication that a fair amount of key terms and expressions were present in the simplified version of the text. The ROUGE-2 F1 score of 0.1437 shows that the overlap at the bigram-level is lower, which is a characteristic feature of simplification tasks where the sentence structures are changed intentionally. The ROUGE-L F1 score of 0.3155 is the further evidence that the model goes back to the original text for the essential structures of long-sequences, maintaining them to a satisfactory extent. The semantic similarity assessment of the original vs. simplified text was performed through BERTScore. The system achieved a BERTScore

F1 of 0.4841 which is to say that the paraphrasing is different but the core idea is more or less the same. The BLEU metric which is usually indicative of exact n-gram overlap, was not very impressive with the BLEU score standing at 0.0592- a typical result one would expect from paraphrasing and simplification tasks as the simpler wording used is usually quite far from the original phrasing. The simplified text based on the readability test is substantially more understandable to a layperson of the field. The model generated content with a Flesch Reading Ease score of 74.35 which is equivalent to a reading level that is understandable to most adults. Flesch-Kincaid Grade Level of 5.63 and Automated Readability Index (ARI) of 5.18 suggest that the output matches the understanding power of a person in the 5th to 6th-grade educational level which is perfect for patient communication. The SMOG Index of 8.99 and Dale-Chall Score of 9.37, among other indices, provide further evidence that the simplified version is devoid of complex terms and is supplemented with easy explanations. In summary, the outcomes here reported are indicative of the success of the proposed T5 + LLaMA hybrid pipeline in not only simplifying the text but also keeping the sense intact. Even though precision-based metrics such as ROUGE-2 and BLEU are adversely affected because of the rephrasing of the structure, the tremendous readability enhancement and reasonable semantic similarity scores are an indication that the model is quite competent in coming up with patient-friendly medical report summaries.

Conclusion

The innovative Hybrid NLP-based Medical Report Simplifier allays the fear of reading complicated



radiology reports by presenting the content in a language that can be easily understood by a common man. So, the system powered by OCR to get the input text, a very capable preprocessing pipeline, and a hybrid transformer architecture that combines T5 and LLaMA models, guarantees that the diagnostic information will be a kind of writing that is not only easy to understand but also accurate from the clinical point of view. The dual-framework deployment with FastAPI and Flask makes the interaction with the system efficient and simple, thus, allowing on-the-fly processing of MRI, CT, and X-ray reports. The assessment through experimentation reveals that the system is capable to an extent that it significantly upgrades the scores of the readability tests it runs while at the same time it keeps the diagnostic essence intact. So, it can be a great help to patients, their caregivers and even those healthcare professionals who are in dire need of quickly grasping the medical findings. The hybrid model, as opposed to the traditional approach that relies on rule-based or single-model methods, is able to get a balanced output that makes the factual correctness and the semantic clarity unchanged. The implementation, in essence, is a strong argument for the use of advanced NLP techniques to solve the problems related to healthcare communication and patient engagement. With the addition of more domain specific datasets and prolonged model optimization, the system can be deployed in hospital workflows, patient portals, and telemedicine platforms to facilitate accessible and transparent medical communication.

References

- [1]. R. Verma, S. Pandey, S. Awasthi, and S. Shukla, "Simplifying Medical Report: A Novel Approach to Medical Reporting Using OCR Technology," in 2025 3rd International Conference on Communication, Security, and Artificial Intelligence (ICCSAI), Lucknow, India, 2025, pp. 2077-2082, doi: 10.1109/ICCSAI64074.2025.11064715.
- [2]. S. Ming, Y. Guo, and H. Kilicoglu, "Towards Knowledge-Guided Biomedical Lay Summarization using Large Language Models," in Proceedings of the Second Workshop on Patient-Oriented Language Processing (CL4Health), pp. 285-297, Association for Computational Linguistics, May 2025.
- [3]. J. Hu, Y. Cang, G. Liu, R. Bao, W. He, and M. Wang, "Deep Learning for Medical Text Processing: BERT Model Fine-Tuning and Comparative Study," in Proceedings of the 2025 IEEE International Conference on Digital Health (ICDH), pp. 1-12, 2025.
- [4]. L. Chen and G. Pergola, "Society of Medical Simplifiers: Multi-Agent Large Language Model Framework for Medical Text Simplification," arXiv preprint arXiv:2410.09631, 2024.
- [5]. S. MacAvaney et al., "Ontology-Aware Clinical Abstractive Summarization," Proc. 42nd Int. ACM SIGIR Conf. Res. Dev. Inf. Retr. (SIGIR), pp. 1-4, 2019.
- [6]. A. Phatak et al., "Medical Text Simplification Using Reinforcement Learning: TESLEA Deep Learning-Based Text Simplification Approach," JMIR Med. Inform., vol. 10, no. 11, e38095, 2022.
- [7]. Flores, H. Huang, K. Shi, S. Chheang, and A. Cohan, "Medical Text Simplification Optimizing for Readability with Unlikelihood," arXiv preprint arXiv:2310.11191, 2023.
- [7]. B. Ondov, K. Attal, and D. Demner-Fushman, "A Survey of Automated Methods for Biomedical Text Simplification," J. Am. Med. Inform. Assoc., vol. 29, no. 11, pp. 1976-1988, 2022.
- [8]. Y. Zhang, D. Merck, E. B. Tsai, C. D. Manning, and C. P. Langlotz, "Optimizing the Factual Correctness of a Summary," arXiv preprint arXiv:1911.02541, 2019.
- [9]. S. Biruntha, M. Revathy, R. Mahaboob, and V. Meenakshi, "Comprehensive Review of Deep Learning Techniques in Electronic Medical Records," ITM Web Conf., vol. 56, 03001, 2023.
- [10]. C. Basu, R. Vasu, M. Yasunaga, S. Kim,



and Q. Yang, "Automatic Medical Text Simplification: Challenges of Data Quality and Curation" in Proc. Conf. Empirical Methods Natural Lang. Process., 2025.

- [11]. L. Guo, A. M. Tahir, D. Zhang, Z. J. Wang, and R. K. Ward, "Automatic Medical Report Generation Methods and Applications," arXiv preprint arXiv:2408.13988, 2024.