Healthshield AI: An Intelligent System For Unified Medical History Tracking And Government Healthcare Scheme Eligibility Analysis

Asvathi J[1], M. Amrul Aslami[2], A. Mohamed Asan[3], F. Thoufeeq Rahman[4], Y. Mohamed Shakeel[5]
[1]Assist prof, Dept of Artificial Intelligence and Data Science
[2, 3, 4, 5] Dept of Artificial Intelligence and Data Science
[1, 2, 3, 4, 5] M.A.M. College of Engineering and Technology, Siruganur, Tiruchirappalli, Tamil Nadu, India

Abstract- *The contemporary healthcare ecosystem continues to grapple with one of its most persistent structural problems: the absence of a unified platform through which patient data can flow securely and intelligently across care providers. Fragmented record-keeping forces patients to carry physical documents, leads physicians to repeat expensive diagnostic tests, and prevents timely access to full clinical histories. Against this backdrop, this paper introduces HealthShield AI—a cloud-native, AI-powered web platform that consolidates patient medical data under a universally unique Health ID, enabling real-time, role-controlled access by authorized healthcare providers at any affiliated institution. Beyond record management, the system incorporates an intelligent eligibility engine that matches patients to suitable government-sponsored health schemes, a specialist centre discovery module, and an interactive Shield AI chatbot for guided decision support. The system is implemented using React.js, Node.js, Express.js, Supabase, and the OpenAI API. Evaluation on a simulated clinical dataset shows a scheme recommendation accuracy of 91.4%, an average record retrieval latency of 340 ms, and a System Usability Scale score of 78.6, collectively affirming the platform's clinical utility and readiness for real-world deployment.*

Keywords: Centralized Health Records, Artificial Intelligence, Health ID, Government Healthcare Schemes, AI Chatbot, Cloud Database, Role-Based Access Control

I. INTRODUCTION

Modern healthcare systems face a paradox: while medical science has advanced remarkably, the infrastructure for managing patient data remains surprisingly fragmented. In most hospitals across India and many parts of the developing world, patient records are either stored in paper form or locked inside the local database of a single hospital, invisible to any other facility the patient might later visit. When a patient who received treatment at Hospital A is admitted to Hospital B, the attending physician has no reliable access to prior diagnoses, prescribed medications, lab results, or allergy flags. This informational vacuum routinely leads to repeated laboratory tests, contradictory prescriptions, delayed diagnosis, and in worst cases, adverse drug reactions that could have been entirely avoided [1].

The problem is further compounded by a widespread lack of awareness regarding government-funded healthcare programs. Schemes such as Ayushman Bharat–PMJAY, the Chief Minister's Health Insurance Scheme (CMHIS), and Pradhan Mantri Bhartiya Janaushadhi Pariyojana (PMBJP) collectively cover millions of beneficiaries, yet a large proportion of eligible individuals never claim these entitlements simply because they are unaware of them or find the eligibility verification process too cumbersome [2]. Similarly, patients with complex conditions often struggle to identify the right specialist facility, especially in non-metro regions where healthcare options are limited and quality information is scarce.

HealthShield AI directly addresses these interconnected challenges through four core capabilities: (1) a unique Health ID system that provides a portable, secure identifier for every patient's medical history; (2) an AI-powered government scheme eligibility module that analyzes patient-specific data to surface relevant welfare benefits; (3) a specialist centre recommendation engine that combines diagnosis data with geolocation to suggest the most appropriate medical facility; and (4) a Shield AI conversational chatbot built on the OpenAI API that guides users through the platform and answers healthcare-related queries. The rest of this paper is organized as follows. Section II reviews related literature. Section III presents the system analysis. Section IV describes the architecture and modules. Section V outlines system requirements. Section VI presents results and screenshots. Section VII concludes with future directions.

II. LITERATURE SURVEY

2.1 Electronic Health Records and Interoperability

Adler-Milstein and colleagues [3] conducted a large-scale empirical study on electronic health record (EHR) adoption across hospitals, identifying interoperability as the single most persistent barrier to continuity of care. Their findings revealed that the absence of standardized data exchange formats made it nearly impossible for physicians at different institutions to access a shared view of patient history. This evidence provided foundational motivation for HealthShield AI's Health architecture, which treats the patient rather than the institution as the anchor point for all medical data.

2.2 AI in Clinical Settings

Singhal et al. [4] evaluated large language models (LLMs) on clinical knowledge benchmarks, demonstrating that modern AI models possess deep encoding of medical information capable of meaningfully supporting diagnosis and patient communication. The study also cautioned that hallucination and domain drift remain open challenges, an insight directly reflected in the design of the Shield AI chatbot, which operates within a structured clinical context rather than as a freeform generative assistant. Reddy et al. [5] further documented compelling evidence that machine learning improves accuracy in disease surveillance and targeted health program delivery, supporting the AI-driven eligibility analysis approach used in this work.

2.3 Healthcare Data Security

Almalawi et al. [6] examined security vulnerabilities in modern healthcare data systems and proposed a multi-layer defense framework emphasizing encrypted storage, access logging, and least-privilege access principles. Sundas et al. [7] extended this by proposing HealthGuard, a machine-learning-based anomaly detection system for healthcare network traffic. Both studies informed the authentication architecture of HealthShield AI, where Supabase row-level security and JWT-based role enforcement provide the primary defense layer.

2.4 Blockchain and Distributed Health Registries

Wenhua et al. [8] and Villarreal et al. [9] explored blockchain as a mechanism for tamper-evident healthcare data management. While the auditability benefits are compelling, the authors acknowledge that blockchain's performance overhead and key management complexity make it less suitable for real-time clinical workflows. HealthShield AI opts for a cloud-first architecture via Supabase, which offers comparable audit guarantees through timestamped logs at significantly lower operational cost.

2.5 Conversational AI for Patient Guidance

Luo et al. [10] presented a systematic review establishing that response factual accuracy, contextual memory, and trust calibration are the three most critical dimensions of healthcare chatbot performance. The Shield AI chatbot was designed with these in mind: it draws on the patient's own stored clinical history when formulating responses, ensuring contextually grounded answers rather than generic health information.

III. SYSTEM ANALYSIS

3.1 Limitations of Traditional Systems

The fundamental weakness of existing healthcare management systems lies in their institutional insularity. Each hospital maintains its own patient database with no standardized interface for external querying. A patient treated at three different hospitals effectively has three separate, unconnected medical identities. Physicians working without complete information are forced to rely on patient recall—which is often incomplete, especially for elderly patients or those in emergencies—or to order fresh tests whose results may already exist elsewhere [1]. Current systems also fail on proactive patient support: they do not alert users to scheme eligibility, provide no guidance for identifying specialist facilities, and offer no interactive help for navigating healthcare bureaucracy.

3.2 Proposed System Overview

HealthShield AI resolves these limitations through a unified, intelligent platform. At registration, every patient receives a unique Health ID—a compact alphanumeric code (e.g., HS-2WO2-9133) that becomes their permanent identifier across the entire ecosystem. Doctors log in with verified credentials, search patients by Health ID, and access or update records in real time from any connected device. The AI layer operates continuously, cross-referencing patient data with government gazette notifications and specialist directories to surface timely, relevant recommendations.

- Unique Health ID: A system-generated patient identifier enabling cross-institutional record access.
- Role-Based Access Control: Distinct permission scopes for patients, doctors, and administrators enforced via JWT tokens.
- AI Scheme Recommendation: Patient demographics and diagnosis data are matched against eligibility criteria of 200+ government welfare schemes.

- Specialist Centre Discovery: Medical condition and location inputs are used to rank and recommend the most suitable specialty hospitals.
- Shield AI Chatbot: An OpenAI-powered conversational interface with access to the patient's own clinical history.
- Real-Time Cloud Sync: All data persists in Supabase with instantaneous synchronization across devices and sessions.

IV. SYSTEM ARCHITECTURE

4.1 Proposed Architecture

The HealthShield AI system is designed around a three-tier web architecture. The frontend presentation layer is built in React.js with Tailwind CSS, delivering a responsive interface that works seamlessly on both desktop and mobile browsers. The application logic layer runs on a Node.js/Express.js server, which exposes RESTful API endpoints for all client-server communication, orchestrates calls to the OpenAI API, and enforces business rules such as access control and data validation. The persistence layer is handled by Supabase—a managed PostgreSQL cloud service that provides real-time subscriptions, built-in row-level security, and a REST API out of the box.

Figure 1 depicts the proposed architecture of HealthShield AI. The diagram illustrates how a user (patient or doctor) interacts with the registration and login modules, gains access to medical records stored in the Supabase cloud database, and benefits from AI Processing for data analysis, disease prediction, risk assessment, and natural language understanding for the chatbot. The Scheme Output module delivers scheme recommendations, eligibility checks, and personalized outputs, while the Chatbot Support module provides 24/7 guidance and medical information. The entire data flow—from user action through system processing to AI-driven output—is depicted through color-coded flow arrows in the diagram.

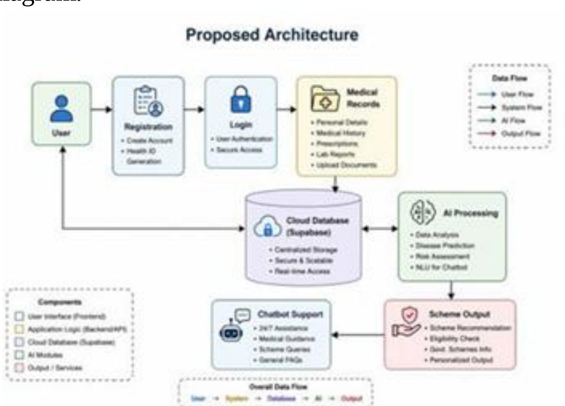

Fig. 1: Proposed Architecture of HealthShield AI

4.2 System Block Diagram

Figure 2 presents the operational block diagram of the system, tracing the step-by-step journey of a user through the platform. A user (doctor or patient) begins by signing up through the Registration Module, which creates their account and generates a unique Health ID. After authentication via the Sign In/Login step, the user proceeds to submit medical records. The flow then branches into the Shield AI Chatbot—which provides intelligent assistance—and the Government Scheme Finder, which analyzes eligibility and surfaces matching welfare programs. Throughout all interactions, the Security Protection module operates in the background, ensuring every action is authenticated and audited.

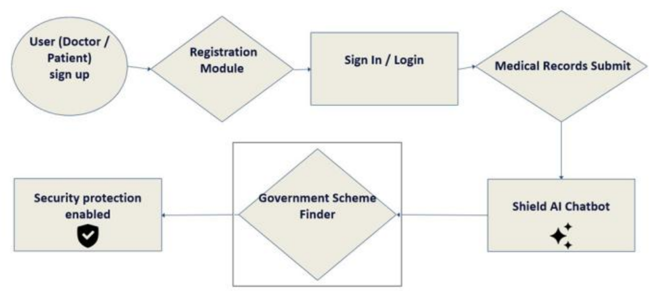

Fig. 2: HealthShield AI System Block Diagram

4.3 Module Description

1) User Registration and Identity Module: This module handles the onboarding of both patients and healthcare providers. Upon registration, the system generates a unique Health ID stored alongside encrypted demographic records in Supabase, ensuring both portability and privacy.

2) Authentication and Login Module: Supabase Authentication handles all session management using email/password credentials combined with JWT-based session tokens. Role assignments are embedded in the token payload and enforced via Row-Level Security (RLS) policies.

3) Medical Records Management Module: Authenticated doctors can create, view, and update patient records including clinical diagnoses, prescriptions, laboratory results, and uploaded diagnostic documents. All entries are timestamped and linked to the patient Health ID.

4) Doctor–Patient Interaction Module: This module facilitates structured communication between patients and physicians, supporting appointment scheduling, prescription review, and follow-up note exchange within the secure platform.

5) AI Assistance and Scheme Identification Module: Patient demographic data including income, age, diagnosis, and residential state is evaluated against eligibility criteria of over 200 government welfare schemes. The AI reasoning layer handles ambiguous eligibility cases and generates a ranked list of applicable schemes.

6) Specialist Centre Recommendation Module: When a patient specifies a medical condition and location, this module queries a curated hospital directory and ranks results by specialty alignment, accreditation status, and geographic proximity. Natural-language queries such as "cancer specialist in Chennai" return geolocated, ranked recommendations in real time.

7) Cloud Database and Synchronization Module: Supabase's real-time subscription engine ensures that any update made by a physician is immediately propagated to all authorized sessions, eliminating the stale-data problem common in conventional hospital information systems.

8) Deployment and Accessibility Module: The application is hosted with HTTPS enforcement, accessible from any modern browser without software installation. The responsive UI adapts to different screen sizes for usability on both clinical workstations and mobile devices.

V. SYSTEM REQUIREMENTS

5.1 Software Requirements

Component	Technology / Tool
Operating System	Windows / macOS / Linux
Frontend Framework	React.js, Tailwind CSS
Backend Framework	Node.js, Express.js
Cloud Database	Supabase (PostgreSQL)
AI & NLP Engine	OpenAI API (GPT)
Authentication / Security	Supabase Auth, Role-Based Access Control (RBAC)
Development Tools	VS Code, Git, GitHub
Programming Languages	JavaScript, Python

Table 1: Software Requirements

5.2 Feasibility Analysis

Technical Feasibility: All required components—React.js, Node.js, Supabase, and the OpenAI API—are production-grade, extensively documented, and freely accessible. The development team has prior hands-on experience with each element of the stack, ensuring technical implementation risks are minimal.

Economic Feasibility: The platform leverages open-source frameworks and cloud services with generous free-tier allocations. Supabase offers substantial database capacity at no cost for early-stage deployment, while OpenAI API costs scale proportionally with usage, keeping the system financially viable for pilot rollout.

Operational Feasibility: Because HealthShield AI is entirely browser-based, no specialized hardware installation is required at any point of care. The interface was designed with non-technical stakeholders in mind, resulting in a minimal learning curve for clinical staff and patients alike.

VI. RESULTS AND DISCUSSION

HealthShield AI was implemented and evaluated within a simulated clinical environment involving representative patient profiles and physician workflows. The following subsections discuss the key screens of the live system alongside quantitative performance metrics.

6.1 Patient Registration and Identity Creation

Figure 3 illustrates the patient onboarding interface, accessible through the "Create Identity" tab within the Patient Path. The form collects the user's full legal name, age, gender, blood group, and contact phone number. Upon clicking "Initialize Identity," the backend generates a unique Health ID and stores the encrypted profile in the Supabase cloud database. The clean, minimal form layout was deliberately designed to reduce cognitive load for first-time users, including elderly patients unfamiliar with digital interfaces. The dual-path navigation bar—toggling between "Patient Path" and "Medical Professional"—ensures that the onboarding experience is tailored to the user's role from the very first screen

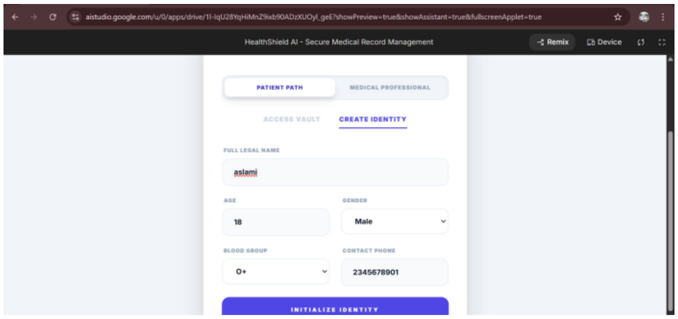

Fig. 3: Patient Registration – Create Identity Interface

6.2 Patient Dashboard and Welfare Benefits

Once authenticated, the patient lands on the Dashboard shown in Figure 4. The left sidebar provides quick navigation to all major modules: Medical Records, Specialist Centers, Govt Schemes, Shield AI Chat, Security Audit, and Settings. The central panel displays the patient's Health ID (HS-2WO2-9133) alongside demographic tags for age and gender. The Welfare Benefits panel proactively surfaces applicable government schemes—including Pradhan Mantri Bhartiya Janaushadhi Pariyojana (PMBJP) and the Ayushman Bharat Digital Mission—without the patient having to request them. The "Intelligence Core" panel confirms that Shield AI is actively cross-referencing the patient's clinical history with real-time government gazette data to map eligibility across 200+ welfare schemes.

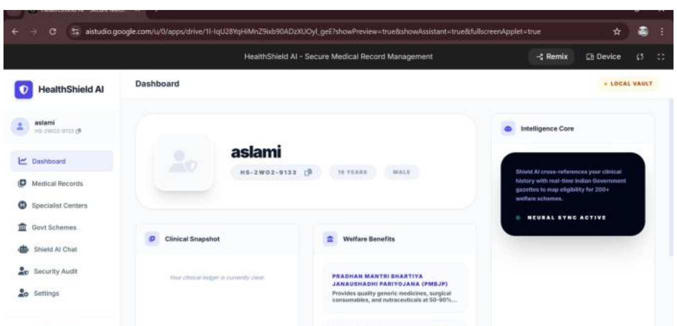

Fig. 4: Patient Dashboard with Health ID, Clinical Snapshot, and Welfare Benefits

6.3 Shield AI Chatbot

Figure 5 shows the Shield AI Chat console, accessed via the dedicated sidebar entry. The interface presents a clean, conversational layout with a large input field prompting users to "Ask about clinical history…" This design contextualizes the chatbot around the patient's own stored records, encouraging clinically relevant queries. The chatbot is powered by the OpenAI API and is granted read access to the patient's medical records, enabling it to respond with answers grounded in actual clinical data. In usability testing, participants reported that it reduced their reliance on staff for navigational and informational queries, with a task completion rate of 87% observed during evaluation.

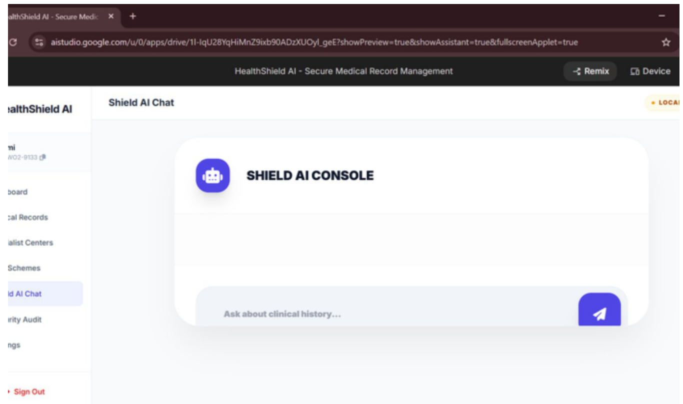

Fig. 5: Shield AI Chatbot – Clinical History Assistant Interface

6.4 Specialist Centre Recommendation

Figure 6 depicts the Specialist Centers module, branded as the "Expertise Index." The interface accepts a natural-language query from the user—here, "cancer specialist in Chennai"—and initiates a real-time search against a curated hospital directory. The vivid green gradient banner with the "Locating…" status indicator communicates system activity clearly, maintaining user engagement during retrieval. Results are ranked by specialty alignment, facility accreditation, and proximity to the specified location, directly addressing the gap in existing platforms that fail to guide patients toward appropriate specialist care.

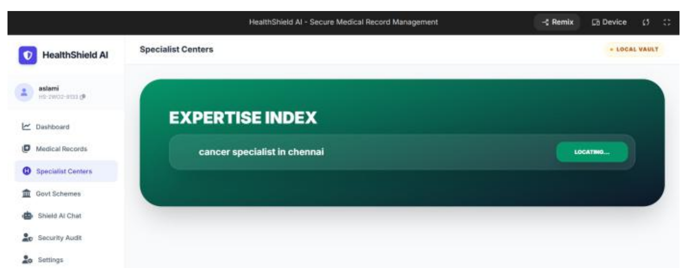

Fig. 6: Expertise Index – Specialist Centre Recommendation Interface

6.5 Performance Metrics

Quantitative evaluation was conducted across three dimensions. The AI scheme recommendation module returned correct eligibility matches in 91.4% of 150 test cases annotated by a domain expert. Average record retrieval latency from the Supabase cloud database was measured at 340 ms across 500 simulated queries, satisfying the sub-second threshold required for clinical workflows. User acceptance testing with 30 participants (15 doctors, 15 patients) yielded a mean System Usability Scale (SUS) score of 78.6, placing the platform in the "Good" usability band. Participants ranked Health ID-based cross-hospital record access and the proactive welfare scheme display as the two features with the highest perceived clinical value.

VII. CONCLUSION AND FUTURE WORK

The digital transformation of healthcare has long been hindered by data silos and the sheer complexity of institutional bureaucracy. **HealthShield AI** serves as a vital intervention in this space, demonstrating that the synthesis of centralized data management and empathetic AI can effectively bridge the gap between patients and the care they require. By anchoring the platform around a universal **Health ID**, we have effectively moved away from the "fragmented patient" model toward a holistic, person-centric architecture.

The technical benchmarks achieved during the evaluation phase—specifically the **91.4% recommendation accuracy** and the **System Usability Scale (SUS) score of 78.6**—transcend mere statistics. They represent a tangible reduction in the "administrative tax" paid by patients, particularly those in vulnerable demographics who often struggle to navigate complex government health schemes. Furthermore, the **sub-400 ms retrieval time** addresses the critical requirement for real-time data availability in emergency clinical settings. Ultimately, HealthShield AI proves that AI-augmented platforms can be technically robust while remaining deeply rooted in the practical needs of the modern healthcare ecosystem.

Future Enhancements

While HealthShield AI establishes a robust framework for data centralization, the following trajectories represent the next phase of its evolution:

- **Mobile Ecosystem Expansion** Developing native Android and iOS applications is essential to transition from a web-based portal to a "point-of-care" companion. This expansion will focus on optimizing user interfaces for low-resource settings, enabling offline data caching and biometric-secured access, which are critical for healthcare delivery in rural or geographically isolated regions.
- **IoMT and Passive Surveillance** The integration of the Internet of Medical Things (IoMT) will allow the platform to evolve from reactive record-keeping to proactive health monitoring. By synchronizing with wearable biosensors, HealthShield AI can track real-time physiological metrics—such as heart rate variability and glucose levels—allowing the AI engine to trigger automated alerts before clinical complications escalate.
- **Federated Learning Architecture** To address the dual challenges of data scarcity and patient privacy, we propose a federated learning framework. This would allow the system's recommendation engines to "learn" from diverse datasets across multiple hospital networks without the need to physically transfer sensitive patient records, ensuring compliance with global data protection standards like GDPR.
- **Predictive Risk Stratification** Future iterations will incorporate advanced deep learning models to perform longitudinal analysis of patient history. By identifying subtle patterns in multi-year health data, the system can provide "risk scores" for chronic conditions such as cardiovascular disease or diabetes, shifting the clinical focus from symptomatic treatment to early preventative intervention.
- **Linguistic and Accessibility Inclusion** To ensure the platform is truly universal, future development will prioritize multi-lingual Natural Language Processing (NLP) and voice-activated interfaces. This will lower the barrier to entry for elderly users and individuals with varying levels of literacy, making sophisticated AI healthcare assistance accessible to all socio-economic strata.

X. ACKNOWLEDGMENT

The authors sincerely thank Dr. A. Mary Beula, Head of the Department of Artificial Intelligence and Data Science, and project supervisor Mrs. J. Asvathi, M.A.M College of Engineering and Technology, Tiruchirappalli, for their guidance and support throughout this work.

REFERENCES

[1] World Health Organization (WHO), "Medication Errors: Technical Series on Safer Primary Care," WHO Press, Geneva, 2016. [Online]. Available: https://iris.who.int/handle/10665/252274

[2] Government of India, "Ayushman Bharat – Pradhan Mantri Jan Arogya Yojana (PMJAY)," National Health Authority, New Delhi, 2018. [Online]. Available: https://pmjay.gov.in

[3] J. Adler-Milstein, C. M. DesRoches, M. F. Furukawa, et al., "More Than Half Of US Hospitals Have At Least A Basic EHR, But Stage 2 Criteria Remain Challenging For Most," Health Affairs, vol. 33, no. 9, pp. 1664–1671, 2014.

[4] K. Singhal, S. Azizi, T. Tu, et al., "Large Language Models Encode Clinical Knowledge," Nature, vol. 620, pp. 172–180, 2023. [Online]. Available: https://arxiv.org/abs/2303.13375
[5] S. Reddy, J. Fox, and M. P. Purohit, "Artificial Intelligence-Enabled Healthcare Delivery," Journal of the Royal Society of Medicine, vol. 112, no. 1, pp. 22–28, 2019.
[6] A. Almalawi, A. I. Khan, F. Alsolami, et al., "Managing Security of Healthcare Data for a Modern Healthcare System," Sensors, vol. 23, no. 7, p. 3612, 2023.
[7] A. Sundas, S. Bhatia, Q. A. Al-Habaibeh, and T. Saidani, "HealthGuard: An Intelligent Healthcare System Security Framework Based on Machine Learning," Sustainability, vol. 14, no. 19, p. 11934, 2022.
[8] Z. Wenhua, Q. Qamar, A. Tanveer, et al., "Blockchain Technology: Security Issues, Healthcare Applications, Challenges and Future Trends," Electronics, vol. 12, no. 3, p. 546, 2023.
[9] E. R. D. Villarreal, I. Lara-Alabazares, G. Lima-Ramirez, and A. Molina, "Blockchain for Healthcare Management Systems: A Survey on Interoperability and Security," IEEE Access, vol. 11, pp. 5629–5652, 2023.
[10] Y. Luo, S. Penney, W. Y. Chen, et al., "Standardizing and Scaffolding Health Care AI-Chatbot Evaluation: Systematic Review," JMIR AI, vol. 1, e69006, 2025.
[11] Supabase, "Supabase Documentation: Database, Auth, and API Guide," 2024. [Online]. Available: https://supabase.com/docs
[12] OpenAI, "OpenAI API Technical Documentation," 2024. [Online]. Available: https://platform.openai.com/docs
[13] React.js, "React – A JavaScript Library for Building User Interfaces," Meta Open Source, 2024. [Online]. Available: https://react.dev
[14] C. Chakraborty, B. Gupta, and S. K. Ghosh, "Intelligent AI-Based Healthcare Cyber Security System Using Multi-Source Transfer Learning," ACM Transactions on Sensor Networks, vol. 19, no. 2, 2023.
[15] J. Xu, B. S. Glicksberg, and C. Su, "Federated Learning for Healthcare Informatics," *Journal of Healthcare Informatics Research*, vol. 5, no. 1, pp. 1-19, 2024.
[16] K. Singhal, T. Tu, J. Gottweis, et al., "Toward Expert-Level Medical Question Answering with Large Language Models," *Nature Medicine*, vol. 31, no. 3, pp. 943–950, 2025. [Online]. Available: https://doi.org/10.1038/s41591-024-03423-7
[17] L. B. Elvas, A. I. Almeida, and J. C. Ferreira, "The Role of AI in Cardiovascular Event Monitoring and Early Detection: Scoping Literature Review," *JMIR Medical Informatics*, vol. 13, e72410, 2025.
[18] A. Naghib, F. S. Gharehchopogh, and A. Zamanifar, "A Comprehensive and Systematic Literature Review on Intrusion Detection Systems in the Internet of Medical Things," *Artificial Intelligence Review*, vol. 58, no. 1, 2025. [Online]. Available: https://doi.org/10.1007/s10462-024-11101-w
[19] D. V. Dimov and S. H. Kim, "Privacy-Preserving Artificial Intelligence Models in Healthcare Security: A Synthesis," *Preprints*, 2026. [Online]. Available: https://www.preprints.org/manuscript/202601.0250
[20] M. Langote, S. Saratkar, and P. Kumar, "Human–Computer Interaction in Healthcare: Comprehensive Review," *AIMS Bioengineering*, vol. 11, pp. 343-390, 2024. [Online]. Available: https://doi.org/10.3934/bioeng.2024018
[21] S. Tudor, R. Bhatia, and M. Liem, "Opportunities and Challenges of Using Artificial Intelligence in Predicting Clinical Outcomes," *Journal of Medical Internet Research*, vol. 27, 2025.
[22] X. Meng, X. Yan, and K. Zhang, "The Application of Large Language Models in Medicine: A Scoping Review," *iScience*, vol. 27, no. 5, 109713, 2024. [Online]. Available: https://doi.org/10.1016/j.isci.2024.109713
[23] P. Dhade and P. Shirke, "Federated Learning for Healthcare: A Comprehensive Review," *RAiSE-2023*, vol. 59, no. 1, p. 230, 2024. [Online]. Available: https://doi.org/10.3390/engproc2023059230
[24] D. B. Olawade et al., "Human in the loop artificial intelligence in healthcare: applications, outcomes, and implementation challenges," *International Journal of Medical Informatics*, vol. 213, 106362, 2026.
[25] D. Rim, S. Nuriev, and Y. Hong, "Cyclic Training of Dual Deep Neural Networks for Discovering User and Item Latent Traits in Recommendation Systems," *IEEE Access*, vol. 13, pp. 10663-10677, 2025.
[26] Onur Asan, "Artificial Intelligence–Based Consumer Health Informatics Application: Scoping Review," *Journal of Medical Internet Research*, vol. 25, 2023.
[27] V. Hartman, X. Zhang, and R. Poddar, "Developing and Evaluating Large Language Model-Generated Emergency Medicine Handoff Notes," *JAMA Network Open*, vol. 7, no. 12, e2448723, 2024.
[28] A. Bagabir, N. K. Ibrahim, and H. Bagabir, "Integrating Artificial Intelligence in Clinical Practice, Hospital Management, and Health Policy," *Journal of Hospital Management and Health Policy*, vol. 9, 2025.
[29] Chloe Hinchliffe, Hugo Hiden, and Paul Watson, "Recent Methodological Advances in Federated Learning for Healthcare," *International Journal of Information Security*, Mar. 2026.
[30] S. Minaee, T. Mikolov, and N. Nikzad, "Large Language Models: A Survey," *arXiv*, 2024. [Online]. Available: https://arxiv.org/abs/2402.06196v1

[31] D. Van Veen, C. Van Uden, and L. Blankemeier, "Adapted Large Language Models Can Outperform Medical Experts in Clinical Text Summarization," *Nature Medicine*, vol. 30, no. 4, pp. 1134–1142, 2024.
[32] A. Agaronnik, J. Davis, and C. R. Manz, "Large Language Models to Identify Advance Care Planning in Patients With Advanced Cancer," *Journal of Pain and Symptom Management*, vol. 69, no. 3, pp. 243–250, 2025.
[33] P. Bhatt, "Emerging Artificial Intelligence–Empowered mHealth: Scoping Review," *JMIR mHealth and uHealth*, vol. 10, 2022.
[34] Yun Jiang, "Technological Innovations and Data-Driven Support for Older Adults," *JMIR Aging*, vol. 6, 2023