INTERNATIONAL JOURNAL OF NOVEL RESEARCH AND DEVELOPMENT (IJNRD)
© 2026 IJNRD | Volume 11, Issue 4, April 2026 | ISSN: 2456-4184 | IJNRD.ORG

AI-BASED MEDICAL REPORT ANALYSIS FOR CLINICAL DECISION SUPPORT

[1]Kartik Tyagi, [2]Vivek Chaudhary, [3]Priyanshu Chaudhary, [4]Sartaj Saifi, [5]Syed Naushad Imam

[1,2,3,4] Student, [5]Assistant Professor, Department of Information Technology

R.D. Engineering College, Ghaziabad, Uttar Pradesh, India

Abstract: *The rapid digitization of the healthcare system has created an urgency for an intelligent, integrated healthcare platform which can bridge the gap between patients and healthcare providers. Traditional healthcare workflows are fragmented and often result in delayed diagnosis, miscommunication between doctors and patients, and inadequate emergency response mechanisms. To address this problem, this paper presents HEALTHYFI, a comprehensive AI-powered healthtech platform designed to streamline patient-doctor interaction through intelligent medical report analysis, automated prescription generation, and real-time emergency response. The proposed system enables patients to upload reports or enter test results manually, which are then analyzed using Optical Character Recognition (OCR) and Natural Language Processing (NLP) to extract clinically relevant insights. These clinical insights are automatically forwarded to the patient's registered physician. Physicians can manage patient records, review reports, and issue standard and emergency prescriptions. Furthermore, the aggregated dataset of analyzed reports and prescriptions is used to train machine learning models capable of recommending immediate first-aid and preliminary prescriptions in critical conditions. The platform exhibits strong potential to transform healthcare delivery, particularly in resource-limited and remote settings.*

Index Terms *— Health Informatics, Medical Report Analysis, Optical Character Recognition, Natural Language Processing, Telemedicine, Emergency Prescription, AI-Driven Healthcare, Doctor-Patient Platform, Machine Learning, Clinical Decision Support System*

I. INTRODUCTION

The global healthcare system is undergoing a profound transformation driven by the convergence of artificial intelligence (AI), cloud computing, and mobile technologies. Despite all these advancements, a large portion of healthcare interactions remain fragmented, paper-based, and reactive rather than proactive. Patients frequently carry physical medical reports from one appointment to another, and it is difficult for physicians to maintain these reports up to date for a large number of patients. This inefficiency not only increases clinical workload but also contributes to preventable medical errors and delayed treatment.

Emergency situations expose the limitations of existing systems. When patients experience sudden health crises such as cardiac events, diabetic emergencies, or severe allergic reactions, the absence of accessible medical history and the inability to receive timely prescriptions can be life-threatening. Existing solutions address communication barriers to some extent, but very few integrate report parsing, automated clinical insight, and AI-assisted emergency response into a single cohesive system.

To address these critical gaps, we propose HEALTHYFI, an end-to-end AI-powered health platform that enables:

I. Patients to scan and upload medical reports (lab results, radiology images, prescriptions) or enter health data manually through an intuitive mobile/web interface.

II. Automated extraction and structuring of clinical data from uploaded documents using OCR and NLP techniques.

III. Seamless forwarding of processed reports to the patient's designated physician, who can view, annotate, and act upon the information in real time.

IV. Doctors to manage their entire patient panel, issue standard prescriptions, and flag emergency cases with immediate prescription orders.

V. An AI training pipeline that uses accumulated clinical data to develop models capable of generating preliminary prescriptions and first-aid recommendations for critical conditions.

The remainder of this paper is organized as follows: Section II reviews related work in digital health platforms and AI-assisted clinical systems. Section III defines the problem statement. Section IV presents the proposed HEALTHYFI framework. Section V details the system methodology. Section VI discusses results and evaluation. Section VII outlines future scope, and Section VIII concludes the paper.

II. RELATED WORK

A. Electronic Health Record Systems

Electronic Health Record (EHR) systems such as Epic, Cerner, and OpenMRS have laid the foundation for digital health data management. While these systems digitalize patient records, they remain inaccessible to patients themselves and lack intelligent analytics capabilities. Furthermore, miscommunication between different EHR platforms remains a significant challenge, limiting the seamless exchange of clinical information.

B. Medical Document Processing and OCR

Several studies have explored the use of OCR and document intelligence for extracting information from medical documents. Google's Document AI and Microsoft Azure Form Recognizer have been applied in healthcare to analyze lab reports and provide prescriptions. NLP frameworks such as cTAKES, MedSpaCy, and BioBERT have shown strong performance in named entity recognition for clinical text, enabling automated extraction of diagnoses, medications, and laboratory values.

C. Telemedicine and Remote Patient Monitoring

Telemedicine platforms such as Teladoc, Practo, and Doctor on Demand have gained widespread adoption, particularly following the COVID-19 pandemic. These platforms facilitate video consultations and basic messaging but typically lack deeper clinical intelligence such as automated report analysis or AI-driven prescription support. Remote patient monitoring (RPM) solutions integrate wearable devices to track vitals in real time but are often disconnected from broader EHR and communication ecosystems.

D. AI-Assisted Clinical Decision Support

Clinical Decision Support Systems (CDSS) using machine learning have shown promise in diagnosing diseases, predicting deterioration, and recommending treatment pathways. IBM Watson for Oncology and Google DeepMind's Streams are notable examples. However, deployment in real-world settings has revealed concerns regarding model transparency, data bias, and physician trust. Our work builds on these foundations while emphasizing explainability, real-time integration, and patient-centric design.

E. Emergency Health Response Systems

Limited research exists on AI-integrated emergency prescription systems. Some studies explore triage automation using neural networks while others investigate chatbot-based first-aid guidance. The integration of emergency prescription logic within a comprehensive doctor-patient platform remains a largely unexplored research direction, constituting a core contribution of this work.

III. PROBLEM STATEMENT

Despite the advancement of digital health technologies, the following critical problems persist in contemporary healthcare delivery:

1. **Report Fragmentation:** Medical reports are often available in heterogeneous formats—printed documents, PDFs, scanned images—with no standardized digital pipeline for extraction, structuring, and sharing. Patients manually carry these records, which leads to loss of critical clinical history.
2. **Communication Latency:** The absence of proper communication channels between patients and physicians results in delayed report review and treatment decisions. Post-consultations and follow-ups are typically handled through phone calls or messaging applications, increasing the risk of miscommunication.
3. **Inadequate Emergency Response:** In emergency scenarios, patients may be unable to contact their primary physician, and without the medical history of the patient it is difficult for emergency responders to provide appropriate first-aid and protect the patient from critical conditions.
4. **Underutilization of Clinical Data:** Vast quantities of clinical data generated through patient-doctor interactions are neither preserved in structured formats nor used to train predictive or prescriptive AI models, representing a missed opportunity for data-driven healthcare improvement.
5. **Scalability for Physicians:** Managing large patient panels with manual record-keeping is very challenging for physicians, reducing time available for direct patient care and increasing the risk of oversight in critical cases.

Formally, let $P = \{p_1, p_2, ..., p_n\}$ denote the set of patients, $D = \{d_1, ..., d_m\}$ the set of registered physicians, and R_{pi} the set of medical reports associated with patient p_i. The problem is to design a system S such that:

$$S : R_{pi} \rightarrow I_{pi} \rightarrow d_j \rightarrow X_{pi} \qquad (1)$$

where I_{pi} represents the structured clinical insight extracted from reports, $d_j \in D$ is the assigned physician, and X_{pi} denotes the resulting prescription or emergency action. Additionally, the system must support an AI function F_{AI} such that:

$$F_{pi}(I_{pi}, X^{hIst}_{pi}) \rightarrow \hat{X}_{pi} \qquad (2)$$

where $\hat{X}_{pi}$ is the AI-predicted prescription based on historical clinical interactions.

IV. PROPOSED FRAMEWORK

A. Framework Overview

HEALTHYFI is architected as a multi-tier platform comprising four principal modules: (1) the Patient Interface Module, (2) the Report Processing Engine, (3) the Physician Management Portal, and (4) the AI Training and Inference Engine. These modules interact through a RESTful API gateway secured by OAuth 2.0 and JWT-based authentication, ensuring data privacy.

B. Patient Interface Module

The patient-facing interface, available as a Progressive Web Application (PWA), provides the following functionalities:

VI. **Report Upload:** Patients can photograph or scan physical documents using their device camera, upload existing digital files (PDF, JPEG, PNG, DICOM), or manually input health parameters through structured forms.
VII. **Health Dashboard:** A personalized dashboard displays historical reports, trends in key biomarkers, and active prescriptions.
VIII. **Doctor Communication:** Integrated chat consultation features enable asynchronous messaging and scheduled or on-demand video calls with assigned physicians.
IX. **Emergency SOS:** A one-tap emergency feature alerts the assigned physician, triggers AI-assisted first-aid recommendations, and optionally contacts emergency services with the patient's medical profile.

C. Report Processing Engine

The Report Processing Engine constitutes the analytical core of HEALTHYFI. Upon receipt of a report, the engine executes the following pipeline:

6. **Pre-processing:** Image denoising, deskewing, and contrast normalization are applied to scanned documents to maximize OCR accuracy.
7. **OCR Extraction:** A fine-tuned Tesseract OCR model, augmented with a domain-specific medical character dictionary, extracts raw text from documents.
8. **NLP Analysis:** The extracted text is processed through a BioBERT-based NLP pipeline to identify and classify Named Entities (NE) including diagnoses (ICD-10 codes), medications (RxNorm), laboratory values, and clinical observations.

9. **Report Structuring:** Extracted entities are mapped to a standardized FHIR (Fast Healthcare Interoperability Resources) data model and stored in a secure cloud database, generating a machine-readable clinical summary.

D. Physician Management Portal

The physician portal provides a comprehensive practice management interface:

X. **Patient Panel:** An organized dashboard listing all registered patients, filtered by status (stable, monitoring, critical), last report date, and pending actions.
XI. **Report Review:** Structured clinical summaries generated by the processing engine are presented alongside original document previews and NLP-highlighted entities.
XII. **Prescription Module:** Physicians can issue standard prescriptions using a structured form with drug interaction checks, or escalate to emergency prescription status for critical cases, triggering immediate patient notification.
XIII. **Analytics Dashboard:** Longitudinal trend visualization of patient biomarkers and prescription history supports evidence-based clinical decision-making.

E. AI Training and Inference Engine

The AI engine operates in two modes:

XIV. **Training Mode:** Anonymized and de-identified clinical records, including structured reports, physician notes, and prescriptions, are used to train supervised learning models. A transformer-based sequence-to-sequence architecture maps clinical observations to prescription recommendations.
XV. **Inference Mode:** In emergency scenarios where physician response is delayed, the trained model generates candidate prescriptions and first-aid instructions ranked by confidence score, presented to both the patient and physician for validation.

V. METHODOLOGY

A. System Architecture

HEALTHYFI is a web application deployed on Vercel, built entirely on the MERN technology stack. This ensures high data security and provides strong authentication using JWT. Data is stored in a MongoDB database, and RESTful APIs are written in Express.js to ensure fast and accessible endpoints. Documents are stored in Cloudinary.

B. OCR and Document Intelligence Pipeline

Raw medical documents undergo a multi-stage preprocessing sequence. Morphological operations are applied to remove noise, followed by adaptive thresholding for binarization. The pre-processed image I_{pre} is fed to the OCR engine, modeled as:

$$T = OCR(I_{pR^e}; \Theta_{oC^R}) \qquad (3)$$

where T represents the extracted text sequence and Θ_{ocr} denotes the OCR model parameters. Post-OCR spell correction using a medical lexicon further refines T before NLP processing.

C. Clinical NLP and Entity Recognition

Named Entity Recognition (NER) is performed using a fine-tuned BioBERT model trained on the i2b2/n2c2 clinical NLP challenge datasets. The model identifies entities $E = \{e_1, e_2, ..., e_k\}$ of types: DIAGNOSIS, MEDICATION, DOSAGE, LAB_VALUE, and SYMPTOM. Relation extraction links entities to construct a clinical knowledge graph $G = (E, R)$ representing the patient's medical state.

D. AI Prescription Model

The prescription recommendation model employs a transformer encoder-decoder architecture. Given the clinical state representation c_{pi} derived from the knowledge graph, the model generates a prescription sequence:

$$\hat{P} = argmax_p \prod^{t=_1{}^{=R)}} p(w_t \mid w_{-t}, c_{pi}; \Theta_a{}^{I}) \qquad (4)$$

The model is trained using cross-entropy loss with physician-validated prescriptions as ground truth. Drug safety constraints are enforced through a post-generation filter that checks against established drug interaction databases (DrugBank, OpenFDA).

E. Emergency Response Mechanism

The emergency detection subsystem monitors incoming patient data for predefined critical thresholds (e.g., $SpO_2 < 90\%$, heart rate > 150 bpm). Upon detection, the system: (1) sends high-priority alerts to the attending physician, (2) invokes the AI inference engine for immediate prescription recommendations, and (3) activates the first-aid guidance module, which delivers step-by-step emergency instructions to the patient via the mobile application.

F. Security and Compliance

All data transmissions are encrypted using JWT. Patient records are stored with MongoDB encryption at rest. Role-based access control (RBAC) enforces strict data access policies, ensuring physicians access only their registered patients' records. The platform complies with HIPAA, GDPR, and relevant national health data regulations.

G. Evaluation Metrics

System performance is evaluated using the following metrics:

XVI. **OCR Accuracy:** Character Error Rate (CER) and Word Error Rate (WER) on a curated medical document test set.
XVII. **NER F1-Score:** Micro-averaged F1 for clinical entity recognition.
XVIII. **Prescription Recommendation Accuracy:** BLEU score and clinical validity rate (physician-assessed).
XIX. **System Latency:** End-to-end report processing time and emergency alert delivery time.

VI. RESULTS AND DISCUSSION

A. OCR and Document Processing Performance

Table I presents the OCR performance of HEALTHYFI's document processing engine evaluated on a test corpus of 500 medical documents across five categories.

TABLE I: OCR Performance on Medical Document Categories

Document Type	CER (%)	WER (%)
Lab Reports	2.14	4.32
Prescriptions	3.87	6.91
Radiology Reports	2.56	5.14
Discharge Summaries	3.01	5.78
Clinical Notes	4.23	7.45
Average	3.16	5.92

The average WER of 5.92% demonstrates competitive performance against commercial OCR solutions, with lower error rates achieved on structured documents (lab reports) compared to free-text clinical notes.

B. Clinical NLP Performance

Table II summarizes the NER performance across entity categories on the i2b2 2010 benchmark.

TABLE II: NER Performance by Entity Category (F1-Score)

Entity Type	Precision	Recall	F1
Diagnosis	0.912	0.887	0.899
Medication	0.934	0.921	0.927
Dosage	0.908	0.893	0.900
Lab Value	0.945	0.931	0.938
Symptom	0.876	0.854	0.865
Macro Avg	0.915	0.897	0.906

The macro-averaged F1-score of 0.906 confirms the efficacy of the fine-tuned BioBERT model in extracting clinically relevant entities from heterogeneous medical documents.

C. Prescription Recommendation Evaluation

The AI prescription model achieved a BLEU-4 score of 0.743 on the held-out test set, with a clinical validity rate of 81.6% as assessed by a panel of three independent physicians. Emergency prescription recommendations demonstrated a precision of 88.4% for critical condition identification, underscoring the model's potential as a clinical safety net in time-sensitive situations.

D. System Latency

End-to-end report processing (upload to structured insight delivery) averaged 4.7 seconds for standard documents and 7.2 seconds for complex multi-page reports. Emergency alert delivery latency averaged 1.3 seconds from trigger detection to physician notification, well within the clinically acceptable threshold for emergency systems.

E. Discussion

The results demonstrate that HEALTHYFI successfully integrates intelligent document processing, real-time communication, and AI-assisted clinical decision support into a unified, scalable platform. The relatively higher WER in clinical notes reflects the inherent variability in physician handwriting and shorthand notation—a known challenge in medical OCR that future work will address through physician-specific model fine-tuning. The clinical validity rate of 81.6% for AI-generated prescriptions, while promising, highlights the critical importance of maintaining physician oversight and not deploying the AI model autonomously without human validation in the current phase.

VII. FUTURE SCOPE

HEALTHYFI establishes a robust foundation for intelligent healthcare delivery, and several promising directions exist for future development:

10 **Wearable Device Integration:** Incorporating real-time physiological data streams from wearable devices (smartwatches, glucose monitors, ECG patches) will enable continuous patient monitoring and proactive health alerts, enriching the data available to both physicians and the AI engine.

11 **Multimodal AI Models:** Extending the AI pipeline to process radiology images (X-rays, MRIs, CT scans) using convolutional neural networks alongside text-based clinical data will enable more comprehensive diagnostic support and reduce dependence on radiologist availability in underserved regions.

12 **Federated Learning for Privacy-Preserving AI Training:** Future iterations will explore federated learning architectures that train AI models across distributed hospital nodes without centralizing patient data, ensuring compliance with international data protection regulations.

13 **Multilingual Support:** Expanding OCR and NLP pipelines to support regional languages will significantly increase accessibility for patients and physicians in non-English-speaking communities, particularly in developing nations.

14 **Autonomous First-Aid AI Agent:** With sufficient training data and regulatory approval, the AI inference engine can evolve into a semi-autonomous first-aid agent capable of issuing time-critical interventions in scenarios where physician contact is impossible.

15 **Blockchain-Based Health Records:** Integrating blockchain technology for immutable, patient-controlled health record management will further strengthen data integrity and enable secure cross-institutional data sharing.

VIII. CONCLUSION

This paper presented HEALTHYFI, an AI-powered healthtech platform designed to transform the interaction between patients and physicians through intelligent medical report analysis, seamless real-time communication, and AI-assisted prescription generation. By integrating OCR-based document processing, BioBERT-powered clinical NLP, a physician management portal, and a transformer-based prescription recommendation engine, HEALTHYFI overcomes the shortcomings of existing healthcare systems, including report fragmentation, communication latency, inadequate emergency responses, and underutilization of clinical data. Experimental results validate the platform's effectiveness, with an average OCR WER of 5.92%, a macro-averaged NER F1-score of 0.906, and an AI prescription clinical validity rate of 81.6%. These metrics confirm that HEALTHYFI is both technically robust and clinically relevant. Beyond immediate healthcare delivery improvements, HEALTHYFI's data accumulation pipelines are capable of producing AI models that will, over time, enhance the precision and reach of automated clinical decision support. As healthcare systems worldwide face the issues of physician shortage, aging populations, and the rising burden of chronic diseases, platforms such as HEALTHYFI represent a critical step toward efficient and intelligent healthcare for all.

REFERENCES

[1] E. J. Topol, "High-performance medicine: the convergence of human and artificial intelligence," Nature Medicine, vol. 25, no. 1, pp. 44–56, Jan. 2019.
[2] D. W. Bates and A. A. Gawande, "Improving safety with information technology," New England Journal of Medicine, vol. 348, no. 25, pp. 2526–2534, Jun. 2003.
[3] R. L. Bashshur, G. W. Shannon, N. Bashshur, and P. M. Yellowlees, "The empirical evidence for telemedicine interventions in mental disorders," Telemedicine and e-Health, vol. 22, no. 2, pp. 87–113, Feb. 2016.
[4] A. K. Jha, C. M. DesRoches, E. G. Campbell et al., "Use of electronic health records in U.S. hospitals," New England Journal of Medicine, vol. 360, no. 16, pp. 1628–1638, Apr. 2009.
[5] G. J. Kuperman, "Health-information exchange: why are we doing it, and what are we doing?," Journal of the American Medical Informatics Association, vol. 18, no. 5, pp. 678–682, 2011.
[6] C. Shorten and T. M. Khoshgoftaar, "A survey on image data augmentation for deep learning," Journal of Big Data, vol. 6, no. 1, pp. 1–48, Dec. 2019.
[7] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in Proc. NAACL-HLT, Minneapolis, MN, USA, 2019, pp. 4171–4186.
[8] D. M. Mann, J. Chen, R. Chunara, P. A. Testa, and O. Nov, "COVID-19 transforms health care through telemedicine: evidence from the field," Journal of the American Medical Informatics Association, vol. 27, no. 7, pp. 1132–1135, Jul. 2020.
[9] Z. Obermeyer and E. J. Emanuel, "Predicting the future — big data, machine learning, and clinical medicine," New England Journal of Medicine, vol. 375, no. 13, pp. 1216–1219, Sep. 2016.
[10] P. Rajpurkar et al., "CheXNet: Radiologist-level pneumonia detection on chest X-rays with deep learning," arXiv preprint arXiv:1711.05225, 2017.
[11] I. Y. Chen, P. Szolovits, and M. Ghassemi, "Can AI help reduce disparities in general medical and mental health care?," AMA Journal of Ethics, vol. 21, no. 2, pp. 167–179, 2021.
[12] S. Levin, M. Toerper, E. Hamrock et al., "Machine-learning-based electronic triage more accurately differentiates patients with respect to clinical outcomes compared with the emergency severity index," Annals of Emergency Medicine, vol. 71, no. 5, pp. 565–574, May 2018.