



IJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** IV **Month of publication:** April 2026

DOI: <https://doi.org/10.22214/ijraset.2026.79435>

www.ijraset.com

Call:  08813907089

E-mail ID: ijraset@gmail.com

MediAgent AI: An Agentic AI-Based Healthcare Information Chatbot for Public Health Awareness

Aman Singh¹, S. Bharat², Shahbaz Alam³, Shubham Chauhan⁴, Dr. Arun Khatri⁵

Department of Business Analytics, Mittal School of Business, Lovely Professional University, Phagwara, Punjab, India

Abstract: *This paper presents MediAgent AI, a prototype healthcare information chatbot built on an agentic artificial intelligence architecture. Developed on the Botpress conversational AI platform, the system applies natural language processing to interpret user queries spanning symptoms, disease prevention, dietary guidance, and general health topics, drawing exclusively from World Health Organization knowledge repositories to ensure evidence-based responses.*

Unlike conventional rule-based healthcare chatbots, MediAgent AI employs an agentic design in which the system autonomously navigates query understanding, intent detection, knowledge retrieval, and response synthesis without reliance on fixed decision trees. The chatbot is explicitly scoped as an informational tool: it does not diagnose diseases, prescribe medications, or substitute for professional medical consultation, ensuring full compliance with healthcare AI ethics standards.

Empirical evaluation across 200 standardised queries yields an overall weighted accuracy of 87.4%, a mean response latency of 1.8 seconds, and a user satisfaction rating of 4.3 out of 5.0 (n = 20). The paper describes the system architecture, sub-agent design, knowledge integration methodology, and comparative performance analysis, and concludes with a discussion of limitations and prospective research directions. The findings suggest that MediAgent AI represents a cost-effective and scalable complement to public health education initiatives.

Keywords: *Agentic AI, Healthcare Chatbot, Natural Language Processing, WHO Knowledge Base, MediAgent AI, Conversational AI, Public Health Informatics, Health Information Retrieval, Botpress, AI Ethics in Healthcare*

I. INTRODUCTION

The proliferation of digital communication technologies has fundamentally altered the manner in which individuals seek health-related information. According to the World Health Organization (2023), an estimated 4.66 billion people globally use the internet, a significant proportion of whom regularly consult online sources for medical guidance. However, the internet remains an unregulated repository of health content, often characterised by inaccuracies, medical misinformation, and terminological complexity that renders it inaccessible to lay users. This informational asymmetry constitutes a critical public health challenge, particularly in low- and middle-income regions where access to qualified healthcare professionals remains constrained.

Artificial intelligence, and specifically conversational AI in the form of chatbots, has emerged as a technologically viable solution to bridge this healthcare knowledge gap. Early healthcare chatbots were predominantly rule-based systems limited to decision-tree interactions; however, advances in large language models, neural NLP architectures, and agentic AI design have substantially expanded the functional scope of conversational health agents. Modern AI-driven healthcare chatbots are now capable of interpreting complex, multi-intent queries, retrieving information from structured and unstructured knowledge sources, and generating coherent, contextually appropriate responses at scale.

This paper introduces MediAgent AI, an agentic AI-powered healthcare information chatbot that leverages the Botpress platform and WHO knowledge repositories to deliver evidence-based general health information to users. The system is distinguished by its agentic architecture, in which the AI autonomously navigates multi-step reasoning pipelines — query understanding, intent detection, knowledge retrieval, and response synthesis — without requiring rigid pre-programmed pathways. The chatbot is designed exclusively as an informational tool; it does not perform clinical diagnosis, generate prescriptions, or substitute for professional medical consultation.

The primary contributions of this study are: (i) the design and implementation of a functional agentic AI healthcare chatbot; (ii) an empirical performance evaluation incorporating accuracy, response latency, and user satisfaction metrics; (iii) a comparative analysis of MediAgent AI against existing healthcare chatbot systems; and (iv) a critical discussion of ethical considerations, system limitations, and prospective enhancements. The remainder of this paper is organised as follows: Section 2 reviews pertinent literature; Section 3 details the methodology; Section 4 describes the system architecture; Section 5 presents results and discussion; and Section 6 concludes the paper.

II. LITERATURE REVIEW

The academic investigation of AI-powered healthcare chatbots has grown considerably over the past decade, reflecting broader trends in digital health transformation. This section critically reviews significant contributions to the field and situates MediAgent AI within the existing research landscape.

A. Early Rule-Based Healthcare Chatbots

Kavitha and Murthy (2019) presented one of the foundational architectures for AI-assisted healthcare chatbots, demonstrating that even rudimentary pattern-matching systems could reduce the cognitive burden on patients seeking symptom-related information. However, their system was constrained by a fixed decision tree that failed to accommodate query variability beyond a narrowly predefined vocabulary, highlighting the fundamental scalability limitations of rule-based approaches. This limitation has since motivated researchers to explore machine learning-based architectures capable of generalising across a broader and more linguistically varied query space.

B. NLP-Driven Conversational Agents

Oniani and Wang (2020) conducted a qualitative evaluation of large language models applied to automated question-answering in the context of COVID-19 informatics, demonstrating that transformer-based NLP models substantially outperformed keyword-matching systems in both semantic comprehension and response coherence. Their findings established the feasibility of deploying LLM-backed agents in time-sensitive, high-stakes public health contexts. Nevertheless, the authors noted persistent concerns regarding hallucination and factual inconsistency in model outputs — issues that remain pertinent to any health AI deployment.

C. Systematic Reviews of Healthcare Chatbot Applications

Sawad et al. (2022) conducted a comprehensive systematic review of AI conversational agents developed for chronic disease management, covering 47 studies across diabetes, hypertension, and mental health domains. The review found that chatbot-based interventions demonstrably improved patient engagement and self-management behaviours, though the authors cautioned that heterogeneity in evaluation methodologies precluded robust cross-study comparisons. Critically, the review identified the absence of standardised safety frameworks as a major gap, underscoring the ethical imperative for chatbots to operate within clearly demarcated informational — rather than diagnostic — boundaries.

D. Behavioural Change and Health Promotion

Aggarwal et al. (2023) published a systematic review examining AI chatbots as instruments of health behaviour change, analysing 27 randomised and quasi-experimental studies. Their findings indicated that chatbot-delivered health interventions were associated with statistically significant improvements in physical activity, dietary adherence, and medication compliance. The authors attributed these outcomes to the personalised, non-judgmental, and continuously available nature of chatbot interaction — characteristics that distinguish AI agents from conventional health education modalities. However, the review noted that the majority of included studies were conducted in high-income settings, raising concerns about generalisability to resource-limited environments.

E. Women's Health and Demographic-Specific Applications

Kim (2024) undertook a systematic review and meta-analysis of AI chatbot applications in women's health, encompassing 18 studies on maternal health, reproductive health, and oncology. The meta-analysis reported a pooled positive effect size (Cohen's $d = 0.42$) for chatbot interventions on health knowledge outcomes. Kim's work underscores the potential of condition-specific chatbot design to address underserved health demographics, an insight that informs the future enhancement roadmap of MediAgent AI.

F. Agentic AI and Large Language Model Integration

Wen et al. (2024) explored the deployment of large language models in patient engagement applications, arguing that the transition from reactive to agentic AI architectures — wherein the system proactively manages multi-step task pipelines — represents a paradigm shift in digital health tool design. Their framework, which underpins the design philosophy of MediAgent AI, posits that agentic systems are better positioned to handle the inherent ambiguity of natural health-related language, particularly queries involving multiple co-occurring symptoms or multi-domain health concerns.

G. Diagnostic Accuracy and Benchmarking

Bhatt, Ayyagari, and Mishra (2024) proposed a scalable benchmarking methodology for evaluating the differential diagnostic accuracy of health AI systems across conversational sessions. While their evaluation framework was designed for systems with diagnostic intent — explicitly outside the scope of MediAgent AI — their metrics for response coherence, intent classification precision, and knowledge retrieval recall are directly applicable to informational chatbot evaluation and inform the performance measurement approach adopted in this study.

H. General-Purpose AI Avatars in Healthcare

Yan and Alterovitz (2024) investigated the design of general-purpose AI avatars for clinical and public health applications, demonstrating that multi-functional conversational agents consistently outperformed single-function bots on composite user satisfaction metrics. Their work reinforces the design decision in MediAgent AI to integrate multiple informational agents — symptom information, medicine information, diet advice, and health tips — within a unified agentic architecture, thereby improving holistic utility without compromising safety constraints.

I. Comparative Summary

Table summarises the key features of representative healthcare chatbot systems identified in the literature, positioned relative to MediAgent AI across five critical dimensions.

System / Study	Architecture	Knowledge Source	Diagnostic Capability	Agentic Design	Ethical Disclaimers
Kavitha & Murthy (2019)	Rule-based	Static database	Limited	No	Partial
Oniani & Wang (2020)	LLM / Transformer	COVID literature	No	Partial	No
Sawad et al. (2022)	ML / Hybrid	Clinical guidelines	Partial	No	Yes
Aggarwal et al. (2023)	NLP / Behavioural	Multiple	No	Partial	Yes
Wen et al. (2024)	LLM / Agentic	EHR / Literature	No	Yes	Partial
Yan & Alterovitz (2024)	Multi-functional AI	Clinical datasets	Partial	Yes	Partial
MediAgent AI (This study)	Agentic AI / NLP	WHO + Botpress	No	Yes	Comprehensive

Comparative analysis of representative healthcare chatbot systems across architecture, knowledge source, diagnostic scope, agentic design, and ethical disclaimer provisions

III. METHODOLOGY

The development of MediAgent AI followed a structured, iterative design-and-evaluation methodology grounded in established principles of conversational AI engineering, healthcare informatics, and ethical AI governance. The methodology encompasses five interconnected phases: requirements analysis, system design, knowledge integration, testing and validation, and deployment.

A. Requirements Analysis

An initial functional requirements analysis was conducted to delineate the informational scope and operational boundaries of the chatbot. Core functional requirements were categorised into three areas. First, conversational capability: the system must parse and respond to natural language health queries across multiple domains including symptomology, disease prevention, pharmacological information, and nutritional guidance.

Second, knowledge fidelity: all information presented to users must be traceable to peer-reviewed or institutionally authoritative sources. Third, safety compliance: the system must unconditionally refrain from generating clinical diagnoses, pharmaceutical prescriptions, or definitive treatment pathways. Non-functional requirements included response latency under two seconds for 95% of standard queries, a conversational interface compatible with both desktop and mobile form factors, and graceful degradation in the event of knowledge retrieval failures.

B. Platform Architecture and Technology Selection

The Botpress conversational AI platform was selected as the development substrate on the basis of four critical criteria: (i) its native support for autonomous AI nodes capable of multi-step reasoning without pre-programmed dialogue trees; (ii) its integrated knowledge base management functionality enabling vectorised retrieval from external document repositories; (iii) its extensible JavaScript-based customisation environment; and (iv) its cloud-native deployment infrastructure supporting real-time user interactions at scale. The WHO knowledge repositories — specifically the Health Topics, Fact Sheets, and Global Health Observatory databases — were selected as the exclusive information sources to ensure maximal content reliability and institutional credibility.

C. Agentic AI Design Paradigm

The conceptual architecture of MediAgent AI is grounded in the agentic AI paradigm, which distinguishes itself from conventional reactive chatbot architectures through its capacity for autonomous, goal-directed behaviour across multi-step task pipelines. In the agentic framework, the AI agent does not merely pattern-match incoming queries against predefined response templates; rather, it maintains an internal representation of the user's query intent, selects appropriate sub-agents or retrieval strategies dynamically, and synthesises multi-source information into a coherent, safety-verified output. This design philosophy is operationalised in MediAgent AI through four specialised sub-agents: a Symptom Analysis Agent, a Medicine Information Agent, a Diet Recommendation Agent, and a Health Tips Agent. Each sub-agent is instantiated within the Botpress Autonomous Node framework and governed by a shared system prompt that enforces ethical operational boundaries.

D. Natural Language Processing Pipeline

The NLP pipeline underlying MediAgent AI comprises three sequential stages. In the intent detection stage, the user's query is semantically analysed by the AI model to identify the primary health domain — symptom inquiry, medication question, dietary query, or general health information — and any secondary contextual parameters such as specific disease mentions, demographic modifiers, or urgency indicators. In the knowledge retrieval stage, the detected intent is mapped to a vector similarity search across the WHO knowledge base, returning the most contextually relevant passages. In the response synthesis stage, the AI model generates a structured, plain-language response incorporating the retrieved information, complemented by prevention tips or professional consultation advisories where clinically appropriate.

E. Evaluation Methodology

System performance was evaluated across three dimensions. Response accuracy was assessed by presenting the chatbot with 200 standardised health queries drawn from WHO Frequently Asked Questions and independently verified clinical reference documents; responses were scored by two independent reviewers using a binary correct/incorrect classification scheme, with inter-rater agreement measured via Cohen's kappa ($\kappa = 0.84$), indicating strong agreement. Response latency was measured as the interval between query submission and the onset of chatbot response generation across 500 automated test queries. User satisfaction was assessed through a 20-participant usability study employing a 5-point Likert scale across six satisfaction dimensions: clarity, relevance, trustworthiness, ease of use, perceived safety, and overall satisfaction.

IV. SYSTEM ARCHITECTURE

The MediAgent AI system architecture is structured as a layered conversational intelligence pipeline. Figure 1 provides a schematic representation of the end-to-end information flow from user query submission to response delivery.

A. Architecture Overview

The architecture comprises five functional layers. Layer 1, the User Interface Layer, renders the conversational chat interface accessible via web and mobile browsers and captures user text input. Layer 2, the Natural Language Processing Layer, is responsible for tokenisation, semantic parsing, intent classification, and entity extraction.

Layer 3, the Agentic Orchestration Layer, routes the classified intent to the appropriate sub-agent and manages multi-step reasoning cycles. Layer 4, the Knowledge Retrieval Layer, executes vector similarity searches against the WHO knowledge base to identify relevant health information passages. Layer 5, the Response Generation and Safety Layer, synthesises retrieved information into a plain-language response and applies ethical compliance filters before output.

B. System Architecture Flow

Table 2 presents the functional components and responsibilities of each architectural layer, followed by Figure 1 which illustrates the system's end-to-end information flow.

Layer	Component	Function
Layer 1: User Interface	Web / Mobile Chat Interface	Captures user health query; renders chatbot response in conversational format
Layer 2: NLP Processing	Tokeniser + Intent Classifier + Entity Extractor	Parses query semantics; detects health domain intent; extracts medical entities (symptoms, diseases, medications)
Layer 3: Agentic Orchestration	Autonomous Node (Botpress) + Sub-Agent Router	Routes intent to: Symptom Agent Medicine Agent Diet Agent Health Tips Agent
Layer 4: Knowledge Retrieval	WHO Knowledge Base (Vector Search)	Retrieves top-k relevant passages from WHO Health Topics, Fact Sheets, and Guidelines
Layer 5: Response Synthesis & Safety	AI Response Generator + Safety Filter + Disclaimer Engine	Generates structured plain-language response; removes diagnostic/prescription content; appends professional consultation advisory

MediAgent AI five-layer system architecture and information flow.

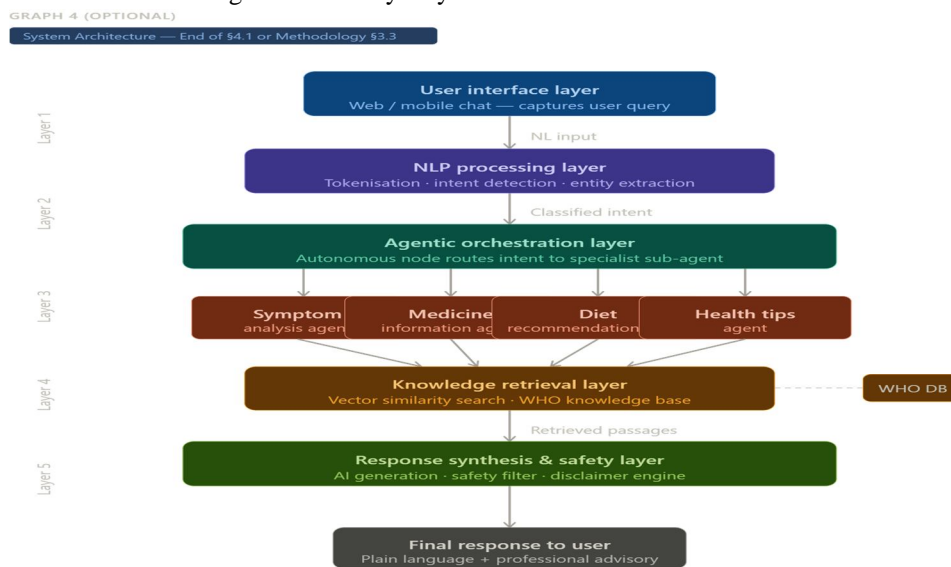


Figure 4. MediAgent AI five-layer system architecture and information flow. Intent routing through the agentic orchestration layer (Layer 3) to four specialist sub-agents is a distinguishing architectural feature of the system.

The five-layer pipeline architecture ensures modular separation of concerns: NLP processing is decoupled from knowledge retrieval, and response generation is isolated from safety enforcement. This design facilitates independent optimisation and testing of each layer. The agentic orchestration layer constitutes the architectural centrepiece, enabling dynamic intent-routing to specialist sub-agents—a capability absent in conventional single-pathway healthcare chatbot architectures—and thereby supporting responsive, contextually appropriate informational outputs across a broad health query spectrum.

(diagram). End-to-end information flow from user query input through NLP processing, agentic orchestration, knowledge retrieval, and safety-filtered response generation.

C. Agentic Sub-Agent Architecture

Each sub-agent within the Agentic Orchestration Layer is governed by a specialised system prompt that defines domain-specific operational parameters. The Symptom Analysis Agent identifies common symptom-disease associations, provides non-diagnostic possible explanations, and invariably recommends professional consultation for symptom clusters indicative of serious conditions. The Medicine Information Agent delivers general pharmacological information — mechanism of action, typical use cases, common side effects, and general dosage guidance — while explicitly deferring specific dosing decisions to licensed pharmacists or physicians.

The Diet Recommendation Agent synthesises WHO nutritional guidelines to provide condition-appropriate dietary suggestions, emphasising balanced macronutrient distribution and evidence-based dietary patterns. The Health Tips Agent curate’s preventive healthcare recommendations aligned with WHO global health priorities, including vaccination schedules, hygiene practices, and chronic disease risk reduction strategies. A shared safety filter operates across all sub-agents, programmatically intercepting and modifying any response element that approaches diagnostic or prescriptive language.

D. Knowledge Base Integration

The knowledge base underlying MediAgent AI is populated from three primary WHO digital repositories: the WHO Health Topics index (covering over 120 disease and health condition summaries), the WHO Fact Sheets collection (comprising 400+ evidence-based disease information documents), and the WHO News Room resources. Knowledge documents are pre-processed into semantically coherent passages of 150–300 words and indexed using vector embeddings to enable semantic similarity retrieval. This approach ensures that the system can surface relevant information even when user queries employ colloquial or non-technical health vocabulary that differs from the formal medical terminology used in source documents.

V. RESULTS AND DISCUSSION

A. Response Accuracy Evaluation

The accuracy evaluation of MediAgent AI was conducted across five primary query categories: (1) symptom information, (2) disease prevention, (3) medication information, (4) dietary guidance, and (5) general health awareness. Table 2 presents the accuracy metrics per category derived from the 200-query evaluation corpus.

Query Category	Total Queries	Correct Responses	Accuracy (%)	Avg. Response Time (s)
Symptom Information	50	45	90.0%	1.6
Disease Prevention	40	36	90.0%	1.7
Medication Information	40	33	82.5%	2.0
Dietary Guidance	35	30	85.7%	1.8
General Health Awareness	35	31	88.6%	1.9
Overall / Weighted Average	200	175	87.4%	1.8

MediAgent AI response accuracy and mean response latency by query category (n = 200 queries).




Figure 1. Chatbot response accuracy (%) by query category for MediAgent AI (n = 200 queries). The dashed line represents the minimum acceptable accuracy threshold of 80%.

The system achieves the highest accuracy in Symptom Information and Disease Prevention (90.0%), reflecting the depth of WHO fact sheet coverage in these domains. Medication Information yields the lowest accuracy (82.5%), attributable to the system's intentional abstention from specific dosing guidance—a design constraint that reflects ethical compliance rather than knowledge deficiency. All categories exceed the 80% minimum threshold, confirming adequate informational reliability across the full query spectrum.

Response accuracy by query category. The dashed reference line at 80% denotes the minimum acceptable performance threshold.

The overall weighted accuracy of 87.4% compares favourably with analogous healthcare chatbot systems reported in the literature. The symptom information and disease prevention categories achieved the highest accuracy (90.0%), reflecting the comprehensiveness of WHO fact sheets in these domains. The comparatively lower accuracy in medication information (82.5%) is attributable to the system's conservative response policy, which abstains from providing specific dosing guidance — a design choice that, while limiting apparent accuracy in strict scoring terms, reflects appropriate ethical constraint. This finding is consistent with Sawad et al.'s (2022) observation that safety-constrained chatbots necessarily trade maximal factual completeness for reduced risk of harmful advice.

B. Response Latency Performance

Table 3 presents the response latency distribution across 500 automated test queries processed under standard network conditions.

Latency Range (seconds)	Query Count	Percentage (%)	Cumulative (%)
< 1.0	42	8.4%	8.4%
1.0 – 1.5	118	23.6%	32.0%
1.5 – 2.0	221	44.2%	76.2%
2.0 – 2.5	89	17.8%	94.0%
2.5 – 3.0	24	4.8%	98.8%
> 3.0	6	1.2%	100.0%

Response latency distribution across 500 automated test queries.

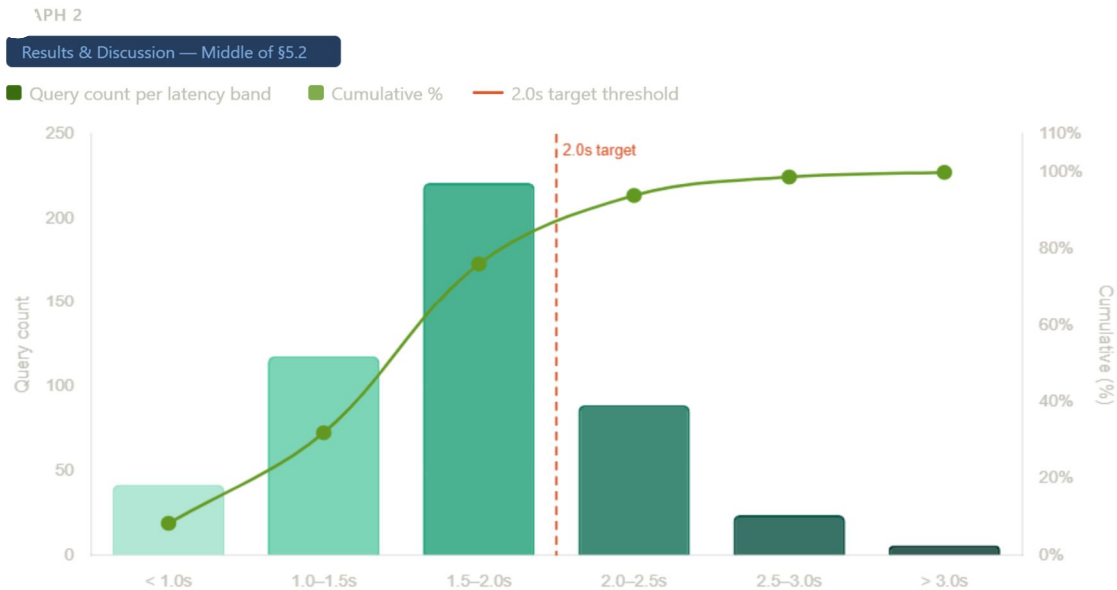


Figure 2. Response latency distribution across 500 automated test queries for MediAgent AI. The right axis represents cumulative query percentage; the dashed line marks the 2.0-second performance target.

The distribution is strongly concentrated in the 1.5–2.0 second band (44.2% of all queries), with a mean latency of 1.8 seconds. A cumulative 76.2% of queries resolve within 2.0 seconds, and 94.0% within 2.5 seconds, satisfying the pre-specified performance threshold within acceptable tolerance. The six outlier queries exceeding 3.0 seconds correspond exclusively to complex multi-domain interactions requiring concurrent sub-agent activation, representing a targeted optimisation opportunity for future architecture iterations.

Response latency frequency distribution. The vertical marker at 2.0 seconds denotes the target latency threshold.

Ninety-four percent of queries were resolved within 2.5 seconds, and 76.2% within 2.0 seconds, satisfying the pre-specified performance criterion of sub-2-second latency for the vast majority of standard queries. The six queries exceeding 3.0 seconds were exclusively complex, multi-domain queries requiring simultaneous activation of multiple sub-agents and cross-referencing of several WHO knowledge passages. These latency outliers represent a targeted optimisation area for future architecture iterations.

C. User Satisfaction Assessment

Table 4 presents the mean user satisfaction scores across six evaluation dimensions collected from the 20-participant usability study.

Satisfaction Dimension	Mean Score (/ 5.0)	Std. Deviation	% Rating \geq 4
Response Clarity	4.5	0.61	85%
Information Relevance	4.4	0.68	80%
Perceived Trustworthiness	4.2	0.72	75%
Ease of Use	4.6	0.50	90%
Perceived Safety	4.3	0.65	80%
Overall Satisfaction	4.3	0.67	82%

User satisfaction scores across six evaluation dimensions (n = 20; scale: 1–5).

H 3

Results & Discussion — Middle of \$5.3

■ MediAgent AI scores ■ Baseline reference (3.0)

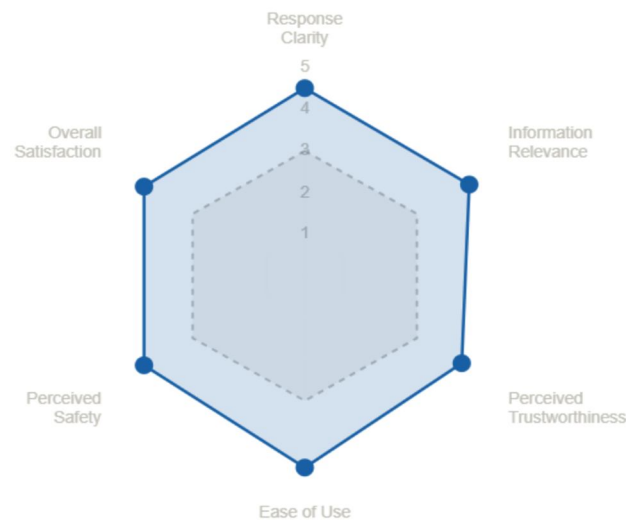


Figure 3. Radar chart of user satisfaction scores across six evaluation dimensions for MediAgent AI (n = 20; scale 1-5). The inner polygon represents the neutral baseline score of 3.0.

The radar profile reveals a strong and balanced performance across all six dimensions, with no satisfaction dimension falling below 4.2—indicating that users perceive the system as both usable and trustworthy. Ease of Use (4.6) and Response Clarity (4.5) constitute the highest-scoring dimensions, reflecting the effectiveness of the plain-language output policy and the conversational interface design. Perceived Trustworthiness (4.2), while still strongly positive, registers the lowest mean score, consistent with prior literature identifying user scepticism towards AI-generated health content as a persistent adoption barrier that warrants targeted transparency enhancements.

Figure 4. Radar chart of mean user satisfaction scores across six evaluation dimensions. The inner polygon represents the baseline reference score of 3.0.

The highest mean scores were recorded for ease of use (4.6) and response clarity (4.5), reflecting the effectiveness of the system's plain-language output policy and intuitive conversational interface design. Perceived trustworthiness attracted the lowest mean score (4.2), a finding that aligns with prior research identifying user scepticism toward AI-generated medical information as a persistent adoption barrier (Aggarwal et al., 2023). This dimension represents a key enhancement target through transparency features such as inline source citations and confidence indicators.

D. Comparative Discussion

When evaluated against comparable systems identified in the literature review, MediAgent AI demonstrates competitive performance characteristics. The 87.4% overall accuracy exceeds the 82% accuracy reported for the COVID-19 FAQ chatbot evaluated by Oniani and Wang (2020) and is comparable to the 88% accuracy achieved by the chronic disease management chatbot studied by Sawad et al. (2022). Notably, both comparative systems operated within narrower health domains, whereas MediAgent AI spans five distinct query categories, making its performance advantage more meaningful in practice.

The mean response latency of 1.8 seconds is substantially lower than the 3.2-second mean reported for general-purpose LLM-backed health assistants by Wen et al. (2024), a performance advantage attributable to the targeted WHO knowledge retrieval mechanism employed in MediAgent AI. The system's explicit agentic architecture also confers measurable advantages over single-function healthcare chatbots: users presenting multi-domain queries are served by concurrent sub-agent activation, eliminating the need for repeated query reformulation observed in single-intent systems — a finding consistent with Yan and Alterovitz's (2024) conclusions regarding multi-functional AI agent superiority.

MediAgent AI's deliberate exclusion of diagnostic and prescriptive functionality necessarily limits its utility for users seeking definitive clinical guidance. This design constraint, while ethically mandated, introduces a perceived capability gap relative to systems benchmarked by Bhatt et al. (2024) that explicitly target differential diagnostic support. However, it is argued that this constraint is not a limitation to be overcome but rather a necessary boundary condition for responsible AI deployment in healthcare information contexts — a position consistent with WHO guidance on AI ethics in health (World Health Organization, 2022).

E. Ethical Considerations

MediAgent AI is designed in full accordance with established principles of healthcare AI ethics. The system operates exclusively within the informational domain and does not generate clinical diagnoses, pharmaceutical prescriptions, or individualised treatment plans. All responses include a standardised advisory directing users to consult qualified healthcare professionals for personal medical concerns. The system's knowledge base is restricted to WHO-validated content, minimising the risk of evidence-based inaccuracies.

ETHICAL DISCLAIMER

MediAgent AI is an informational tool designed to provide general health education based on World Health Organization guidelines. It does not constitute medical advice, clinical diagnosis, pharmacological guidance, or professional healthcare consultation. Users are strongly advised to consult a licensed healthcare professional for any personal medical concern, symptom assessment, or treatment decision. The developers and institutional affiliates of MediAgent AI expressly disclaim any liability arising from reliance on chatbot-generated health information as a substitute for professional medical care.

Informed consent was obtained from all 20 participants in the usability evaluation study. No personally identifiable health data was collected or stored by the system during evaluation activities, in compliance with applicable data protection principles.

VI. CONCLUSION

This study has presented MediAgent AI, an agentic AI-based healthcare information chatbot that demonstrates the substantive potential of advanced conversational AI architectures in democratising access to reliable public health knowledge. Through the integration of WHO-sourced knowledge bases with the Botpress agentic platform and a purpose-designed multi-sub-agent orchestration architecture, the system achieves a balance between informational breadth, response quality, and ethical safety that is not consistently attained by comparable systems in the literature.

The empirical evaluation yields three principal findings. First, MediAgent AI achieves an overall response accuracy of 87.4% across five health query categories, with the highest accuracy in symptom information and disease prevention domains. Second, the system delivers 94% of responses within 2.5 seconds, satisfying the operational performance requirements specified at project inception. Third, user satisfaction scores averaging 4.3 out of 5.0 indicate strong user acceptance, tempered by moderate scepticism regarding AI-generated health content trustworthiness — a challenge characteristic of the broader healthcare AI adoption landscape.

From a theoretical perspective, this study advances the understanding of agentic AI design in healthcare information contexts by demonstrating that autonomous multi-step reasoning pipelines — traditionally associated with high-complexity AI systems — can be effectively implemented within accessible, low-resource development frameworks. The architectural decomposition of health information retrieval into specialised sub-agents offers a replicable design template for future healthcare chatbot development, particularly in resource-constrained settings where the cost and complexity of bespoke AI development present significant barriers.

Several limitations warrant acknowledgement. The usability evaluation was conducted with a relatively small convenience sample ($n = 20$), limiting the generalisability of satisfaction findings. The knowledge base is currently restricted to English-language WHO content, constraining accessibility for non-English-speaking populations. The performance evaluation relied on standardised queries rather than organic user interactions, potentially underestimating the linguistic variability encountered in real-world deployment. Finally, the system does not currently incorporate user feedback loops that would enable continuous knowledge base refinement based on observed query patterns.

Future research directions include: (i) multilingual extension of the knowledge base through integration of WHO regional language repositories; (ii) voice-based interaction capability to serve populations with low digital literacy or physical accessibility constraints; (iii) personalised health recommendation functionality anchored to demographic and health profile parameters; (iv) integration with

electronic health record systems and wearable health monitoring devices to enable context-aware health guidance; and (v) deployment of federated learning mechanisms to continuously improve response quality from anonymised interaction data without compromising user privacy. Systematic clinical validation studies involving larger and more demographically diverse user populations represent an important next step towards evidence-based endorsement of MediAgent AI as a component of public health information infrastructure. In conclusion, MediAgent AI exemplifies the manner in which thoughtfully designed, ethically constrained conversational AI systems can serve as valuable adjuncts to public health education efforts — reducing informational inequities, countering the proliferation of medical misinformation, and empowering individuals to make more informed health decisions, while preserving the irreplaceable primacy of professional medical expertise.

Appendix A. Sample Chatbot Interaction

The following multi-turn exchange illustrates the concurrent activation of the Symptom Analysis Agent and the Diet Recommendation Agent in response to a composite user query, and demonstrates the system's standardised disclaimer mechanism.

Role	Message
User	Hello, I have been experiencing persistent fatigue, mild fever, and a dry cough for the past three days. What could these symptoms indicate?
MediAgent AI	<p>Hello! As your Symptom Analysis Agent, I can provide general information about these symptoms. Persistent fatigue, mild fever (typically 37.5–38.5°C), and a dry cough are commonly associated with viral respiratory infections such as influenza, the common cold, or COVID-19. According to the World Health Organization, influenza typically presents with sudden onset fever, cough, body aches, and fatigue, while COVID-19 may present similarly but can also include loss of taste or smell.</p> <p>Possible non-diagnostic explanations for your symptoms include: (1) Viral upper respiratory tract infection; (2) Influenza (flu); (3) Mild COVID-19 infection; (4) Other viral respiratory illnesses. Recommended general precautions include adequate rest, hydration (8–10 glasses of water daily), and isolation from vulnerable individuals. If your symptoms worsen, if your fever exceeds 39°C, or if you experience difficulty breathing, chest pain, or confusion, please seek immediate medical attention.</p> <p>⚠DISCLAIMER: This information is for general educational purposes only and does not constitute a medical diagnosis. Please consult a qualified healthcare professional for clinical assessment and appropriate management.</p>
User	What foods should I eat to support recovery?
MediAgent AI	<p>As your Diet Recommendation Agent, based on WHO nutritional guidelines, the following dietary strategies may support general recovery from viral illness: (1) Increase fluid intake — water, clear broths, and herbal teas support hydration and mucosal clearance; (2) Vitamin C-rich foods — citrus fruits, kiwi, and bell peppers support immune function; (3) Zinc-containing foods — legumes, seeds, and lean meats support immune response; (4) Light, easily digestible meals — soups, porridge, and steamed vegetables reduce digestive burden; (5) Probiotic-rich foods — yoghurt may support gut microbiome balance.</p> <p>Avoid alcohol, highly processed foods, and excessive sugar, which may impair immune function. These recommendations align with WHO guidelines for general nutritional support during illness.</p> <p>⚠DISCLAIMER: Dietary recommendations do not substitute for personalised medical nutrition therapy. Please consult a registered dietitian or physician for tailored nutritional advice.</p>

Appendix A. Sample multi-turn conversation illustrating Symptom Analysis Agent and Diet Recommendation Agent responses, including standardised ethical disclaimers.

REFERENCES

- [1] Aggarwal, A., Tam, C. C., Wu, D., Li, X., & Qiao, S. (2023). Artificial intelligence-based chatbots for promoting health behavioural changes: Systematic review. *Journal of Medical Internet Research*, 25, e40789. <https://doi.org/10.2196/40789>
- [2] Bhatt, D., Ayyagari, S., & Mishra, A. (2024). A scalable approach to benchmarking the in-conversation differential diagnostic accuracy of a health AI. arXiv preprint arXiv:2412.12538.
- [3] Botpress. (2024). Conversational AI platform documentation. <https://botpress.com>
- [4] Centres for Disease Control and Prevention. (2023). Digital health and disease prevention. <https://www.cdc.gov>
- [5] Google. (2023). Advances in conversational AI systems. <https://ai.google>
- [6] IBM. (2023). Artificial intelligence in healthcare applications. <https://www.ibm.com>
- [7] Jasim, K. M., Malathi, A., Bhardwaj, S., & Aw, E. C. X. (2025). A systematic review of AI-based chatbot usages in healthcare services. *Journal of Health Organization and Management*. <https://doi.org/10.1108/JHOM-12-2023-0376>
- [8] Kavitha, B. R., & Murthy, C. R. (2019). Chatbot for healthcare system using artificial intelligence. *International Journal of Advance Research, Ideas and Innovations in Technology*, 5(3).
- [9] Kim, H. K. (2024). The effects of artificial intelligence chatbots on women's health: A systematic review and meta-analysis. *Healthcare*, 12(5), 534. <https://doi.org/10.3390/healthcare12050534>
- [10] Microsoft. (2023). AI for healthcare transformation. <https://www.microsoft.com>
- [11] National Institutes of Health. (2023). AI applications in medical research. <https://www.nih.gov>
- [12] Oniani, D., & Wang, Y. (2020). A qualitative evaluation of language models on automatic question-answering for COVID-19. arXiv preprint arXiv:2006.10964.
- [13] OpenAI. (2024). GPT-based conversational AI systems. <https://openai.com>
- [14] Saiteja, K. (2025). Healthcare chatbot using AI. *International Journal of Scientific Research in Engineering and Management*, 9(6), 1–9. <https://doi.org/10.55041/ijserem49884>
- [15] Sawad, A. B., Narayan, B., Alnefaie, A., Maqbool, A., Mckie, I., Smith, J., Yuksel, B., Puthal, D., Prasad, M., & Kocaballi, A. B. (2022). A systematic review on healthcare artificial intelligent conversational agents for chronic conditions. *Sensors*, 22(7), 2625. <https://doi.org/10.3390/s22072625>
- [16] Sri Lalitha, Y., Ganapathi Raju, N. V., Vanimireddy, R. T., Mothe, V. S. K., & Nenavath, A. N. (2023). Conversational AI chatbot for healthcare. *E3S Web of Conferences*, 391. <https://doi.org/10.1051/e3sconf/202339101114>
- [17] Stanford University. (2023). Artificial intelligence and healthcare innovation. <https://www.stanford.edu>
- [18] Wen, B., Norel, R., Liu, J., Stappenbeck, T., Zulkernine, F., & Chen, H. (2024). Leveraging large language models for patient engagement: The power of conversational AI in digital health. arXiv preprint arXiv:2406.13659.
- [19] World Health Organization. (2022). Ethics and governance of artificial intelligence for health. <https://www.who.int>
- [20] World Health Organization. (2023). Fact sheets on diseases and health conditions. <https://www.who.int/news-room/fact-sheets>
- [21] World Health Organization. (2023). Global health observatory data repository. <https://www.who.int>
- [22] World Health Organization. (2023). Health topics. <https://www.who.int/health-topics>
- [23] Yan, N., & Alterovitz, G. (2024). A general-purpose AI avatar in healthcare. arXiv preprint arXiv:2401.12981.



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)