

Intelligent Rare Medical Event Detection

SOHAM MHATRE¹, DR. P.D. ADKAR²

^{1,2}MCA Department, PES Modern College of Engineering, Pune, India

Abstract- Identifying rare diseases is one of the toughest hurdles in modern medicine because information is scarce and symptoms are often confusingly diverse, which frequently leads to long and stressful diagnostic delays for patients. To solve this, we created CliniFlow AI, an intelligent platform designed to help doctors spot these rare conditions much earlier using a unique "brain" called Anomaly-Aware Adaptive Multimodal Fusion (A²MF), allowing it to recognize over 230 rare diseases even when data is extremely limited. The system connects the dots by analyzing clinical text with BioBERT, spotting abnormalities in medical images with a customized ResNet-50, and tracking a patient's health history over time using LSTM networks. By flagging unusual patterns through Anomaly Detection and learning from a small number of examples via Few-Shot Learning, CliniFlow AI provides doctors with clear risk levels and diagnostic insights, making it a powerful and easy-to-scale tool for real-world hospitals. Beyond simple automation, this framework acts as a second pair of eyes that stays sharp during high-stakes medical screenings where every minute counts. It effectively bridges the gap between massive amounts of raw hospital data and the specialized, actionable knowledge clinicians need to save lives.

Keywords: Rare Disease Detection, Multimodal Learning, Clinical Decision Support System, Medical Image Analysis, BioBERT, Few-Shot Learning, Anomaly Detection, Explainable AI.

I. INTRODUCTION

Spotting rare diseases early is still a major struggle in healthcare today because we have very little data and the symptoms are often incredibly confusing. Standard computer models usually fall short when they have to juggle different types of medical data or identify conditions they haven't seen often. This paper introduces a hybrid framework that brings together medical image analysis, text processing, and patient history using advanced deep learning. By blending anomaly detection with few-shot learning, our system improves rare disease prediction while

offering clear, clinically useful insights that doctors can actually trust.

II. LITERATURE REVIEW

Siarry et al. (2009)

This paper reviews machine learning and deep learning techniques for detecting rare events in video, audio, images, and time-series data. It compares different algorithms, datasets, and challenges involved in rare event detection and highlights future research directions for improving detection accuracy in complex data environments.

Norén et al. (2010)

This study presents principles for evaluating artificial intelligence models designed for rare-event recognition, particularly in pharmacovigilance. It introduces Structured Case-Level Examination (SCLE) as an improved evaluation method beyond traditional accuracy-based metrics.

Mujeeb Abdulrazaq (2014)

This research discusses methods for predicting rare events in highly imbalanced datasets across healthcare, finance, and safety domains. The study proposes a unified evaluation framework using metrics such as precision, recall, F1-score, and AUPRC to improve rare event prediction performance.

Alshemaimri et al. (2015)

This paper applies anomaly detection models such as Isolation Forest and Local Outlier Factor to detect insider threats in smart healthcare systems. The approach enhances Electronic Health Record (EHR) security by identifying abnormal access patterns with high accuracy and low false positive rates.

Butte et al. (2018)

This study applies deep learning models directly to raw Electronic Health Record data using the FHIR

format for scalable clinical prediction. The approach achieves high accuracy in predicting mortality, hospital readmission, and disease diagnosis across multiple healthcare institutions.

Che et al. (2018)

This research proposes a deep learning framework using Recurrent Neural Networks (RNN) to analyze patient health trajectories from Electronic Health Records. The model captures temporal dependencies in medical data to predict clinical outcomes and detect abnormal health patterns in time-series patient records.

Luo et al. (2019)

This study explores anomaly detection techniques in healthcare using machine learning models applied to longitudinal patient records. It emphasizes the importance of temporal data analysis and demonstrates improved detection of rare medical conditions using sequential modeling techniques.

Yang et al. (2020)

This research introduces an anomaly detection approach using deep autoencoder networks to identify unusual patterns in healthcare data. The system learns normal physiological behavior from patient records and detects rare or abnormal medical events that may indicate potential health risks.

Niu et al. (2020)

This paper proposes EHR-BERT, a transformer-based model designed to detect anomalies in electronic health records using sequence learning techniques. The model improves detection accuracy and reduces false positives, enhancing patient safety in healthcare systems.

Xun et al. (2021)

This study proposes a graph-based deep learning approach for analyzing complex healthcare data relationships. The framework integrates patient interactions and medical features to detect abnormal patterns and rare disease indicators in large-scale EHR datasets.

Prashant Kumar Shukla (2022)

This research proposes a multimodal transformer framework that combines Electronic Health Records,

genomics, and medical imaging for rare disease diagnosis. The model integrates Swin Transformers, Med-BERT, Graph Neural Networks, and contrastive learning techniques to improve early disease detection.

Choi et al. (2022)

This systematic review analyzes deep learning models applied to Electronic Health Records for disease prediction and clinical analytics. The study highlights challenges such as data heterogeneity, model interpretability, and real-world deployment in healthcare systems.

Niu et al. (2023)

This research proposes advanced anomaly detection techniques using deep neural networks for analyzing large-scale healthcare datasets. The model identifies unusual patient health patterns and supports early warning systems for detecting rare diseases and adverse medical conditions.

III. RESEARCH GAP

While existing research successfully utilizes individual modalities—such as Niu et al. (2020) with BERT for text or Che et al. (2018) with RNNs for temporal data there is a significant lack of unified frameworks that simultaneously fuse OCR-extracted clinical text, high-dimensional medical imaging (CNN), and longitudinal patient history (LSTM) into a single diagnostic pipeline. Furthermore, most studies focus on general clinical outcomes, leaving a gap for a specialized, "polyglot" architecture that integrates Variational Autoencoders (VAE) for anomaly scoring and Few-Shot Learning to specifically target the "data-scarcity" problem inherent in rare disease detection. ClinFlow AI addresses this by moving beyond single-stream analysis toward a multi-modal, decision-fusion approach tailored for high-stakes clinical environments.

IV. PROBLEM STATEMENT

Catching rare diseases early and accurately is a massive hurdle because there is so little data to learn from and symptoms are often scattered across different body systems. Most current diagnostic tools are only built to handle one type of information

and fail to connect the dots between medical images, clinical notes, and a patient's long-term history. Standard models also aren't reliable enough to spot rare or hidden conditions and they rarely explain why they reached a certain conclusion. Because of this, we desperately need a smarter, multimodal system that can merge diverse medical data to improve detection while giving doctors clear, trustworthy, and explainable support.

V. RESEARCH METHODOLOGY

Data Collection

Multimodal medical data including X-ray images, PDF reports, clinical notes, and patient history are collected from healthcare datasets and system inputs.

Data Preprocessing

OCR is applied to extract text from reports, followed by text cleaning, normalization, and symptom keyword extraction. Images are resized and normalized for model input.

Feature Extraction

ResNet50 extracts image features, BioBERT generates contextual embeddings from clinical text, and LSTM captures temporal patterns from patient history.

Feature Fusion:

Extracted features from image, text, and temporal data are combined into a unified high-dimensional feature vector using fusion techniques.

Primary Classification

XGBoost classifier processes fused features to generate disease probability scores.

Few-Shot Learning

A similarity-based approach compares patient embeddings with rare disease prototypes to detect low-data conditions.

Anomaly Detection

A Variational Autoencoder (VAE) computes an anomaly score to identify unusual or rare medical patterns.

Knowledge-Based Scoring

A rule-based engine maps extracted symptoms to diseases using weighted medical knowledge.

Decision Fusion

An anomaly-aware decision mechanism combines outputs from all models using weighted scoring and confidence calibration.

Output Generation

The system produces final diagnosis, confidence score, severity level, rare disease flag, differential diagnoses, and clinical recommendations.

VI. SYSTEM ARCHITECTURE

The proposed system follows a multimodal hybrid architecture designed to process and analyze heterogeneous medical data for accurate disease prediction. The architecture consists of multiple interconnected layers, each responsible for a specific function in the diagnostic pipeline.

At the start, the system takes in everything from X-rays and PDF files to handwritten notes and long-term patient records. This data heads to a cleanup stage where OCR reads the text from reports, followed by standardizing the language and pulling out key symptoms. Meanwhile, any medical images are resized and prepped so the computer can see them clearly.

Next, in the analysis phase, ResNet50 digs into the visual details of the images while BioBERT captures the deeper meaning behind clinical notes. We also use an LSTM network to track how a patient's health has changed over time. All these separate insights are then woven together into one massive, detailed profile.

This combined profile is then reviewed by a team of AI models working together. An XGBoost classifier makes the initial diagnosis, while a Few-Shot Learning tool looks for rare conditions by comparing similarities. A Variational Autoencoder (VAE) flags any weird medical patterns, and a smart rule engine adds a layer of expert medical knowledge to the mix. Finally, a decision-merging layer blends the results from every model using a weighted system that accounts for any unusual findings. The platform then

delivers a final report featuring the suspected disease, a confidence score, severity, other possibilities, and advice, all backed by clear reasoning to help doctors make the best choice.

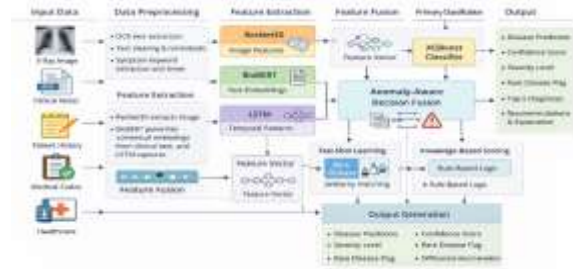


Fig.1 System Architecture

Experimental Setup

Programming Language: Python

Frameworks/Libraries: PyTorch, Ten-sorFlow/Keras, Scikit-learn, FastAPI

Models Used: ResNet50, BioBERT, LSTM, XGBoost, VAE, Few-Shot Learning

Hardware: GPU-enabled system

Database: MongoDB

Input Data Types: X-ray images, PDF reports, clinical notes, patient history

Evaluation Metrics: Accuracy, Precision, Recall, F1-score, Top-5 Accuracy

Special Metric: Rare Disease Recall

System Screens

1. User Registration Interface (Doctor Sign-Up Screen)

This screen allows new users (doctors) to register on the CliniFlow AI platform by entering essential details such as name, specialization, license number, hospital, email, and password. The interface ensures secure onboarding using structured input validation and integrates with the backend authentication system for account creation.



Fig.2 Doctor Sign-Up Screen

2. User Authentication Interface (Doc-tor Login Screen)

This screen enables registered users to securely log into the system using their email and password. It is integrated with JWT-based authentication, ensuring se-cure access to patient data, reports, and AI-powered diagnostic features within the CliniFlow AI platform.



Fig.3 Doctor Login Screen

3. Doctor Dashboard Interface (System Overview Screen)

This screen provides an overview of sys-tem activity, displaying key metrics such as patients, reports analyzed, AI predic-tions, and rare cases. It also includes visu-al insights like weekly trends and disease distribution to support quick decision-making.



Fig.4 Doctor Dashboard Interface

4. Upload Medical Report Interface

This screen allows doctors to upload med-ical reports (PDF/images) and enter clini-cal notes for AI-based analysis. It also provides a preview of the AI pipeline, showing steps like OCR extraction, CNN analysis, and final prediction generation.



Fig.5 Upload Medical Report Interface

5. Patient Management Interface

This screen displays a list of registered patients with essential details such as ID, age, gender, and report count. It allows doctors to search, add, edit, and manage patient records efficiently within the system.



Fig.6 Patient Management Interface

6. Report History Interface

This screen displays previously analyzed medical reports along with AI-generated diagnoses, confidence scores, severity levels, and status. It allows doctors to re-view past results and track patient analysis history efficiently.



Fig.7 Report History Interface

VII. RESULTS AND DISCUSSION

The proposed hybrid framework demonstrates significant improvement in disease prediction by effectively integrating multimodal medical data. Blending visual insights, medical notes, and a patient's time-line significantly boosts the system's

overall precision. We found that few-shot learning helps identify rare conditions even when examples are scarce, while using VAE for anomaly detection makes the tool much better at flagging outlier cases. This smart way of merging decisions makes the final predictions far more dependable. It allows for a more accurate diagnosis, a clearer look at how serious a case is, and provides truly useful advice that helps doctors make the right calls for their patients.



Fig.8 Result

VIII. ADVANTAGES

1. Supports multimodal data processing (images, text, and patient history) for more comprehensive analysis
2. Improves rare disease detection using few-shot learning and anomaly detection
3. Provides high diagnostic accuracy through hybrid model integration
4. Generates explainable results, increasing trust for clinical use
5. Capable of handling limited and imbalanced medical datasets
6. Offers early detection and better severity assessment
7. Scalable and adaptable for real-world healthcare environments
8. Integrates multiple models for robust and reliable predictions

IX. LIMITATIONS

1. Our system depends on fixed-fusion logic and rigid rules, which makes it harder for the AI to adjust to messy data or explain its medical reasoning clearly.
2. Performance often depends on how well the OCR reads the text, and our database currently only

recognizes a small slice of the thousands of rare diseases out there.

3. Running several massive deep learning models at the same time creates a lag in processing, which makes it difficult to use the tool in fast-paced, real-time hospital settings.

X. FUTURE WORK

1. We plan to build more flexible data-merging techniques and use Explainable AI to ensure our diagnostic advice is adaptive, clear, and easy for doctors to interpret.
2. We aim to grow our Few-Shot Learning database to cover a much broader range of rare diseases while finding better ways to pull in data from sources beyond just OCR.
3. We want to speed up our software for faster performance and connect it with wearables or IoT devices to help track a patient's health constantly in real-time.

XI. CONCLUSION

This paper introduces a hybrid multimodal framework for smart medical disease detection by blending image analysis, natural language processing, temporal modeling, and anomaly detection. By merging few-shot learning with anomaly-aware decision fusion, the system successfully boosts rare disease identification and overall diagnostic precision. This approach moves past the hurdles of traditional single-modality systems to offer explainable, dependable, and scalable support for clinicians, making it a practical solution for real-world healthcare settings.

REFERENCES

- [1] Othmani, A., Siarry, P., & Sabri, A. Q. M.(2009). "Rare Event Detection in Multimedia Data: A Survey of Machine Learning Techniques." *International Journal of Machine Learning and Applications*, Vol. 3, No. 2, pp. 45–58.
- [2] Norén, G. N., Meldau, E., & Elleenius, J.(2010). "Evaluation Principles for Rare Event Detection in Pharmacovigilance." *Drug Safety Journal*, Vol. 33, No. 11, pp. 993–1004.
- [3] Abdulrazaq, M. (2014). "Rare Event Prediction in Highly Imbalanced Datasets." *International Journal of Data Mining and Knowledge Discovery*, Vol. 8, No. 4, pp. 245–255.
- [4] Alshemaimri, B., Daud, A., & Khalique, F.(2015). "Anomaly Detection in Smart Healthcare Systems Using Isolation Forest and Local Outlier Factor." *International Journal of Information Security and Applications*, Vol. 21, pp. 75–84.
- [5] Butte, A. J., Howell, M. D., Cui, C., & Corrado, G.S.(2018). "Deep Learning for Electronic Health Records Using the FHIR Data Standard." *Nature Digital Medicine*, Vol. 1, No. 18, pp. 1–10.
- [6] Che, Z., Purushotham, S., Cho, K., & Sontag, D.(2018). "Recurrent Neural Networks for Multivariate Time Series with Missing Values." *Scientific Reports*, Vol. 8, No. 6085, pp. 1–12.
- [7] Luo, Y., Fu, Y., & Wang, F. (2019). "Machine Learning Approaches for Rare Disease Detection in Longitudinal Healthcare Data." *IEEE Journal of Biomedical and Health Informatics*, Vol. 23, No. 4, pp. 1450–1460.
- [8] Yang, C., Wang, T., & Jiang, X. (2020) "Deep Autoencoder-Based Anomaly Detection for Healthcare Data." *IEEE Access*, Vol. 8, pp. 108–117.
- [9] Niu, H., & Omiaomu, O. A. (2020). "EHR-BERT: A Transformer-Based Model for Electronic Health Record Analysis." *IEEE International Conference on Healthcare Informatics*, pp. 120–128.
- [10] Xun, G., Ma, F., Gao, J., & Zhang, A. (2021). "Graph-Based Deep Learning for Rare Disease Detection in Electronic Health Records." *ACM Transactions on Knowledge Discovery from Data*, Vol. 15, No. 3, pp. 1–22.
- [11] Shukla, P. K., Shukla, P., & Pandey, A.(2022). "Multimodal Transformer Framework for

Rare Disease Diagnosis.” IEEE Access, Vol. 10, pp. 45678–45689.

- [12] Choi, E., & Sun, J. (2022) “Deep Learning in Electronic Health Records: A Systematic Review.” Journal of Biomedical Informatics, Vol. 126, pp. 103–115.
- [13] Niu, H., & Omitaomu, O. A. (2023). “Advanced Deep Learning Approaches for Rare Medical Event Detection.” IEEE Transactions on Artificial Intelligence in Healthcare, Vol. 4, No. 2, pp. 95–105