International Journal of Health Technology and Innovation
An Official Journal of **Kalam Institute of Health Technology** | ISSN - **2583-8547**
RESEARCH ARTICLE

Deployment of LLM-Powered Healthcare Information Kiosks in Primary Health Centers in India: Technical, Legal, and Ethical Perspectives

Rishi Manoj[1], Gourav Singh[2] Siva Shankara Raju Vysyaraju [2], Navni Krishna[3], Darsana G[4]

[1]Design & Development Flinders Biomedical Enterprises Hyderabad, India
[2]Techpod Andhra Pradesh MedTech Zone Visakhapatnam, India
[3]School of Legal Studies Cochin University of Science and Technology (CUSAT) Kochi, India
4Biomedical Engineering Chennai Institute of Technology Chennai, India

***Received:** 19th January , 2026; **Revised:** 2st February, 2026; **Accepted:** 24th March, 2026; **Available Online:** 04th April, 2026*

ABSTRACT

Access to accurate health information in India is often limited by language barriers, digital divides, and a shortage of qualified healthcare professionals at the Primary Health Centers (PHCs) and community awareness hubs. This paper proposes the deployment of localized Large Language Models (LLMs) to enable multilingual, conversational health communication for patients and caregivers. The solution combines interactive, low-cost kiosks installed at PHCs and public health awareness centers with a web-based application for at-home access, allowing individuals to inquire about common diseases, symptoms, preventive care, and treatment guidance in their native languages. The system is designed to operate offline using locally hosted LLMs in areas with poor internet connectivity, while seamlessly switching to cloud APIs when connectivity is available. By bridging language and access gaps, this approach aims to enhance early disease awareness, patient engagement, and healthcare equity, supporting the government's vision for digitally inclusive primary healthcare in India. The paper discusses the architecture, hardware specifications for kiosks, language model selection and fine-tuning strategies, as well as privacy and data security considerations necessary for safe and responsible AI deployment in public health environments.

Keywords: Large Language Models (LLMs), Primary Health Centers (PHCs), Multilingual Healthcare Communication, Low-Cost Interactive Kiosks, Offline AI Deployment, Public Health Awareness

International Journal of Health Technology and Innovation (2026)

How to cite this article: Manoj R, Singh G, Vysyaraju SSR, Krishna N, Darsana G. Deployment of LLM-Powered Healthcare Information Kiosks in Primary Health Centers in India: Technical, Legal, and Ethical Perspectives. International Journal of Health Technology and Innovation. 2026;5(1):56-67.

Doi: 10.60142/ijhti.v5i01.08

Source of support: Nil.

Conflict of interest: None

INTRODUCTION

Primary Health Centers (PHCs) form the backbone of India's public healthcare system, serving as the first point of contact between communities and formal medical services, especially in rural and semi-urban regions [1]. These centers are responsible for delivering essential healthcare services, including preventive care, maternal and child health, and basic disease management. Despite their critical role, PHCs face persistent structural and operational challenges that limit their ability to effectively disseminate healthcare information to the population they serve.

One of the most significant gaps in healthcare delivery in India is the high doctor–patient ratio [2], which places immense pressure on medical professionals and restricts the time available for patient education and counseling. This constraint is further compounded by linguistic diversity [3], as patients often speak regional or local languages that differ from those commonly used in medical documentation or consultations. As a result, critical health information related to disease prevention, symptoms, and treatment adherence may not be effectively communicated or understood. Additionally, low health literacy levels in many communities hinder individuals' ability to access, interpret, and act upon available healthcare information, leading to delayed diagnosis and preventable health complications.

Advances in Large Language Models (LLMs) present

**Author for Correspondence:* rishimanoj@ieee.org

an opportunity to address these challenges by enabling natural, multilingual, and conversational access to healthcare information [4]. When integrated into low-cost, interactive kiosks, LLMs can serve as scalable digital assistants within PHCs, allowing patients to seek reliable health information in their native languages without increasing the workload of healthcare staff. The inclusion of a complementary web-based platform further extends access to individuals outside physical healthcare facilities, supporting continuity of health awareness at home.

The primary objective of this work is to design and propose an LLM-powered healthcare information kiosk system that enhances disease awareness, health education, and patient engagement at PHCs and community health centers [5]. The system is intended to function in both offline environments using locally hosted models and online settings via cloud-based APIs, ensuring usability in regions with limited or unreliable internet connectivity.

The scope of this study is strictly limited to health information dissemination and awareness. The proposed system does not provide medical diagnoses, treatment prescriptions, or emergency decision-making support. Instead, it is designed as a supplementary tool to assist patients in understanding health conditions and to encourage timely consultation with qualified healthcare professionals. Ethical considerations, including data privacy, responsible AI usage, and clear communication of system limitations, are integral to the proposed framework [6]

LITERATURE REVIEW

Existing Healthcare Kiosks & Telemedicine Systems

Health kiosks have been studied as low-cost, point-of-care tools that improve access to information, screening and simple monitoring tasks [7]. Scoping and systematic reviews conclude kiosks are multifunctional and cost-effective when integrated with health systems, but highlight implementation barriers such as usability, maintenance, connectivity, and integration with clinical workflows.

In India specifically, telemedicine and kiosk-like interventions have been used to extend primary care (maternal/child health, chronic disease follow-up, remote consultations) especially after the expansion of ISRO/telemedicine initiatives and the COVID-19 acceleration of digital care; reviews of Indian telemedicine interventions emphasize promise but note uneven infrastructure, workforce, and governance readiness [8,9]

AI & Chatbot Applications in Healthcare

A growing body of work evaluates AI chatbots for health education, symptom triage, mental-health support, and administrative tasks. Recent systematic and narrative reviews document that chatbots can increase reach and engagement, but they also repeatedly flag limitations: variable accuracy, hallucinations or misleading guidance, limited clinical validation, and safety/regulatory concerns when chatbots stray toward diagnostic or prescriptive behaviour. Standardization and evaluation frameworks for healthcare chatbots are emerging (proposals and early frameworks aim to scaffold safety, transparency, and monitoring) [10,11]

Multilingual Health Information Systems in India

Language and literacy barriers are a recognized bottleneck for digital health uptake in India. Recent projects and academic work have explored NLP and multilingual systems for patient-facing health content, including voice/text translators, localized health chatbots, and efforts from research labs and startups to build Indian language corpora and models. These initiatives demonstrate technical feasibility of multilingual conversational agents but report gaps in coverage across low-resource languages, culturally appropriate phrasing, and evaluation with real users in rural PHC settings. Industry solutions for localisation (translation/localisation vendors and platforms) exist but are often proprietary and not designed for offline, kiosk-based deployment [12,13]

Ethical, Regulatory & Governance Context

India's biomedical and AI governance bodies have begun to release guidance specific to AI in healthcare — stressing ethics, data privacy, risk classification, and the need for human oversight and validation in clinical contexts. National guidelines and early policy documents emphasize that AI tools intended for health should be transparent about limitations and should not replace qualified clinical judgment. Recent news and reviews also show real harms can arise when conversational AI provides misleading medical content, reinforcing the need for conservative scope and strong governance for patient-facing systems [6,14]

Gaps in Current Solutions

From the literature above we identify four recurring gaps directly relevant to deploying LLM-powered kiosks at PHCs in India:

Language & literacy coverage

many deployed solutions and commercial platforms lack robust support for India's many low-resource languages and dialects, and they often provide literal translations that fail to account for cultural and health-literacy differences.

Offline/edge deployment

most advanced AI chatbots assume persistent connectivity or cloud APIs; there is limited published work on resilient, locally hosted LLMs tuned for constrained hardware and intermittent networks typical of PHCs.

Clinical safety and evaluation

reviews stress that few conversational health systems have undergone rigorous clinical validation or prospective user-studies in real PHC environments; evaluation frameworks are nascent.

Governance & ethical safeguards

while policy guidance is emerging, implementation-level governance (consent, data minimization, clear non-diagnostic disclaimers, logging and escalation pathways) is inconsistently applied in reported systems. Recent incidents where AI produced harmful medical assertions underline this gap.

The literature shows that kiosks and telemedicine can extend access, and that chatbots/LLMs offer novel multilingual, conversational capabilities — but published work still lacks practical, evaluated examples of locally deployed, multilingual LLM systems designed for PHCs that combine offline operation, rigorous safety evaluation, and enforceable ethical governance [15]. This gap motivates our study: designing, implementing and evaluating a low-cost, LLM-powered, multilingual kiosk + web-app solution that is explicitly non-diagnostic, locally deployable, and aligned with emerging Indian AI/health guidelines.

SYSTEM OVERVIEW OF THE PROPOSED KIOSK

Overall System Architecture

The proposed system follows a modular AI-driven conversational architecture, as illustrated in Fig.1, integrating user interaction, language processing, response generation, and multimodal output delivery [16]. The system is designed to function as an interactive healthcare information kiosk deployed at Primary Health Centers (PHCs), with optional web-based accessibility for remote users. The architecture ensures scalability, multilingual capability, and support for both offline and online operational modes.

The interaction begins at the User Input Layer, where patients provide queries through voice or text using a microphone-enabled interface or touchscreen keyboard. The interface is implemented using a lightweight graphical framework (e.g., PyQt), which manages input capture and display rendering [17,18]. Voice inputs are processed through speech recognition modules, while text inputs are directly forwarded for processing. The system is designed to support regional Indian languages to address linguistic diversity in rural and semi-urban areas.

Following input capture, the query is transmitted to a Language Detection Module, which identifies the user's language. This component ensures appropriate routing of the query for processing in the detected language or performs translation if required. Language detection enables seamless multilingual interaction and enhances accessibility for users with limited proficiency in English or Hindi [19,20]

The processed query is then passed to the Natural Language Processing (NLP) L**ayer**, powered by a Large Language Model (LLM). This layer performs contextual understanding of the health-related query and generates an appropriate informational response. The LLM module is configured to operate in two modes: (i) a locally hosted model for offline environments where internet connectivity is limited, and (ii) an API-based cloud model for enhanced performance when connectivity is available. The model is constrained to provide awareness-oriented, non-diagnostic health information aligned with verified public health guidelines [21,22]

Once a response is generated, it is forwarded to the Text-

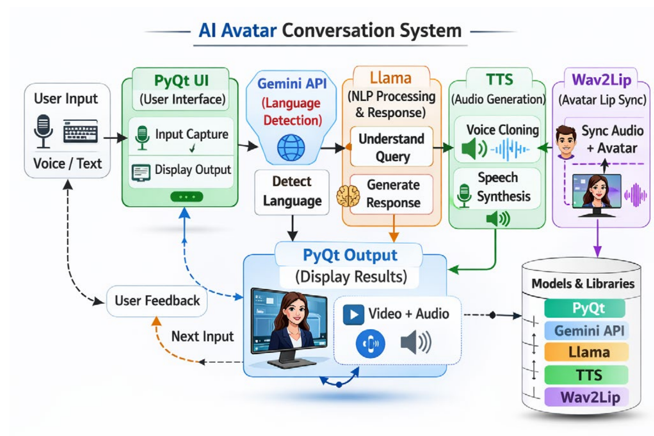

Fig.1: System Architecture of the Kiosk

to-Speech (TTS) Module, which converts textual output into natural-sounding speech. The system may incorporate optional voice personalization features to improve user engagement. For enhanced interaction, the speech output can be synchronized with a digital avatar using lip-synchronization techniques, enabling audiovisual communication. This multimodal presentation improves comprehension, particularly for low-literacy users [23]

The final output is rendered through the Display and Output Layer, which presents synchronized video and audio responses on the kiosk screen. The system also incorporates a feedback loop, allowing users to ask follow-up questions or refine their queries, thereby enabling conversational continuity. Additionally, an administrative backend manages system logs, content updates, model configuration, and usage monitoring while maintaining privacy-preserving data practices.

Overall, the architecture emphasizes modular integration of user interface components, multilingual processing, LLM-based reasoning, speech synthesis, and audiovisual rendering to create an inclusive and scalable healthcare information delivery system.

Stakeholders

The proposed LLM-powered healthcare information kiosk involves multiple stakeholders whose coordinated participation is essential for effective deployment and sustained impact.

Patients

Patients constitute the primary users of the system. The kiosk enables them to access reliable information regarding common diseases, symptoms, preventive measures, maternal and child health, vaccination schedules, hygiene practices, and chronic disease management. By providing responses in their native language through conversational voice interaction, the system reduces communication barriers and improves health awareness among rural and underserved populations [24]

Family Members and Caregivers

In many Indian households, healthcare decisions are collectively made by family members. The kiosk supports caregivers by providing accessible information related to disease prevention, home-based care practices, nutrition, and early symptom recognition. This promotes informed decision-making at the household level and encourages timely medical consultation [25]

Healthcare Staff: Doctors, nurses, and community health workers at PHCs benefit indirectly from the system. By handling frequently asked informational queries, the kiosk reduces repetitive counseling workload and allows healthcare professionals to focus on diagnosis and treatment. Additionally, it can reinforce public health campaigns by delivering standardized, guideline-based awareness content consistent with national health programs.

Administrators and Public Health Authorities

Administrative stakeholders oversee deployment, maintenance, compliance, and performance monitoring of the system. They are responsible for ensuring alignment with government healthcare policies, updating knowledge resources, managing system analytics, and enforcing data privacy and ethical AI standards. Their role is critical in maintaining the reliability, accountability, and sustainability of the kiosk infrastructure [26,27]

In summary, the proposed system establishes a collaborative ecosystem in which AI-driven multilingual communication enhances healthcare information dissemination while supporting patients, families, healthcare providers, and administrators within the primary healthcare framework.

ABDM Compliance and ABHA Integration

To ensure interoperability with India's national digital health ecosystem, the proposed kiosk architecture can be extended to support integration with the Ayushman Bharat Digital Mission (ABDM) framework. This includes optional linkage with the Ayushman Bharat Health Account (ABHA) system for identity-enabled healthcare services.

The kiosk can incorporate a consent-driven authentication mechanism wherein users may voluntarily provide their ABHA ID to access or link health records. Upon authentication, the system can interface with ABDM-compliant APIs such as Health Information Exchange (HIE) and Health Locker services. This enables secure retrieval of patient health records, subject to user consent and regulatory compliance.

The integration layer can be designed as an independent module within the system architecture, ensuring that core kiosk functionality remains operational even in the absence of ABDM connectivity. Fig. 2 illustrates the proposed architecture for the integration of ABDM and ABHA. All data exchanges must adhere to standardized Electronic Health Record (EHR) formats and consent management protocols defined under ABDM guidelines.

It is important to note that such integration is proposed as a future enhancement and requires strict adherence to national data protection regulations, interoperability standards, and certification requirements prior to deployment.

MECHANICAL AND PHYSICAL DESIGN OF THE KIOSK

Physical Form Factor and Ergonomics

The proposed kiosk is designed as a freestanding, vertically oriented interactive terminal suitable for installation within Primary Health Centers (PHCs) and community health awareness facilities. As illustrated in Fig. 3, the proof-of-concept (PoC) model features a stable base platform, dual-supported vertical columns, and an angled display enclosure to ensure optimal visibility and user comfort. As illustrated in Fig. Y, the proof-of-concept (PoC) model features a stable base platform, dual-support vertical columns, and an angled display enclosure to ensure optimal visibility and user comfort. The display panel is positioned at a slight tilt to reduce glare and improve readability under varying lighting conditions commonly observed in rural and semi-urban healthcare settings.

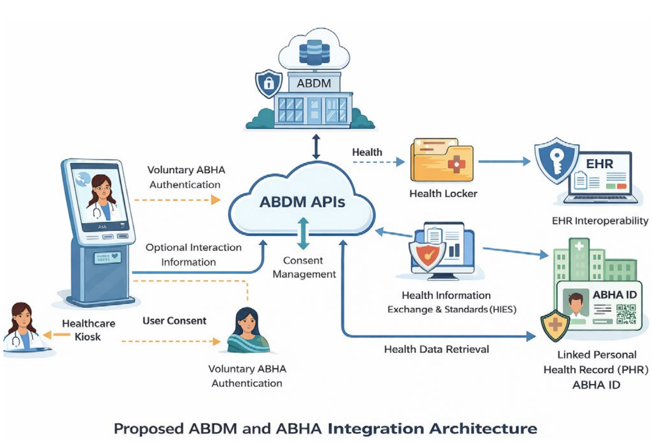

Fig.2: Proposed ABDM and ABHA Integration Architecture

Fig.3: Mechanical Design Mockup of the Kiosk

The overall structure prioritizes compactness, durability, and ease of maintenance. The enclosure is constructed using lightweight yet robust materials such as powder-coated mild steel or reinforced polymer composites to withstand continuous public use. The base plate is dimensioned to prevent tipping while maintaining portability for relocation within PHC premises. Rounded edges and concealed wiring enhance safety and aesthetic appeal, while the front-access service panel allows easy hardware maintenance.

Ergonomically, the screen height is selected to accommodate an average standing adult user while remaining accessible to shorter individuals. The angled interface reduces neck strain and facilitates natural eye-level interaction. The touch interface is designed for minimal force activation, ensuring ease of use for elderly individuals.

ACCESSIBILITY CONSIDERATIONS

Inclusivity is a core design principle of the kiosk system. The mechanical and interface design accounts for diverse user demographics, including elderly individuals and differently abled users.

For elderly users, the system incorporates: Large-font display options and high-contrast UI themes, Voice-based interaction to minimize dependency on typing, Clear audio output with adjustable volume levels, and simple navigation with minimal menu complexity.

The slightly inclined screen positioning reduces glare and enhances readability, particularly for users with visual impairments. For differently abled users, the kiosk design includes: Adequate lower clearance beneath the display for wheelchair accessibility, Touchscreen placement within reachable height range as per accessibility guidelines, Voice input and audio output to support visually impaired individuals, Optional headphone jack for private listening in noisy environments.

The system can be further extended with tactile markers or braille labels on key interface buttons in future iterations.

Hardware Components

The hardware configuration of the kiosk is optimized for cost-efficiency, scalability, and operational reliability in low-resource environments. A capacitive touchscreen display is recommended to provide adequate screen space for multilingual text, avatar-based interaction, and instructional visuals. The display supports FHD resolution for clarity while maintaining energy efficiency. A noise-reducing unidirectional microphone captures voice inputs effectively, even in moderately noisy PHC environments. Integrated stereo speakers deliver clear audio output for conversational responses. Volume control is adjustable to accommodate varying ambient noise levels. The kiosk operates on standard AC power supply with integrated surge protection. An optional battery backup (UPS) ensures short-term continuity during power fluctuations, which are common in rural regions. Connectivity options include Wi-Fi and Bluetooth.

Deployment Suitability for Rural and Semi-Urban PHCs

The kiosk is specifically engineered for deployment in resource-constrained healthcare environments. Its modular design enables straightforward transportation, installation, and servicing. The use of energy-efficient components reduces operational costs, while offline LLM capability ensures uninterrupted functionality in areas with limited internet access [28]

The ruggedized structure tolerates frequent public interaction and environmental variations such as dust and moderate humidity. Maintenance requirements are minimal, and administrative access panels allow authorized personnel to perform updates and diagnostics without extensive technical expertise.

Overall, the mechanical and physical design of the kiosk aligns with the operational realities of rural and semi-urban PHCs, ensuring durability, inclusivity, cost-effectiveness, and ease of deployment. The proof-of-concept demonstrates feasibility for scalable implementation across diverse healthcare settings in India.

LLM-BASED HEALTHCARE INFORMATION MODULE

The LLM-Based Healthcare Information Module constitutes the core intelligence layer of the proposed kiosk system. It is designed to provide reliable, multilingual, and awareness-oriented health information while strictly avoiding diagnostic or prescriptive outputs. The module operates within predefined safety boundaries and leverages curated public health knowledge sources to ensure consistency, accuracy, and regulatory compliance. Its architecture integrates domain-restricted prompting, response validation mechanisms, and uncertainty-aware output generation.

Medical Condition Overview (Non-Diagnostic)

The kiosk provides general explanations of common diseases such as diabetes, hypertension, tuberculosis, malaria, anemia, and respiratory infections. Responses include definitions, causes, risk factors, and basic pathophysiology presented in simplified, regionally understandable language. The system explicitly avoids confirming or ruling out diagnoses and encourages consultation with healthcare professionals.

Symptoms and Early Warning Signs

Users may inquire about typical symptoms associated with specific conditions. The system presents early warning signs, red-flag indicators requiring immediate medical attention, and general guidance on when to visit a PHC. Clear disclaimers are included to prevent self-diagnosis.

Standard Treatment Procedures

The kiosk explains standard treatment pathways as defined in national or state-level public health programs. This may include duration of treatment (e.g., for tuberculosis), follow-up protocols, or vaccination schedules. The information is descriptive rather than prescriptive.

Medicine Information

The system provides awareness-level information regarding commonly available medications at PHCs, including: General purpose of the medicine, Availability under public health schemes, Broad age-based dosage ranges (non-prescriptive and clearly labeled as informational only), Common side effects requiring medical attention. The system avoids recommending specific dosages tailored to individuals and directs users to consult healthcare staff for personalized prescriptions.

Home Remedies and Supportive Care

The kiosk offers safe, culturally accepted, supportive measures for minor conditions (e.g., hydration, rest, steam inhalation for mild cold symptoms). All recommendations are limited to low-risk practices and are accompanied by safety cautions.

Preventive Measures

Preventive healthcare information includes vaccination awareness, nutrition advice, sanitation practices, maternal and child health guidelines, and lifestyle modifications for non-communicable diseases. This category aligns strongly with PHC awareness objectives.

Communicable Disease Safety

For infectious conditions, the system provides: Basic isolation guidance, Steps to protect family members, Hygiene and sanitation recommendations, and masking and ventilation practices where applicable. Such responses are structured to promote public safety without creating panic or misinformation.

RESPONSE GENERATION AND CONTROL

The structure of the system architecture is illustrated with Fig.4. The response generation mechanism of the proposed system is designed to ensure safe, controlled, and domain-restricted dissemination of healthcare information. The objective is to maintain informational integrity while preventing diagnostic, prescriptive, or unsafe outputs.

The interaction process begins with the user's activation of the microphone interface. Upon successful initiation, the system captures voice input and converts it to text using an automatic speech recognition module. A language detection component then identifies the user's language to ensure accurate contextual processing. The detected language is retained for the duration of the session to maintain conversational consistency and reduce repetitive detection overhead.

Following language identification, the processed query is forwarded to the LLM for intent recognition and contextual understanding. The model is guided using structured prompt templates that restrict responses to predefined healthcare awareness categories. These prompts explicitly instruct the model to provide general informational content, avoid diagnostic confirmation, refrain from prescribing medication, and encourage consultation with qualified healthcare professionals when necessary. This structured prompting strategy serves as the primary mechanism for domain control in the absence of external retrieval support.

To enhance safety, guardrail mechanisms are integrated into the response pipeline. Generated outputs are evaluated against predefined constraints to detect medically inappropriate advice, overconfident assertions, unsupported claims, or deviation from the informational scope. If a response violates safety constraints, the system triggers a fallback message that redirects the user to seek assistance from healthcare staff. This layered validation approach reduces the likelihood of unsafe or misleading outputs.

Hallucination mitigation is addressed through domain limitation and controlled output formatting. The model is restricted to operate within a bounded healthcare-awareness context, thereby minimizing speculative elaboration. Response length constraints and template-based structuring further reduce the risk of fabricated statistics, unverified treatment

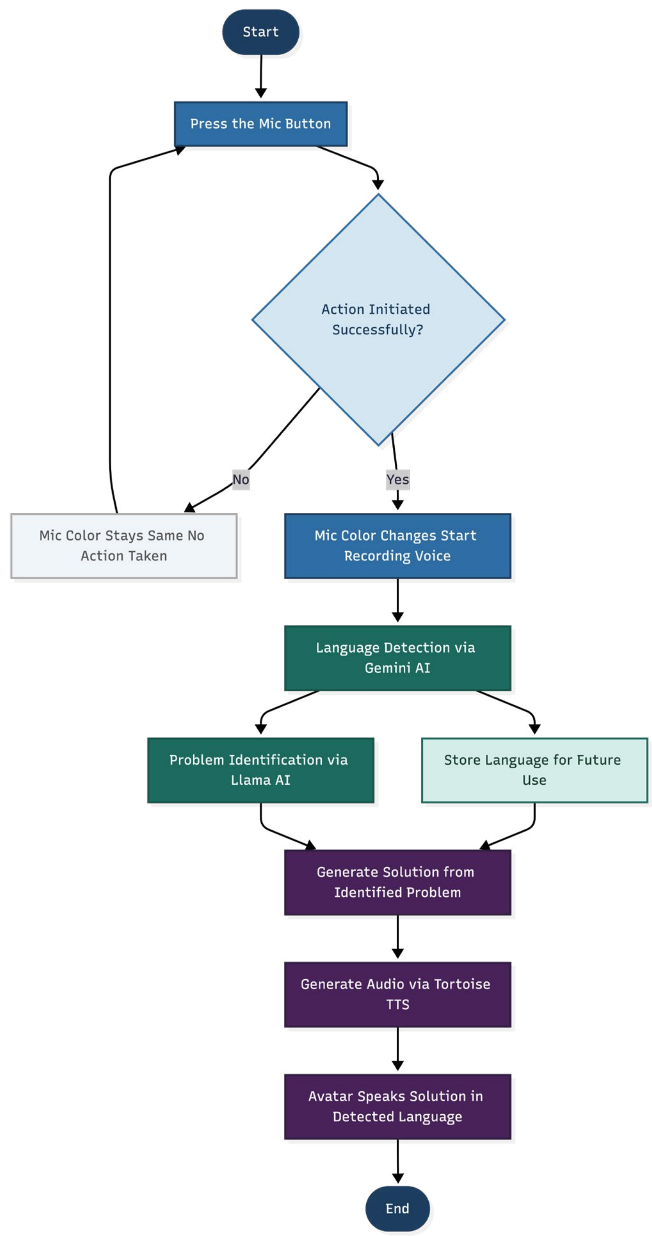

Fig.4: Flow Diagram of the System Algorithm

claims, or unsupported dosage recommendations. Periodic evaluation using curated test queries ensures continued alignment with public health communication standards.

To prevent overconfidence in ambiguous scenarios, the system incorporates uncertainty-aware response behavior. When the input query lacks sufficient context or exhibits ambiguity, the model may request clarification before generating a final response. In situations where confidence is inherently low, the system explicitly communicates uncertainty and emphasizes the need for professional medical consultation. This calibrated approach ensures that the kiosk functions as a supportive informational tool rather than a substitute for clinical judgment.

Finally, the validated textual response is converted into speech using a text-to-speech module and delivered through the audiovisual interface. This controlled end-to-end pipeline—from voice capture to guarded response synthesis—ensures that healthcare information dissemination remains accurate, ethically compliant, and aligned with the intended non-diagnostic scope of the system.

USER INTERACTION FLOW

The user interaction process implemented in the proposed LLM-powered healthcare kiosk is structured to be intuitive, multilingual, and accessible to individuals with varying literacy levels. The operational sequence is represented in the accompanying flow diagram.

The interaction begins with system activation. Upon approaching the kiosk, the user initiates a session by pressing a start button or activating the microphone interface. At the beginning of the session, the system either requests explicit language selection from a predefined list of supported regional languages or automatically detects the spoken language from the initial voice input. The selected or detected language is retained for the duration of the session to ensure conversational consistency.

Following language initialization, the user submits a healthcare-related query through either voice input or text entry via the touchscreen interface. For voice-based interaction, the system captures audio and converts it into text using an automatic speech recognition module. The recognized text is then forwarded to the language model for intent recognition and contextual understanding. For text-based input, the query is directly processed without requiring speech conversion.

Once the query has been analyzed and a validated informational response generated, the output is delivered through an avatar-based audiovisual interface. The textual response is converted into speech using a text-to-speech engine, and the on-screen digital avatar synchronizes lip movements with the synthesized audio. This multimodal delivery enhances comprehension, particularly for users with limited literacy, and provides a conversational experience rather than a static text-based interaction.

The system supports multi-turn dialogue through follow-up clarification prompts. If ambiguity is detected in the user's query, the model may request additional details before producing a final response. Users may also submit subsequent queries within the same session, allowing iterative clarification and progressive understanding of healthcare topics. Session memory is limited to contextual continuity and does not retain personally identifiable information beyond operational requirements.

Sensitive or emergency-related queries are handled through predefined escalation logic. If the system detects references to severe symptoms, life-threatening conditions, or crisis-related keywords, it generates a high-priority advisory response emphasizing immediate medical consultation. The output directs the user to available healthcare staff or emergency services within the PHC. Detailed emergency intervention instructions are not provided beyond basic awareness guidance. A structured referral mechanism ensures appropriate escalation to human healthcare professionals when necessary. If a query falls outside the informational scope of the kiosk, involves personalized medical advice, or exceeds system confidence thresholds, the response clearly recommends consultation with qualified healthcare staff. In PHC deployments, the kiosk may provide directional guidance or notify personnel to assist the user.

The overall interaction flow is designed to maintain a balance between accessibility, safety, and clinical boundaries, reinforcing the kiosk's role as a supportive health information tool rather than a substitute for professional medical care.

PILOT STUDY AND PROOF OF CONCEPT EVALUATION

To evaluate the feasibility and effectiveness of the proposed kiosk system, a pilot study can be conducted across selected Primary Health Centers (PHCs). The proof-of-concept (PoC) deployment can involve a representative sample of users from rural and semi-urban populations.

A structured sampling methodology can be adopted, with an initial cohort of approximately 50–100 participants across 2–3 PHCs. The evaluation can include diverse demographic groups, including elderly users, caregivers, and individuals with varying literacy levels.

The study design can incorporate both quantitative and qualitative evaluation metrics. Quantitative measures may include response accuracy, system usability scores, interaction completion rates, and language comprehension effectiveness. Qualitative feedback can be collected through user surveys and interviews to assess perceived usefulness, trust, and ease of interaction.

Pre- and post-interaction assessments can be conducted to measure improvements in health awareness levels. The results of such a pilot study can provide empirical evidence for system effectiveness and inform further optimization prior to large-scale deployment.

ETHICAL CONSIDERATIONS

The deployment of large language model (LLM)–driven healthcare kiosks in clinical environments raises critical ethical considerations spanning regulatory compliance, data governance, security, and institutional accountability. Compliance should be embedded at the architectural level to safeguard patient rights, rather than treated as a procedural formality. Core ethical requirements include robust data protection, meaningful informed consent, transparency in system functioning, and clearly defined accountability structures [29]

Human oversight remains essential. Despite technical safeguards such as encryption and secure storage, clinical supervision is required to monitor outputs, validate recommendations, and intervene in cases of inaccurate or potentially harmful responses. Automated systems must operate under structured human governance.

Privacy and data governance are central concerns. Patient data must be processed lawfully, with strict purpose limitation and data minimization. Unauthorized secondary use of health data—such as for model retraining—without explicit informed consent compromises patient autonomy. Clear communication of data rights and retention policies is therefore necessary.

Data provenance and intellectual property integrity also require attention. The use of improperly sourced or copyrighted training materials can create legal and ethical risks while undermining trust in system outputs. Transparent documentation of data sources and lawful processing practices is essential.

Finally, safety, reliability, bias mitigation, and transparency must guide system design. LLMs must be grounded in validated medical knowledge, continuously audited for fairness, and subject to institutional accountability in cases of adverse outcomes. Ethical integration demands a multilayered governance framework combining technical safeguards with sustained human oversight.

LEGAL AND REGULATORY PERSPECTIVE (INDIA)

The deployment of LLM-enabled healthcare information kiosks in Primary Health Centres (PHCs) in India must be assessed within the existing constitutional and statutory framework. Although India does not yet have a dedicated AI regulation statute, multiple intersecting legal regimes collectively govern data protection, healthcare delivery, intellectual property, consumer rights, and liability.

The Digital Personal Data Protection Act, 2023 (DPDP Act) forms the central pillar of compliance. Health data constitutes personal data—and in many cases, sensitive in nature—requiring lawful processing based on informed consent, purpose limitation, and implementation of reasonable security safeguards. For LLM-operated kiosks in PHCs, this necessitates clear patient notice regarding data collection, storage, processing, and potential secondary use (e.g., model updates). Non-compliance may attract regulatory penalties and civil liability.

The Information Technology Act, 2000, particularly Section 43A and related rules on sensitive personal data, reinforces obligations concerning reasonable security practices and cybersecurity safeguards. Given the digital architecture of LLM systems, robust protection against unauthorized access, breaches, and misuse of patient information is legally essential.

Technology-enabled healthcare must also align with the Telemedicine Practice Guidelines, 2020, issued by the Ministry of Health and Family Welfare. These guidelines prescribe standards for consent, confidentiality, and professional accountability in digital health interactions. Although LLM kiosks are not substitutes for registered medical practitioners, their integration into PHCs must not dilute these safeguards. Additionally, compliance with the Ayushman Bharat Digital Mission (ABDM) framework and Electronic Health Record (EHR) Standards is necessary to ensure interoperability, standardized data handling, and secure digital health infrastructure.

If LLM kiosks perform diagnostic or clinical decision-support functions, they may fall within the regulatory ambit of the Drugs and Cosmetics Act, 1940 and the Medical Devices Rules, 2017, particularly where software qualifies as a medical device. This would trigger requirements relating to regulatory approval, quality assurance, and post-market oversight.

Liability considerations further implicate the Consumer Protection Act, 2019, under which patients availing digital health services qualify as consumers. Erroneous or misleading AI-generated outputs may give rise to claims for deficiency of service, necessitating clear contractual allocation of responsibility among developers, deployers, and healthcare institutions.

Intellectual property compliance is governed by the Copyright Act, 1957 and the Patents Act, 1970. The use of copyrighted medical literature or proprietary datasets for model training without authorization may result in infringement. Questions also arise regarding ownership of AI-generated outputs within publicly funded healthcare systems.

At the constitutional level, the right to informational privacy recognized by the Supreme Court in Justice K.S. Puttaswamy (Retd.) & Anr v. Union of India & Ors (2017) affirms that state interference with personal data must satisfy the tests of legality, necessity, and proportionality. Accordingly, any AI deployment in public healthcare must conform to these constitutional standards [30]

In sum, while India lacks a consolidated AI statute, the combined operation of data protection law, digital health regulations, consumer protection norms, intellectual property statutes, and constitutional principles establishes a structured—though evolving—legal framework governing the use of LLM-powered kiosks in PHCs.

LIABILITY AND ACCOUNTABILITY FRAMEWORK

The integration of AI systems within healthcare environments introduces complex questions of accountability and legal responsibility. While the proposed kiosk is designed as a non-diagnostic informational tool, its potential influence on user behavior necessitates clear delineation of responsibilities.

Responsibility for system outputs can be distributed across multiple stakeholders, including developers, deploying healthcare institutions, and regulatory authorities. Developers are responsible for ensuring model safety, validation, and adherence to defined operational constraints. Healthcare institutions are responsible for appropriate deployment, supervision, and user guidance. Regulatory bodies play a critical role in defining compliance standards and oversight mechanisms.

In scenarios where AI-generated information may contribute to delayed care or adverse outcomes, the absence of clearly defined liability frameworks presents a significant challenge. Therefore, it is essential to establish regulatory guidelines that define accountability, compensation mechanisms, and legal recourse pathways.

The system must explicitly communicate its limitations to users, reinforcing that it does not replace professional medical advice. Clear disclaimers, escalation protocols, and human-in-the-loop mechanisms are necessary to mitigate risks and ensure ethical deployment.

SECURITY AND PRIVACY ARCHITECTURE

The security and privacy architecture of the proposed LLM-powered healthcare kiosk is designed to ensure responsible handling of user interactions within Primary Health Centers (PHCs). Given the sensitivity of health-related information, the system follows a privacy-by-design approach that prioritizes minimal data exposure, secure processing, and controlled access. The architecture integrates technical safeguards at the data collection, processing, storage, and monitoring levels to maintain confidentiality, integrity, and accountability [31]

Data Collection Minimization

The system adheres to the principle of data minimization by collecting only the information necessary to process user queries. The kiosk does not require user registration, personal identifiers, or medical history for standard informational interactions. Queries are processed in real time without mandating storage of personally identifiable information (PII). Any optional feedback mechanisms are designed to operate anonymously unless explicit consent is obtained. By limiting the scope of collected data, the system reduces privacy risks and aligns with ethical AI deployment standards.

On-Device vs. Cloud Processing

To enhance data security and ensure functionality in low-connectivity environments, the architecture supports both on-device (offline) and cloud-based (online) processing modes. In offline mode, speech recognition, language detection, and LLM inference are executed locally on the kiosk's compute unit. This minimizes external data transmission and reduces exposure to network-based vulnerabilities.

In online mode, when cloud APIs are utilized for enhanced performance, only the processed textual query (not raw audio unless required) is transmitted over secure communication channels. Data transmission is encrypted using industry-standard protocols such as TLS. The hybrid design allows PHCs to balance performance optimization with privacy preservation depending on connectivity conditions.

Anonymization of User Interactions

All interactions are treated as anonymous by default. The system avoids linking queries to identifiable individuals. If session logs are maintained for system improvement or analytics, they are stripped of any direct or indirect identifiers before storage. Metadata such as timestamps or language preferences may be retained in aggregated form for operational optimization, but without user traceability. This anonymization strategy prevents profiling and unauthorized tracking of individuals.

Secure Storage and Transmission

Data storage, when required for administrative or analytical purposes, is secured through encryption at rest using standardized cryptographic mechanisms. Access to stored logs and configuration data is restricted to authorized administrative personnel through role-based access control (RBAC).

For cloud-based interactions, all communications between the kiosk and remote servers are encrypted during transmission. Firewalls and secure network configurations are implemented to prevent unauthorized access. Physical security measures, including locked hardware compartments and restricted service panels, further protect on-device components from tampering.

Audit and Logging Mechanisms

The system incorporates structured audit and logging mechanisms to maintain operational transparency and accountability. Logs may include system performance metrics, error reports, model confidence thresholds, and anonymized query categories. These logs enable administrators to monitor system health, detect anomalies, and evaluate compliance with safety constraints.

Audit trails are designed to support oversight without compromising user privacy. Regular review of logs assists in identifying misuse, attempted system exploitation, or deviations from predefined response boundaries. Logging policies are governed by retention limits to prevent unnecessary long-term data storage.

Overall, the security and privacy architecture ensures that the kiosk operates within a controlled, ethical, and privacy-preserving framework. By combining data minimization, hybrid processing models, anonymization techniques, secure communication protocols, and accountable logging practices, the system aligns with responsible AI deployment standards in public healthcare environments [32-34]

RISK STRATIFICATION AND STANDARDS COMPLIANCE

The deployment of AI-assisted healthcare systems necessitates a structured risk management framework aligned with

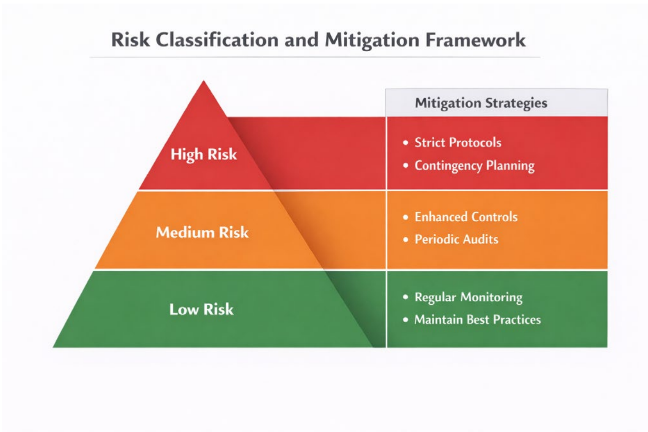

Fig.5: Risk Classification and Mitigation Framework

international standards. The proposed kiosk, although currently limited to informational use, can be evaluated against medical-grade AI compliance frameworks to ensure safety and future scalability.

A risk stratification approach can be adopted to classify potential system risks into categories such as low-risk (informational inaccuracies), medium-risk (misinterpretation of guidance), and high-risk (delayed medical consultation due to over-reliance). Mitigation strategies include response guardrails, uncertainty handling, and mandatory referral mechanisms to healthcare professionals.

The system can be aligned with relevant standards, including ISO 13485:2016 for quality management systems, ISO 14971:2019 for risk management in medical devices, and ISO/DTS 24971-2.2 for AI-specific risk considerations. Additionally, ISO/IEC 42001:2023 can guide AI governance practices, while IEC 62304 can be referenced for software lifecycle management.

Fig. 5 illustrates the risk classification as well as respective risk mitigation frameworks. Although the current implementation is positioned as a non-medical, informational tool, adherence to these standards can facilitate future certification if the system evolves toward clinical decision-support functionalities.

LLM DATA GOVERNANCE AND MULTILINGUAL CLINICAL FRAMEWORK

The reliability and safety of the proposed system depend significantly on robust data governance mechanisms. The LLM module can be supported by a structured data governance framework that ensures the use of verified, policy-aligned, and culturally appropriate healthcare information.

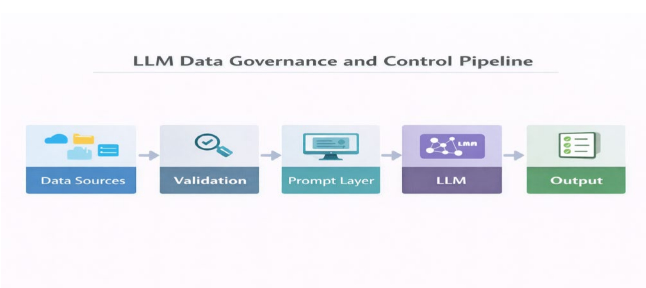

Fig.6: LLM Data Governance and Control Pipeline

The data pipeline can be designed to incorporate curated sources such as government-issued health guidelines, WHO recommendations, and standardized clinical awareness materials. A multi-layer validation process involving domain experts, public health authorities, and periodic audits can be integrated to ensure accuracy and consistency of responses. Fig. 6 illustrates the LLM data governance and control pipeline of the proposed model. In the context of India's linguistic diversity, a multilingual framework can be incorporated to handle multiple languages, dialects, and code-mixed inputs. Rather than relying solely on direct translation, the system can adopt context-aware localization strategies that account for cultural nuances, regional terminology, and varying health literacy levels.

Additionally, governance policies can define restrictions on model behavior, including domain boundaries, prohibited outputs, and escalation protocols. This ensures that the system remains aligned with its non-diagnostic objective while maintaining user trust and safety.

CHALLENGES AND LIMITATIONS

The deployment of LLM-powered healthcare kiosks in Primary Health Centers (PHCs) presents several technical, infrastructural, and regulatory challenges that must be carefully addressed.

Large Language Models are inherently probabilistic and may generate inaccurate or overconfident responses, particularly in complex medical scenarios. Despite guardrails and prompt constraints, the risk of hallucination or misinterpretation cannot be entirely eliminated. Ensuring consistent alignment with public health guidelines remains an ongoing technical challenge.

Many rural and semi-urban PHCs face limitations such as unstable electricity supply, intermittent internet connectivity, limited physical space, and a lack of technical maintenance personnel. These constraints can affect system uptime, performance, and long-term sustainability.

India's linguistic diversity extends beyond officially recognized languages to numerous regional dialects. Accurately detecting and responding in colloquial or mixed-language inputs (code-switching) remains challenging. Variations in pronunciation, accents, and local terminology may impact speech recognition accuracy.

The use of AI in healthcare information delivery raises concerns regarding accountability, liability, and regulatory compliance. Clear boundaries must be maintained to prevent the system from being perceived as a diagnostic tool. Additionally, evolving AI governance policies may impose new compliance requirements.

Scaling the system across multiple PHCs requires standardized hardware deployment, centralized monitoring, regular content updates, and model optimization for resource-constrained devices. Ensuring consistent performance across diverse geographic and infrastructural conditions presents logistical and operational complexities.

These challenges highlight the need for cautious implementation, continuous monitoring, and policy-aligned development to ensure safe and sustainable deployment.

FUTURE SCOPE

The proposed system offers significant opportunities for expansion and enhancement within the public healthcare ecosystem. Future versions may connect with PHC medicine inventory databases to provide real-time information on drug availability and stock status, improving transparency and patient guidance. A supervised interaction mode can be introduced where healthcare professionals monitor or review AI-generated responses, enabling hybrid AI-human collaboration and increasing reliability in sensitive cases [35].

Model updates can be periodically performed using curated healthcare datasets under strict regulatory supervision to ensure alignment with evolving medical guidelines and public health policies. The system can be extended to include specialized modules for antenatal care, immunization tracking awareness, nutrition guidance, and early childhood health education. Future integration with national digital health infrastructures can enable interoperability with government health platforms, supporting unified healthcare data ecosystems while maintaining privacy safeguards. These advancements can strengthen the role of AI-assisted kiosks as scalable, policy-aligned tools within India's primary healthcare framework.

CONCLUSION

This work presents the design and implementation framework of a multilingual, LLM-powered healthcare information kiosk tailored for Primary Health Centers (PHCs). The proposed system integrates voice-based interaction, controlled response generation, offline capability, and safety guardrails to deliver non-diagnostic, awareness-oriented health information through low-cost kiosks and web platforms.

The solution addresses critical gaps in healthcare communication within Indian PHCs, particularly language barriers, limited doctor–patient interaction time, and low health literacy levels. By enabling accessible and conversational health information in regional languages, the system has practical relevance for rural and semi-urban healthcare environments.

The deployment of such systems underscores the importance of responsible AI practices in healthcare. Strict domain boundaries, privacy-preserving design, and clear referral mechanisms are essential to ensure that AI functions as a supportive informational tool rather than a substitute for clinical expertise. With careful implementation and regulatory alignment, LLM-powered kiosks can contribute meaningfully to equitable and inclusive primary healthcare delivery.

REFERENCES

1. Asgari-Jirhandeh N, Zapata T, Jhalani M. Strengthening primary health care as a means to achieve universal health coverage: experience from India. Journal of Health Management. 2021 Mar;23(1):20-30.
2. Potnuru B. Aggregate availability of doctors in India: 2014–2030. Indian journal of public health. 2017 Jul 1;61(3):182-7.
3. Douglass K, Narayan L, Allen R, Pandya J, Talib Z. Language diversity and challenges to communication in Indian emergency departments. International Journal of Emergency Medicine. 2021 Dec;14(1):57.
4. Manoj R, Nandhini G. A comprehensive investigation on leveraging generative AI and large language models in the healthcare domain. In2024 IEEE 12th Region 10 Humanitarian Technology Conference (R10-HTC) 2024 Oct 1 (pp. 1-6). IEEE.
5. Zamanifar A, Faezipour M, Gharehchopogh FS, Farhadi A, editors. Applications of Large Language Models (LLM) in Healthcare Systems: Opportunities, Challenges, and Ethical Considerations.
6. Manoj R, Krishna N, TS MS. A Comprehensive Study on the Integration of Robotic Technology in Medical Applications considering Legal Frameworks & Ethical Concerns. International Journal of Health Technology and Innovation. 2024 Aug 20;3(02):29-37.
7. Edited on behalf of the National Institute of Biomedical Imaging and Bioengineering/National Heart, Lung, and Blood Institute/National Science Foundation Workshop Faculty, Price CP, Kricka LJ. Improving healthcare accessibility through point-of-care technologies. Clinical Chemistry. 2007 Sep 1;53(9):1665-75.
8. Sinha N, Rahul E. India's digital health transformation: advancements and challenges in telemedicine. Medicon Medical Sciences. 2025;8:31-40.
9. Singh M, Agarwal A. Telemedicon 2022 Kerala 10th November to 12th November 2022. Journal of Medical Evidence. 2022 Sep 1;3(3):294-6.
10. Omarov B, Narynov S, Zhumanov Z. Artificial intelligence-enabled chatbots in mental health: A systematic review. Computers, Materials & Continua. 2023 Mar 1;74(3).
11. Casu M, Triscari S, Battiato S, Guarnera L, Caponnetto P. AI chatbots for mental health: A scoping review of effectiveness, feasibility, and applications. Applied Sciences. 2024 Jul 5;14(13):5889.
12. Tejasri K, Shri ML. Bridging Languages in Healthcare: A Comprehensive Review of Multilingual and Code-Switched Chatbot Interactions. IEEE Access. 2026 Feb 12.
13. Munjal G, Agarwal P, Goyal L, Samiran N. Multilingual virtual healthcare assistant. Health Care Science. 2025 Aug;4(4):281-8.
14. Mirzaei T, Amini L, Esmaeilzadeh P. Clinician voices on ethics of LLM integration in healthcare: a thematic analysis of ethical concerns and implications. BMC Medical Informatics and Decision Making. 2024 Sep 9;24(1):250.
15. Kumar S, Balloli V, Ranjit M, Ahuja K, Ganu T, Sitaram S, Bali K, Nambi A. Bridging the Gap: Dynamic Learning Strategies for Improving Multilingual Performance in LLMs. arXiv preprint arXiv:2405.18359. 2024 May 28.
16. Trivedi A, Pant N, Shah P, Sonik S, Agrawal S. Speech to text and text to speech recognition systems-Areview. IOSR J. Comput. Eng. 2018 Mar;20(2):36-43.
17. Summerfield M. Rapid GUI programming with Python and Qt: the definitive guide to PyQt programming. Pearson Education;

2007 Oct 18.

18. Yoon G. Comparison Between Python and Web User Interfaces.
19. Hakkarainen S, Engelbrecht K. Automatic Idiomatic Expression Detection. Comparison Between GPT-4 and Gemini Pro Prompt Engineering & LSTM-RNN Construction.
20. Patel H, Shah KA, Mondal S. Do large language models generate similar codes from mutated prompts? A case study of Gemini Pro. InCompanion Proceedings of the 32nd ACM International Conference on the Foundations of Software Engineering 2024 Jul 10 (pp. 671-672).
21. Mora-Delgado J, Ramos-Ruperto L, Pardilla MJ, Sicilia MÁ, Rodríguez-González A, Sempere JM, Puchades R. Generative AI: foundational models. Natural Language Processing (NLP) and LARGE Language Models (LLM). Revista Clínica Española (English Edition). 2025 Dec 20:502413.
22. Li R, Wang X, Yu H. LlamaCare: an instruction fine-tuned large language model for clinical NLP. InProceedings of the 2024 Joint international conference on computational linguistics, language resources and evaluation (LREC-COLING 2024) 2024 May (pp. 10632-10641).
23. Kusuma S, Krishnan SG, Vikram R, UD S, Reddy V. Text-Based Audio Correction Using Deep Learning Framework. In2024 International Conference on Emerging Research in Computational Science (ICERCS) 2024 Dec 12 (pp. 1-6). IEEE.
24. Paul RJ, Ezz I, Kuljis J. Healthcare information systems: a patient-user perspective. Health Systems. 2012 Dec 1;1(2):85-95.
25. Pandey A, Sharma I. Family ties and health care decision-making in late life: examining financial dependency in health for India. Innovation in Aging. 2024 Dec 31;8(Suppl 1):959.
26. Panteli D, Adib K, Buttigieg S, Goiana-da-Silva F, Ladewig K, Azzopardi-Muscat N, Figueras J, Novillo-Ortiz D, McKee M. Artificial intelligence in public health: promises, challenges, and an agenda for policy makers and public health institutions. The Lancet Public Health. 2025 May 1;10(5):e428-32.
27. Datta K. AI-driven public administration: Opportunities, challenges, and ethical considerations. The Social Science Review. 2024 Nov;2(6):134-9.
28. Novikau A. Online vs. offline LLM inference: Unlocking the best of both worlds in Mobile applications. International Journal of Science and Engineering Applications. 2024;13(12):5-8.
29. Fareed M, Fatima M, Uddin J, Ahmed A, Sattar MA. A systematic review of ethical considerations of large language models in healthcare and medicine. Frontiers in digital health. 2025 Sep 11;7:1653631.
30. Ong JC, Chang SY, William W, Butte AJ, Shah NH, Chew LS, Liu N, Doshi-Velez F, Lu W, Savulescu J, Ting DS. Ethical and regulatory challenges of large language models in medicine. The Lancet Digital Health. 2024 Jun 1;6(6):e428-32.
31. Falahati S, Alizadeh M, Ghazipour F, Safahi Z, Khaledian N, Salmanpour MR. An AI-powered public health automated kiosk system for personalized care: an experimental pilot study. arXiv preprint arXiv:2504.13880. 2025 Apr 2.
32. Shaikh AJ, Takbhate SB, Gawai S, Limbore AM, Patil VA. Revolutionizing healthcare: AI-powered solutions, technological integration for accessible and effective medical service. In2024 International Conference on Advances in Computing Research on Science Engineering and Technology (ACROSET) 2024 Sep 27 (pp. 1-7). IEEE.
33. Montagna S, Ferretti S, Klopfenstein LC, Ungolo M, Pengo MF, Aguzzi G, Magnini M. Privacy-preserving LLM-based chatbots for hypertensive patient self-management. Smart Health. 2025 Jun 1;36:100552.
34. Tahera MA, Sidhu KS, Dass S, Saha S. SoK: Privacy-aware LLM in Healthcare: Threat Model, Privacy Techniques, Challenges and Recommendations. arXiv preprint arXiv:2601.10004. 2026 Jan 15.
35. Fragiadakis G, Diou C, Kousiouris G, Nikolaidou M. Evaluating human-ai collaboration: A review and methodological framework. arXiv preprint arXiv:2407.19098. 2024 Jul 9.