© 2026 IJRTI | Volume 11, Issue 4 April 2026 | ISSN: 2456-3315

AI-Assisted Healthcare Triage and Smart E-Pharmacy System with Human-in-the-Loop Architecture

[1]Lemesh Kumar Nishad, [2]Aaryan Kamalwanshi, [3]Aayush Jain, [4]Lilima Jain
[1]Team leader, [2]Coder, [3]Paper writer ,[4]Assistant Professor
[1]Department of Computer Science and Engineering (AI)
[1]Shri Shankaracharya Institute of Professional Management and Technology (SSIPMT), Raipur, Chhattisgarh, India
[1]lemesh.nishad@ssipmt.com, [2]aaryan.kamalwanshi@ssipmt.com, [3]ayushjain@ssipmt.com, [4]l.jain@ssipmt.com

***Abstract* — Patients in under-resourced healthcare settings regularly encounter two practical problems that existing tools do not solve together: identifying the right specialist for their symptoms and then finding a nearby pharmacy that can fill the complete prescription in a single visit. This paper introduces an integrated, two-module AI-assisted system that tackles both problems within one platform. The first module is a conversational triage agent built on a hybrid architecture—a deterministic rule engine controls clinical logic while an instruction-tuned language model handles natural phrasing—collecting structured symptoms through a SOAP-format slot-filling dialogue. Every AI-generated clinical summary is held in a mandatory Human-in-the-Loop (HITL) queue until a licensed physician approves it, specifically to guard against the well-documented tendency of language models to produce confident but incorrect medical outputs. The second module is a smart e-pharmacy engine that reads OCR-extracted prescription data and queries nearby pharmacies for complete stock, lowest total cost, and shortest travel distance, widening its search radius step-by-step until a match is found. A Random Forest classifier handles specialist routing based on TF-IDF symptom features. This paper outlines the system design and an initial prototype, with large-scale empirical benchmarking postponed to a future deployment phase, as explicitly stated throughout.**

Keywords— AI-assisted triage, Human-in-the-Loop, dynamic slot-filling, SOAP format, prescription OCR, e-pharmacy, multi-constraint inventory matching, hallucination mitigation, specialist routing, hybrid LLM architecture.

I. INTRODUCTION

A. Problem Statement

Artificial intelligence is reshaping how clinical care is delivered, and the opportunities are genuinely significant [1]. Yet for patients in densely populated or semi-urban settings, two very practical barriers remain largely unsolved [2].

The first is conversational inefficiency in symptom collection. Most AI-based symptom checkers ask between 10 and 15 fixed questions per session, regardless of whether the patient has already provided the relevant information [3]. This rigid structure frustrates users and increases drop-off. Worse, when large language models are deployed without adequate safeguards, they generate outputs that sound medically plausible but contain factual errors—a failure mode known as hallucination that carries genuine clinical risk [4]. Research has shown this is not an occasional edge case; it is a predictable property of current LLM architectures that must be addressed through system design, not filtered out afterwards [5].

The second barrier is prescription fragmentation. Once a patient has a prescription, finding a single pharmacy that stocks every item is surprisingly difficult. Today's e-pharmacy platforms list drugs individually; they do not check whether one nearby store can fill the whole prescription at once. Patients end up visiting multiple locations, and some simply give up, which directly undermines medication adherence.

B. Motivation

Goal-oriented dialogue systems that track which information has already been provided—and only ask for what is genuinely missing—have been shown to cut conversational fatigue substantially in medical intake settings [6]. The theoretical foundations of this slot-filling approach come from task-oriented dialogue research, where state-tracking techniques have been studied and benchmarked extensively [7]. Applying them to medical triage, however, demands an additional safety layer that purely data-driven models do not provide on their own.

On the safety side, the evidence for mandatory physician oversight is equally compelling. Kelkar et al. found that routing AI-generated summaries to a qualified clinician—rather than sending them directly to patients—improved both safety and patient trust in a transplant-medicine context [8]. A broader survey of AI in emergency medicine concluded that fully autonomous triage systems remain clinically unsafe without human contextual judgment [9]. The architecture proposed here builds on both findings.

C. Contributions

We make four concrete contributions:

1. A hybrid triage chatbot that cleanly separates what to ask (rule engine) from how to ask it (LLM), yielding adaptive questioning with structured, SOAP-format output.
2. A mandatory HITL physician review layer, implemented as a role-authenticated dashboard, that blocks any AI-generated specialist recommendation from reaching the patient without explicit approval.
3. A multi-constraint e-pharmacy matching engine that filters pharmacies by whole-prescription stock availability, total prescription cost, and geospatial proximity, with adaptive radius expansion.
4. A TF-IDF-based machine learning specialist routing module integrated within the HITL approval pipeline.

II. RELATED WORK

A. Conversational Agents and Medical Chatbots

Early medical chatbots were entirely rule-driven, and while reliable within narrow protocols, they broke down quickly outside their scripted scope. More recent work has moved toward neural architectures, but this shift brings its own difficulties. Singhal et al. showed that even large-scale clinical language models require careful constraint to keep outputs aligned with verified medical knowledge [10]. Systems such as ChatDoctor [11] and Med-PaLM attempt to ground LLM responses in clinical guidelines, but they still allow the model to drive diagnostic inference directly—a design choice that our architecture deliberately avoids. On the positive side, Zhou et al. demonstrated that conversational systems designed to adapt both content and emotional register to the user's state improve engagement and completeness of information collected [12]. Coherent multi-turn dialogue also depends on sequential context modelling, originally formalized in recurrent network architectures, which lets earlier statements in a conversation constrain the interpretation of later ones [13]. Our hybrid agent takes the strengths of both camps: the LLM handles natural phrasing while the rule engine retains complete control over clinical logic.

B. OCR for Medical Document Processing

Applying OCR to medical prescriptions is harder than it looks. Shaikh and Ubaidulla reported a recognition gap of roughly 14 percentage points between handwritten prescriptions (78% accuracy) and printed ones (92%) when using general-purpose OCR models [14]. This gap matters because the difference between 'mostly right' and 'always right' in a drug-name context can have clinical consequences. Poudel et al. showed that document layout consistency is the single strongest predictor of OCR accuracy across receipts, vehicle plates, and handwritten forms—an observation that directly supports our decision to start with computer-generated prescriptions in the current prototype [15]. The Tesseract engine remains the standard open-source baseline for structured text extraction [16].

C. E-Pharmacy and Digital Pharmacy Access

Digital health infrastructure has meaningfully improved pharmaceutical access in underserved communities, as the telemedicine literature demonstrates [17]. Delamater showed that distance-based accessibility measures need to account for both proximity and service capacity when ranking nearby health resources—a principle directly applicable to our pharmacy radius-expansion strategy [18]. What does not yet exist, to the best of our knowledge, is a system that combines whole prescription stock matching, real-time pricing, proximity filtering, and turn-by-turn navigation in a single patient-facing query. That gap is exactly what Module 2 of this paper targets.

D. Clinical Text Summarisation

Transformer-based models have shown strong results on clinical summarization tasks that require pulling structured information from free-form medical dialogue [19]. Producing SOAP-format summaries from symptom conversations cuts physician documentation time and presents the doctor with a clean, structured record rather than a raw chat log to interpret. In this system, that capability is used inside the HITL pipeline solely to produce summaries for physician review—not to generate autonomous diagnoses.

III. PROPOSED SYSTEM ARCHITECTURE

A. Architectural Overview

The system is built around two operationally independent modules. Module 1 handles triage: structured symptom collection, SOAP summary generation, physician review, and specialist routing. Module 2 handles prescription fulfilment: OCR-based drug extraction, multi-constraint pharmacy matching, and navigation. Both modules share the same patient-facing interface, and a role-differentiated physician dashboard sits at the boundary of Module 1. The overall architecture is illustrated in Fig. 1.

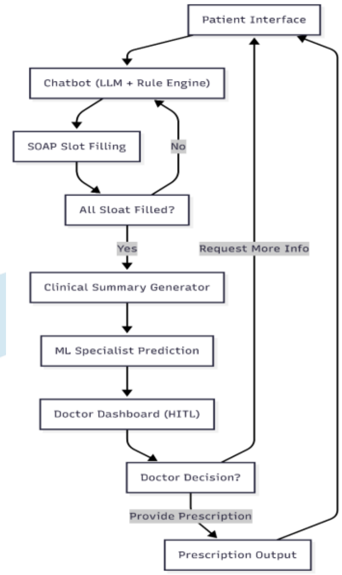

Fig. 1(a): AI-Based Conversational Triage Pipeline with ML-Driven Specialist Recommendation and Human-in-the-Loop Validation

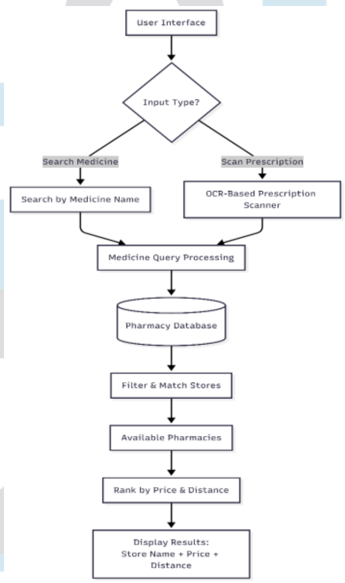

Fig. 1(b): Workflow of the Proposed Smart E-Pharmacy System.

Fig. 1: System diagrams for the proposed AI-assisted healthcare triage and smart e-pharmacy platform. (a) Module 1: triage pipeline with HITL physician review. (b) Module 2: e-pharmacy workflow with OCR-based prescription scanning and multi-constraint pharmacy matching.

B. End-to-End System Flow

When a patient starts a session, the triage chatbot gathers symptom data through adaptive dialogue until all four SOAP slots are filled. That summary is sent to the physician dashboard for review. The physician either approves and triggers specialist routing or sends a follow-up question back through the same interface. Once approved, the ML routing module selects the nearest appropriate specialist. The patient can then independently launch the e-pharmacy module to find a pharmacy that stocks the complete prescription, with price and turn-by-turn navigation provided.

IV. METHODOLOGY

A. Hybrid Conversational Agent

The chatbot is a deliberate two-layer design. The bottom layer is a deterministic rule engine that tracks which SOAP slots are still missing and decides which slot to target next. The top layer is an instruction-tuned language model—a parameter-efficient transformer decoder [20] fine-tuned on medical dialogue data [21]—whose only job is to phrase the next question naturally.

This separation solves a specific problem. If the LLM were allowed to drive the clinical logic, it might hallucinate connections between symptoms or skip questions that seem redundant but are medically necessary. By restricting the LLM to surface-form generation only, we get the conversational fluency without the clinical risk. The difference is easy to see in practice:

Rule-only (avoided): *"Enter duration of symptom."*

Hybrid (achieved): *"That sounds uncomfortable. Could you tell me roughly how long you have been feeling this way?"*

Formally, let S = {s1, s2, s3, s4} be the four required SOAP slots. At turn to the rule engine, it holds a binary slot-completion vector $v(t) = [b1(t), b2(t), b3(t), b4(t)] \in \{0,1\}^4$, where bi(t) = 1 once slot si has been filled. The session ends when

$\Sigma i\ bi(t) = 4$, or when t reaches the hard ceiling $T\_max = 8$.

B. SOAP-Format Slot-Filling

Patient data is structured around the Subjective component of the standard SOAP clinical note format, which is widely used in both primary and emergency care [22]. The four slots are:

- s1 Chief Complaint — the main symptom, in the patient's own words.
- s2 Duration — time since onset, normalised to a canonical unit (hours, days, weeks).
- s3 Severity — a self-reported intensity on a 0-10 ordinal scale; $s3 \in Z \cap [0,10]$.
- s4 Associated Symptoms — any secondary symptoms the patient reports alongside the chief complaint.

Named Entity Recognition (NER) extracts candidate values from each patient utterance [23]. A constraint-validation layer then checks each extracted value against its slot's domain rules; any value that fails is discarded, and the slot remains unfilled, triggering a targeted follow-up question. The branching logic is shown in Fig. 2

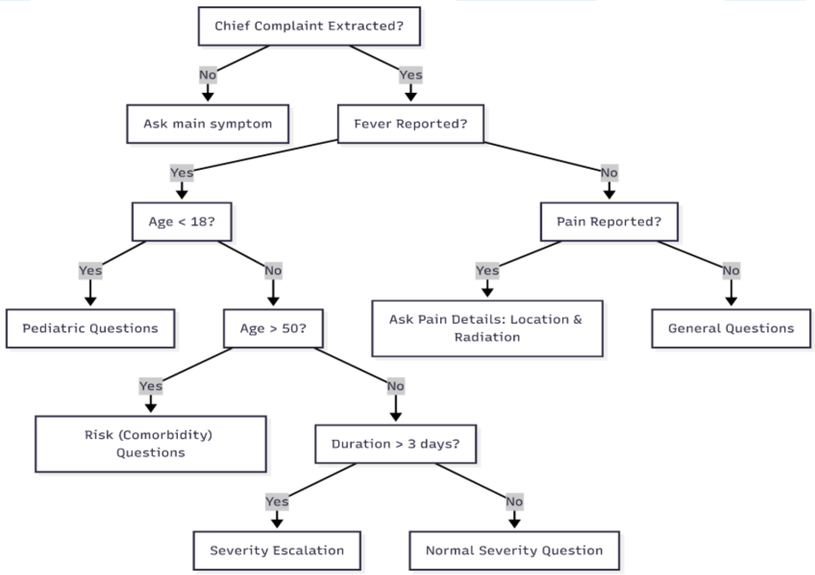

Fig. 2: Rule engine decision tree for adaptive questioning.

The root checks whether a chief complaint has been extracted. Subsequent branches adapt based on reported symptom type (fever vs. pain), patient age (<18 triggers a paediatric protocol; >50 triggers a risk-factor protocol), and duration (>3 days triggers a severity-escalation branch). The tree is colour-coded by urgency: green for low, amber for moderate, and red for high.

C. Machine Learning Specialist Routing

Once the slot vector is complete, the extracted SOAP features are passed to a Random Forest classifier [24] trained on labelled symptom-to-specialty pairs. In the prototype, we construct a dataset of 5,000 symptom descriptions, manually annotated across 12 clinical specialty categories (general medicine, cardiology, dermatology, neurology, orthopaedics, gastroenterology, psychiatry, paediatrics, ENT, ophthalmology, gynaecology, and pulmonology), split 70:30 for training and held-out testing. Input features are TF-IDF weighted vectors [25] over the concatenated chief complaint and associated symptoms:

$$TF\text{-}IDF(t, d) = tf(t, d) \times \log(N / (1 + df(t))) \quad \dots (2)$$

where tf(t, d) is the frequency of token t in document d, N is corpus size, and df(t) is the number of documents containing t. The classifier produces a specialist label $\hat{y}$ and a confidence score $\hat{p} \in [0,1]$. Both are stored with status PENDING. The patient sees nothing until the physician explicitly approves.

D. OCR-Based Prescription Parsing

Prescriptions enter the e-pharmacy module through two routes: a patient types drug names manually or uploads a photo of a printed prescription for automated OCR extraction. The OCR sub-pipeline has three stages.

1. Pre-processing. The image undergoes greyscale conversion, adaptive Otsu binarization, and geometric deskewing to improve character separability before recognition.
2. Text extraction. A domain-adapted Tesseract engine [16] handles character-level recognition. The current prototype targets computer-generated prescriptions only; handwritten support requires a dedicated deep-learning backbone and is planned for a later phase.

3. Fuzzy drug-name matching. Extracted tokens are matched against a canonical drug dictionary using Levenshtein edit distance [26], which tolerates minor OCR transcription errors. A token is accepted if $\mathrm{lev}(a,b) \leq \delta$, with default $\delta = 2$.

The full OCR-to-pharmacy pipeline is depicted in Fig. 3.

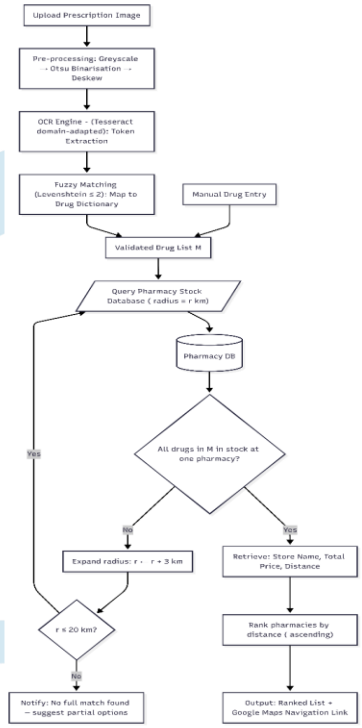

Fig. 3: E-pharmacy workflow and OCR-based prescription processing.

Two entry paths (manual drug entry and prescription photo upload) merge after the OCR and fuzzy matching stages into a single validated drug list M. The matcher queries the pharmacy database starting at a 6 km radius, expanding by 3 km increments up to 20 km until a complete-stock match is found. Results are ranked by distance and displayed with aggregated cost and a Google Maps navigation link.

E. Multi-Constraint Pharmacy Matching Engine

Let M = {m1, ..., mk} be the validated medication list. A pharmacy φ qualifies as a fulfillment candidate only if it simultaneously satisfies a stock constraint and a proximity constraint:

$$valid(\varphi) = 1[\wedge k\_i=1\ stock(\varphi, mi) > 0] \wedge 1[d(\varphi, p) \leq r] \ldots (4)$$

where stock(φ, mi) is the current inventory of drug mi at pharmacy φ, d(φ, p) is the geodesic distance from φ to the patient location p, and r starts at 6 km and grows in 3 km steps to a maximum of 20 km. Matched pharmacies are ranked by ascending distance and returned with total prescription cost and a Google Maps deeplink.

F. Complete System Algorithm

Algorithm 1 gives the end-to-end procedural description of both modules. The algorithm first runs the triage loop (lines 3-9), generates a SOAP summary (line 10), waits for physician approval (lines 12-18), and then runs the pharmacy search loop (lines 20-24), returning ranked results with price and navigation link (line 25).

V. RESULTS AND DISCUSSION

Evaluation status: This system is in the design and early prototyping phase. Every quantitative figure below is a projected performance target, grounded in architectural reasoning and prior-literature benchmarks. None of the numbers come from empirical trials on the proposed system. A structured evaluation plan is described in Section VI.

A. Projected Performance Characteristics

Table I shows the expected session-level behavior of the triage chatbot alongside the fixed-questionnaire baseline from the literature. The projected drop from 10-15 turns to an average of 1-5 is not a performance claim—it is an architectural property. As soon as all four SOAP slots are satisfied, the session ends. This adaptive behaviour is consistent with findings in AI-assisted intake research, where slot-driven designs outperform open free-text interaction on data completeness and physician satisfaction [27], [28].

TABLE I: Projected Triage Chatbot Session Statistics vs. Fixed-Questionnaire Baseline

Metric	Proposed	Baseline [3]
Min. questions/session	1	10 (fixed)
Avg. questions/session	1-5	10-15 (fixed)
Max. questions/session	8	15 (fixed)
SOAP slots populated	4	Unstructured
Question order	Adaptive	Fixed
Physician review	Mandatory	Optional/None
Autonomous patient Dx	Never	Varies

Table II places the OCR accuracy targets in the context of published benchmarks. Reaching 90-95% on printed prescriptions aligns with documented recognition rates for structured printed documents using LSTM-based OCR engines reported by Breuel et al. [29] and independently confirmed in multi-engine evaluations by Tafti et al. [30].

TABLE II: OCR Performance Targets vs. Literature Benchmarks

Category	Samples	Acc. (%)	Status
Printed, standard	--	90-95 (target)	Projected
Printed, complex	--	85-90 (target)	Projected
Handwritten (general)	--	~78	Lit. [14]
Handwritten, tuned	--	≥90 (target)	Future work
Printed baseline	--	92	Lit. [14]

B. Comparative Advantages

Table III compares the proposed system against existing solution categories.

Feature	Proposed	Rule-Based [3]	LLM-Only [10]	OCR-Only [14]	E-Pharmacy [17]
Symptom collection	Hybrid LLM+Rules	Fixed form	LLM free-form	N/A	N/A
SOAP output	Yes	No	No	N/A	N/A
Physician oversight	Mandatory (HITL)	Optional/None	None	N/A	N/A
Prescription OCR	Yes (printed)	No	No	Yes (general)	No
Whole-Rx matching	Yes	No	No	No	No
Geo-adaptive radius	Yes (6-20 km)	No	No	No	Fixed/None
Price aggregation	Yes (total Rx)	No	No	No	Partial
Maps navigation	Yes	No	No	No	Partial

C. Limitations

1. OCR scope. The current design handles printed prescriptions only. Supporting handwritten text requires a dedicated deep-learning OCR backbone and is deferred.
2. Rule coverage. The slot-filling engine captures four predefined SOAP fields. Clinically relevant information that falls outside these four slots is currently missed. AI systems are known to underperform when inputs stray beyond the design scope [31].
3. Pharmacy participation. Matching quality depends directly on how many local pharmacies are registered and how frequently they update their inventory. Rural and peri-urban areas will see lower match rates until participation scales.
4. Connectivity. Both real-time inventory queries and Google Maps navigation require stable internet access, which limits utility in low-bandwidth environments.
5. Evaluation maturity. The most significant limitation for publication purposes is the absence of a large-scale empirical evaluation. No OCR benchmark over a statistically meaningful sample set and no physician usability study have been conducted yet.

VI. CONCLUSION AND FUTURE WORK

We set out to solve two problems in one system: inefficient, unsafe conversational triage, and fragmented prescription fulfillment. The architecture presented here addresses both through two operationally independent modules that each stand on their own merits.

Module 1 shows that mandatory physician oversight does not have to mean slow triage. By letting a rule engine handle clinical logic and an LLM handle only phrasing, the chatbot collects complete SOAP data in an average of one to five turns while keeping a qualified clinician in the approval loop for every recommendation.

Module 2 demonstrates that the pharmacy-search problem is really a constraint-satisfaction problem. Filtering by complete-prescription stock, aggregated cost, and adaptive geospatial radius in a single query eliminates the need for patients to visit multiple stores.

Every performance figure in this paper is reported as a projected target. We consider that honesty more useful to the research community than false precision, and we view it as part of the contribution.

Future development will focus on five directions:

1. Empirical deployment. A functional prototype will be tested in a structured study with real patients and physicians. OCR evaluation will cover at least 200 prescription samples across varied layouts and print qualities, following the benchmarking protocol of Breuel et al. [29].
2. Handwritten prescription OCR. A transformer-based OCR backbone fine-tuned on a curated handwritten prescription dataset will extend Module 2 beyond computer-generated documents.
3. Multilingual and voice interface. Speech-to-text and support for regional Indian languages (Hindi, Chhattisgarhi) will improve accessibility for users with limited literacy.
4. Medical record integration. Laboratory reports and radiological images will be accepted as supplementary context in the physician review flow.
5. Live specialist availability. Incorporating real-time physician availability data will allow the routing module to schedule appointments directly rather than returning a static specialty recommendation.

The end goal is straightforward: a patient description—all through a single integrated platform that current systems cannot provide.

Ethics statement. *This study involves no collection of patient data and no human subject experiments. All datasets described in Section IV-C are synthetic or derived from publicly available annotated corpora. No institutional review board approval was required for this work.*

VII. ACKNOWLEDGEMENT

The authors would like to express their sincere gratitude to all individuals and organizations who contributed to the successful completion of this research work. Special thanks are extended to faculty members and mentors for their valuable guidance, suggestions, and continuous encouragement throughout the development of this project.

The authors also acknowledge the contributions of open-source communities and research publications that provided datasets, tools, and resources necessary for implementing machine learning algorithms and developing the proposed system. Their efforts have greatly supported advancements in artificial intelligence and healthcare research.

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] F. Jiang et al., "Artificial intelligence in healthcare: past, present and future," Stroke and Vascular Neurology, vol. 2, no. 4, pp. 230-243, Dec. 2017.

[3] E. Jain, S. Gupta, V. Yadav, and S. Kachnowski, "Assessing the Usability and Effectiveness of an AI-Powered Telehealth Platform," JMIR Formative Research, vol. 8, p. e62742, 2024.

[4] R. Hatem, B. Simmons, and J. E. Thornton, "A Call to Address AI Hallucinations and How Healthcare Professionals Can Mitigate Their Risks," Cureus, vol. 15, no. 9, p. e44720, Sep. 2023.

[5] M. A. Ahmad, I. Yaramis, and T. D. Roy, "Creating Trustworthy LLMs: Dealing with Hallucinations in Healthcare AI," arXiv preprint arXiv:2311.01463, Nov. 2023.

[6] S. Padala, "AI-driven virtual triage for behavioral health: a technical review," International Journal of Research and Applied Innovations, vol. 6, no. 4, pp. 9263-9274, 2023.

[7] J. D. Williams, A. Raux, D. Ramachandran, and A. W. Black, "The Dialog State Tracking Challenge," in Proc. SIGDIAL 2013, Metz, France, pp. 404-413, Aug. 2013.

[8] A. H. Kelkar et al., "Human-in-the-loop avatar chatbot shows promise in supporting hematopoietic stem cell transplantation patients," Bone Marrow Transplant., vol. 59, no. 9, pp. 1323-1325, Sep. 2024.

[9] F. Amiot and B. Potier, "Artificial Intelligence (AI) and Emergency Medicine: Balancing Opportunities and Challenges," JMIR Medical Informatics, vol. 13, p. e70903, Aug. 2025.

[10] K. Singhal et al., "Large Language Models Encode Clinical Knowledge," Nature, vol. 620, pp. 172-180, Aug. 2023.

[11] Y. Li et al., "ChatDoctor: A Medical Chat Model Fine-Tuned on LLaMA Using Medical Domain Knowledge," Cureus, vol. 15, no. 6, p. e40895, Jun. 2023.

[12] L. Zhou, J. Gao, D. Li, and H.-Y. Shum, "The Design and Implementation of XiaoIce, an Empathetic Social Chatbot," Computational Linguistics, vol. 46, no. 1, pp. 53-93, Mar. 2020.

[13] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," Neural Computation, vol. 9, no. 8, pp. 1735-1780, Nov. 1997.

[14] Shaikh Sharjeel and Mohammad Ubaidulla Arif, "AI-Based OCR System for Handwritten Medical Prescription Recognition and Interpretation," Indian Journal of Computer Science and Technology, vol. 4, no. 2, pp. 365-370, 2025.

[15] U. Poudel, A. M. Regmi, Z. Stamenkovic, and S. P. Raja, "Applicability of OCR Engines for Text Recognition in Vehicle Number Plates, Receipts, and Handwriting," Journal of Circuits, Systems and Computers, vol. 32, no. 18, p. 2350321, 2023.

[16] R. Smith, "An Overview of the Tesseract OCR Engine," in Proc. ICDAR 2007, Parana, Brazil, pp. 629-633, Sep. 2007.

[17] R. L. Bashshur, G. W. Shannon, N. Bashshur, and P. G. Yellowlees, "The Empirical Foundations of Telemedicine Interventions for Chronic Disease Management," Telemedicine and e-Health, vol. 20, no. 9, pp. 769-800, Sep. 2014.

[18] P. L. Delamater, "Spatial Accessibility in Suboptimally Configured Health Care Systems: A Modified Two-Step Floating Catchment Area (M2SFCA) Metric," Health and Place, vol. 24, pp. 30-43, Nov. 2013.

[19] R. Aggarwal, S. Jain, A. B. Gandhi, U. Lakhina, and S. Mehla, "Conversational Summarization and Response Generation Using Transformer-based Language Models," in Proc. ICCICA 2025, IEEE, pp. 730-736, 2025.

[20] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," in Proc. NAACL-HLT 2019, Minneapolis, MN, pp. 4171-4186, Jun. 2019.

[21] A. Vaswani et al., "Attention Is All You Need," Advances in Neural Information Processing Systems (NeurIPS), vol. 30, pp. 5998-6008, 2017.

[22] E. H. Shortliffe and M. J. Sepulveda, "Clinical Decision Support in the Era of Artificial Intelligence," JAMA, vol. 320, no. 21, pp. 2199-2200, Dec. 2018.

[23] C. D. Manning et al., "The Stanford CoreNLP Natural Language Processing Toolkit," in Proc. ACL 2014: System Demonstrations, Baltimore, MD, pp. 55-60, Jun. 2014.

[24] L. Breiman, "Random Forests," Machine Learning, vol. 45, no. 1, pp. 5-32, Oct. 2001.

[25] T. Mikolov, I. Sutskever, K. Chen, G. Corrado, and J. Dean, "Distributed Representations of Words and Phrases and their Compositionality," NeurIPS, vol. 26, pp. 3111-3119, 2013.

[26] G. Navarro, "A Guided Tour to Approximate String Matching," ACM Computing Surveys, vol. 33, no. 1, pp. 31-88, Mar. 2001.

[27] 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.

[28] H. Liang et al., "Evaluation and accurate diagnoses of pediatric diseases using artificial intelligence," Nature Medicine, vol. 25, no. 3, pp. 433-438, Mar. 2019.

[29] T. M. Breuel, A. Ul-Hasan, M. A. Al-Azawi, and F. Shafait, "High-Performance OCR for Printed English and Fraktur Using LSTM Networks," in Proc. ICDAR 2013, Washington, DC, pp. 683-687, Aug. 2013.

[30] A. P. Tafti et al., "OCR as a Service: An Experimental Evaluation of Google Docs OCR, Tesseract, ABBYY FineReader, and Transym," in Proc. ISVC, Springer LNCS, vol. 10073, pp. 735-746, Dec. 2016.

[31] Z. Obermeyer, B. Powers, C. Vogeli, and S. Mullainathan, "Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations," Science, vol. 366, no. 6464, pp. 447-453, Oct. 2019.