© FEB 2026 | IRE Journals | Volume 9 Issue 8 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I8-1714021

Radiology Report Analysis Using Deep Learning and NLP

PRIYADARSHINI S[1], SIVA VENKATESH M[2], EZHIL AADHITHYAN K[3], DARSAN KARTHIC P D[4]
[1] *Department of Information Technology, KPR Institute of Engineering and Technology*
[2, 3] *Department of Artificial Intelligence and Data Science, KPR Institute of Engineering and Technology*
[4] *Department of Biomedical Engineering, KPR Institute of Engineering and Technology*

***Abstract-* Radiology reports contain important diagnostic information but are commonly written in unstructured textual form. This makes automated processing and analysis difficult. This paper presents a framework for Radiology Report Analysis Using Deep Learning and Natural Language Processing (NLP) to extract relevant medical entities and generate concise summaries from radiology reports. The proposed system applies biomedical Named Entity Recognition for identifying anatomical structures and pathological conditions, followed by ontology-based mapping for term standardization. A transformer-based summarization model is used to generate brief clinical summaries. To address privacy concerns, a de-identification module is included to remove sensitive patient information. In addition, a basic consistency check is performed between extracted findings and generated impressions. Experimental observations indicate that the framework can assist in organizing and summarizing radiology report content, supporting clinical review processes.**

***Index Terms-* Radiology Reports, Medical NLP, Deep Learning, Clinical Text Analysis.**

I. INTRODUCTION

A. Background and Motivation

Radiology departments generate a large volume of diagnostic reports on a daily basis. These reports contain valuable clinical information related to anatomical structures, pathological conditions, and diagnostic impressions. Despite their importance, radiology reports are typically stored as unstructured text, limiting their usability for automated analysis, large-scale research, and clinical decision support systems. Manual interpretation of these reports is time-consuming and prone to variability across radiologists.

Recent advances in deep learning and NLP have shown strong potential in analyzing medical text data. Transformer-based models and biomedical NLP frameworks enable automated extraction of key entities and summarization of long clinical narratives. However, challenges remain in ensuring clinical accuracy, explainability, and patient privacy. There is a need for intelligent systems that not only summarize radiology reports but also validate findings, maintain consistency, and provide transparent explanations.

B. Objectives

The main objectives of this research are:

- To design an automated system for extracting clinically relevant entities from radiology reports
- To generate concise and structured summaries of radiology findings
- To validate consistency between extracted findings and generated impressions
- To ensure patient privacy through effective de-identification
- To provide an explainable and interactive interface for clinical users

II. METHODOLOGY AND SYSYTEM

DESIGN

A. System Overview

The proposed system follows a modular pipeline consisting of text preprocessing, entity extraction, ontology mapping, summarization, consistency validation, and visualization. The system accepts radiology reports in multiple formats, including manual text input, CSV datasets, and PDF documents. Extracted information is stored and

© FEB 2026 | IRE Journals | Volume 9 Issue 8 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I8-1714021

reused across different analysis modules through session-based state management.

B. Entity Extraction and Ontology Mapping

Biomedical named entity recognition is performed using a domain-specific NLP model trained on clinical text. The system identifies entities such as anatomical structures, diseases, pathological formations, and diagnostic procedures. Extracted entities are normalized and mapped to standardized biomedical ontology categories to improve semantic consistency and clinical interpretability. Noise removal and entity normalization rules are applied to eliminate non-medical and redundant terms.

C. Summarization and Consistency Validation

A transformer-based text summarization model is employed to generate concise summaries from lengthy radiology reports. The system further validates consistency between extracted findings and generated summaries by comparing clinical concepts present in both sections. The output is classified as consistent, partially consistent, or inconsistent, providing an additional quality assurance layer. Structured radiology-style impression bullets are also generated based on severity and clinical relevance.

D. Privacy Preservation

To ensure patient confidentiality, an enhanced de-identification pipeline is integrated into the system. Personally identifiable information such as patient names, dates, identifiers, and institutional details is automatically masked using rule-based and pattern-based techniques. This ensures compliance with privacy standards while allowing safe data analysis and sharing.

III. IMPLEMENTATION DETAILS

The proposed system for Radiology Report Analysis Using Deep Learning and NLP is implemented as a modular and scalable pipeline that integrates text preprocessing, medical entity extraction, summarization, evaluation, and explainability within a single framework.

Radiology reports are accepted through multiple input formats, including manual text input, CSV datasets, and raw PDF reports. For PDF-based reports, text is extracted using a document parsing module to ensure accurate recovery of clinical content. The extracted text undergoes preprocessing, which includes normalization of spacing and removal of irrelevant artifacts.

To ensure patient privacy, an enhanced de-identification module is applied before analysis. This module uses rule-based pattern matching to anonymize sensitive information such as patient names, dates, identification numbers, contact details, and institutional identifiers. This step ensures compliance with privacy-preserving principles while maintaining clinical relevance.

Named Entity Recognition (NER) is performed using a biomedical language model trained on clinical corpora. The system identifies key medical entities such as anatomical structures, diseases, pathological formations, and diagnostic procedures. Extracted entities are further cleaned and normalized to remove non-clinical tokens and redundant terms. Each entity is then mapped to a standardized biomedical ontology using similarity-based matching, improving semantic consistency across reports.

A transformer-based summarization model is employed to generate concise and context-aware summaries of radiology findings. The model is fine-tuned to prioritize clinically important observations while preserving interpretability. Additionally, a structured impression module converts extracted findings into radiology-style bullet points ordered by clinical severity, mimicking standard reporting practices.

The entire pipeline is deployed through an interactive web-based interface built using Streamlit. This interface allows users to analyze reports, visualize extracted entities, explore ontology relationships through knowledge graphs, and interact with explainability components in real time.

© FEB 2026 | IRE Journals | Volume 9 Issue 8 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I8-1714021

IV. EVALUATION STRATEGY AND MATRICS

To assess the effectiveness and reliability of the proposed system, both quantitative and qualitative evaluation strategies are employed. The evaluation focuses on summary quality, entity coverage, and consistency between extracted findings and generated impressions.

Summary quality is evaluated using ROUGE-based metrics, which measure lexical overlap between the generated summary and the original report text. Although radiology reports do not always have gold-standard summaries, ROUGE scores provide a useful approximation of content preservation and relevance.
Entity-level performance is measured using entity recall, which represents the proportion of extracted medical entities that are reflected in the generated summary. A higher recall indicates that clinically significant findings are effectively captured by the summarization model.

To ensure clinical coherence, a finding–impression consistency validation mechanism is introduced. This module checks whether key pathological and anatomical concepts identified in the findings section are appropriately represented in the generated impression. Based on this analysis, the system categorizes reports as consistent, partially consistent, or inconsistent, providing an additional layer of quality assurance.

Explainability is incorporated through entity importance scoring and visualization techniques. Entity importance is computed based on frequency and contextual relevance, and results are presented using graphical representations to help users understand which findings most influenced the model output.

Finally, a radiologist-in-the-loop feedback mechanism allows clinical experts to review, edit, and comment on AI-generated impressions. This feedback not only supports qualitative validation but also lays the groundwork for future supervised learning and continuous system improvement.

V. SYSTEM ARCHITECTURE AND WORKFLOW ANALYSIS

The overall architecture of the proposed radiology report analysis framework is designed to be modular, scalable, and clinically interpretable. The system follows a sequential workflow that transforms unstructured radiology text into structured, meaningful, and explainable outputs suitable for clinical review.

The first stage of the workflow handles data ingestion. Radiology reports can be provided in multiple formats, including free-text input, CSV datasets, and PDF documents. This flexibility allows the system to be applied across different clinical data sources. For PDF inputs, a document parsing module extracts readable text while preserving the original clinical context.

The second stage focuses on preprocessing and privacy preservation. Text normalization is performed to remove unnecessary formatting artifacts and ensure consistent input for downstream models. An enhanced de-identification module anonymizes patient-specific information such as names, dates, identifiers, and institutional references. This step ensures that sensitive information is protected without affecting the medical content required for analysis.

In the third stage, biomedical Named Entity Recognition is applied to identify clinically relevant entities, including anatomical structures, diseases, pathological formations, and diagnostic procedures. These entities are then normalized and mapped to standardized ontology categories to reduce semantic ambiguity and improve interpretability.

The fourth stage applies transformer-based summarization to generate concise clinical summaries. This is followed by a consistency validation step, which checks alignment between extracted findings and generated impressions. Finally, visualization and explainability components present the results in an interactive format, completing the end-to-end workflow.

VI. EXPLAINABILITY AND TRANSPARENCY CONSIDERATIONS

Explainability is a critical requirement for artificial intelligence systems used in healthcare. Clinicians must be able to understand how and why a model produces specific outputs before trusting it in real-world scenarios. To address this requirement, the proposed system integrates multiple explainability mechanisms.

Entity-level explainability is achieved by analyzing the frequency and relevance of extracted medical entities. Entities that appear frequently or are strongly associated with key diagnostic concepts are assigned higher importance scores. These scores help users identify which findings played a significant role in shaping the generated summary.

Additionally, ontology-based grouping allows clinicians to view extracted entities under meaningful clinical categories such as anatomical structures and disease conditions. This structured representation reduces cognitive load and improves interpretability.

The system also incorporates a finding–impression consistency check. By highlighting missing or under-represented findings, the system provides transparency into potential limitations of the generated summary. This feature supports clinical validation and reinforces user trust in the AI-assisted analysis.

VII. CLINICAL APPLICABILITY AND USE CASES

The proposed radiology report analysis framework has several practical applications in clinical and research settings. One primary use case is assisting radiologists by automatically summarizing lengthy reports, allowing them to focus on critical diagnostic decisions rather than manual documentation.

The system can also support clinical audits and quality assurance by validating consistency between findings and impressions. In teaching hospitals, the framework can be used as an educational tool to help trainees understand standard reporting structures and common diagnostic patterns.

From a research perspective, structured extraction of entities enables large-scale analysis of radiology reports for epidemiological studies and outcome analysis. The privacy-preserving design further allows safe sharing of anonymized data for collaborative research.

Overall, the system serves as a decision-support aid rather than a replacement for radiologists, enhancing efficiency and consistency in radiology workflows.

VIII. RESULT AND FINDINGS

The system was evaluated using sample radiology reports from different anatomical regions. The results show effective extraction of clinically meaningful entities and generation of concise summaries that preserve key diagnostic information. The consistency validation module successfully identified missing or under-represented findings in summaries. Evaluation metrics such as ROUGE scores and entity recall indicate satisfactory performance for automated summarization. The explainability dashboard highlights influential entities, improving transparency and trust in the model's outputs. Overall, the system demonstrates reliable performance in supporting radiology report analysis.

IX. LIMITATIONS AND CHALLENGES

Despite promising results, the proposed system has certain limitations. The accuracy of entity extraction and summarization depends on the quality and variability of radiology report language. Reports with uncommon abbreviations or highly narrative styles may affect model performance.

The current framework operates only on textual reports and does not analyze radiological images directly. As a result, the system relies entirely on the quality of the written report. Additionally, ontology mapping is limited to predefined categories and may not fully capture rare or emerging medical terms.

Another challenge is the absence of gold-standard summaries for evaluation, which limits the scope of quantitative assessment. These limitations highlight areas for improvement in future work.

X. FUTURE EXTENSIONS AND RESEARCH DIRECTIONS

Future work will focus on extending the system's capabilities and improving clinical robustness. One key enhancement is the integration of Retrieval-Augmented Generation (RAG) to allow the system to answer clinical queries using both report content and external medical knowledge sources.

Another important direction is multimodal learning, where radiology images are combined with textual reports to provide more comprehensive analysis. Continuous learning using radiologist feedback will also be explored to refine entity extraction and summarization quality.

Large-scale clinical validation studies will be conducted to assess system performance across different institutions and reporting styles. These future enhancements aim to transform the proposed framework into a reliable and scalable AI assistant for radiology departments.

XI. ETHICAL AND REGULATORY CONSIDERATIONS

Ethical use of artificial intelligence in healthcare requires strict adherence to privacy, transparency, and accountability standards. The proposed system addresses these concerns through built-in de-identification, explainability, and human-in-the-loop validation mechanisms.

The system is designed to assist clinicians rather than replace expert judgment. All outputs are intended to be reviewed by qualified medical professionals before clinical use. Compliance with data protection principles ensures that patient confidentiality is maintained throughout the analysis process.

XII. CONCLUSION

This paper presents an end-to-end system for radiology report analysis using deep learning and NLP techniques. The proposed approach successfully combines entity extraction, summarization, consistency validation, and privacy preservation within an interactive and explainable framework. The system reduces manual effort, improves report interpretability, and supports clinical decision-making. Future work will focus on integrating retrieval-augmented generation, expanding ontology coverage, and incorporating radiologist feedback for continuous model improvement.

ACKNOWLEDGMENT

The authors would like to express their sincere gratitude to the first author and project mentor for continuous guidance, technical direction, and valuable feedback throughout the course of this work. The authors also thank the faculty members of the Department of Artificial Intelligence and Data Science, KPR Institute of Engineering and Technology, for their support and encouragement. In addition, the authors acknowledge the valuable clinical insights provided by radiology professionals, which helped align the proposed system with practical diagnostic workflows.

REFERENCES

[1] Devlin, J., Chang, M. W., Lee, K., and Toutanova, K., "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," *NAACL*, 2019.

[2] Lee, J. et al., "BioBERT: A Pre-trained Biomedical Language Representation Model," *Bioinformatics*, 2020.

[3] Demner-Fushman, D., Chapman, W. W., and McDonald, C. J., "What can natural language processing do for clinical decision support?" *Journal of Biomedical Informatics*, 2009.

[4] Vaswani, A. et al., "Attention Is All You Need," *NeurIPS*, 2017.

[5] Johnson, A. E. W. et al., "MIMIC-III, a freely accessible critical care database," *Scientific Data*, 2016.

[6] Irvin, J., Rajpurkar, P., Ko, M., Yu, Y., Ciurea-Ilcus, S., Chute, C., … Ng, A. Y. (2019). *CheXpert: A large chest radiograph dataset with uncertainty labels and expert comparison.* Proceedings of the AAAI Conference on Artificial Intelligence, 33, 590–597.

[7] Wu, S., Roberts, K., Datta, S., Du, J., Ji, Z., Si, Y., … Xu, H. (2020). *Deep learning in clinical natural language processing: A methodical review.* Journal of the American Medical Informatics Association, 27(3), 457–470.

[8] Zech, J. R., Badgeley, M. A., Liu, M., Costa, A. B., Titano, J. J., & Oermann, E. K. (2018). *Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: A cross-sectional study.* PLoS Medicine, 15(11), e1002683.