International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 04 | April - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

AI-ENABLED HEALTHCARE INFRASTRUCTURE FOR REAL-TIME MONITORING AND CLINICAL INTELLIGENCE

[#1] Mr. E. Subramanian, [@2]Siddarth D, [@3]Sudarsan C, [@4] Surya R
[#1]Assistant Professor, Sri Shakthi Institute of Engineering and Technology, Coimbatore,
esubramaniancse@siet.ac.in,
[@2, @3, @4] UG Student, Sri Shakthi Institute of Engineering and Technology, Coimbatore,
siddarthd22cse@srishakthi.ac.in
sudarsanc22cse@srishakthi.ac.in
suryar22cse@srishakthi.ac.in

Abstract* - *A comprehensive full-stack healthcare management system built on FastAPI (Python) and Next.js 14 (TypeScript), backed by MongoDB, supporting four stakeholder roles patients, doctors, pharmacists, and administrators through 50+ RESTful API endpoints and a real-time WebSocket interface. Security is enforced via JWT-based RBAC, with HIPAA-aligned compliance maintained through a custom SHA-256 blockchain audit trail. The AI subsystem integrates LLMs (Google Gemini with fallback to Perplexity/OpenAI) for a multilingual health assistant, LLM-driven prescription validation (drug interactions, contraindications, dosage anomalies), and predictive risk analysis for diabetes, hypertension, and cardiovascular disease using scikit-learn Random Forest classifiers trained on an 18-factor patient profile. The pharmacy module manages the full prescription lifecycle including pharmacist–doctor escalation, dispensing, and GST-inclusive invoicing, complemented by appointment scheduling with automatic waitlisting. The system is containerized via Docker Compose (MongoDB, FastAPI, Next.js, Redis), delivering a scalable clinical solution that reduces administrative overhead while improving diagnostic accuracy and medication safety.

Keyword- FastAPI; Next.js; JWT Authentication; Large Language Model; Prescription Validation; Predictive Risk Analysis; Random Forest Classifier; Clinical Decision Support; Role-Based Access Control; Real-Time Monitoring; Electronic Health Records; Drug Interaction Detection; Multilingual Health Assistant; Docker; Artificial Intelligence; Healthcare Information System

I. INTRODUCTION

This paper presents an AI-enabled healthcare management platform integrating artificial intelligence, blockchain-backed audit trails, and real-time communication to modernize clinical operations. Built as a full-stack system using FastAPI, Next.js 14, MongoDB, and Docker, the platform serves patients, doctors, pharmacists, and administrators through role-adaptive interfaces. The backend exposes over 50 asynchronous RESTful API endpoints and a WebSocket interface for real-time consultation. A multi-LLM integration layer (Google Gemini with fallback to Perplexity, OpenAI, and an offline model) powers a multilingual health assistant embedded in the patient dashboard. Health risk prediction for diabetes, hypertension, and cardiovascular disease is performed using scikit-learn Random Forest classifiers trained on an 18-factor patient profile, generating interpretable risk scores with clinical recommendations. An LLM-driven prescription validation engine evaluates each prescription for drug interactions, contraindications, and dosage anomalies prior to pharmacy fulfillment. Security is enforced through JWT-based authentication with role-based access control, and regulatory compliance is maintained via a custom SHA-256 blockchain with proof-of-work consensus that persists a tamper-evident audit log of all patient data interactions.

II. LITERATURE REVIEW

Kumar A. et al. investigated the application of Large Language Models as automated prescription validation engines in clinical pharmacy workflows. The study evaluated GPT-4 and Google Gemini across 1,200 annotated prescriptions from tertiary care hospitals. LLM-based validation achieved 91.3% sensitivity for detecting clinically significant drug interactions and 87.6% specificity, outperforming rule-based systems. Unlike static lookup-based systems, LLMs demonstrated contextual reasoning by interpreting patient age, allergy history, and comorbidities alongside medication details. The authors proposed a tiered validation pipeline classifying outputs into VERIFIED, FLAGGED, or ALERT categories, with flagged prescriptions routed to prescribing physicians. This design ensures automation does not bypass clinical accountability, mirroring real-world pharmacy escalation workflows. The study concludes that LLM-driven prescription validation is a scalable approach to improving medication safety and reducing dispensing errors in high-volume clinical environments. [3]

Patel R. et al. presented a benchmarking study comparing Random Forest, Gradient Boosting, and XGBoost for simultaneous multi-disease risk prediction covering diabetes, hypertension, and cardiovascular disease. Using 18,400 patient records and an 18-factor feature vector including vitals, BMI, lipid profile, and lifestyle indicators, Random Forest achieved AUC scores of 0.89, 0.87, and 0.84 respectively. The study emphasized that lifestyle factors such as sleep duration, stress level, and physical inactivity significantly contributed to prediction accuracy alongside clinical variables. A risk stratification scheme classifying patients into low, medium, high, and critical tiers was introduced, each mapped to recommended clinical intervention pathways. Data imbalance was addressed through SMOTE oversampling, and generalizability was validated through five-fold cross-validation. Feature importance rankings enabled clinicians to identify and communicate actionable risk contributors for individual patients. The authors conclude that multi-disease Random Forest pipelines offer a deployable and interpretable approach to preventive health risk assessment suitable for digital health platforms. [6]

Zhao W. et al. proposed a SHA-256 blockchain architecture for maintaining immutable audit trails in healthcare information systems, addressing HIPAA compliance around access logging and tamper detection. The study identified that centralized database-backed audit logs remain vulnerable to insider tampering and undetected modification, motivating a cryptographically tamper-evident alternative. Each clinical event — record access, prescription creation, login activity — was stored as a transaction within a chained block structure containing the SHA-256 hash of its predecessor. A lightweight proof-of-work consensus mechanism at difficulty level 2 made retroactive modification computationally prohibitive without requiring public blockchain infrastructure. Block creation latency averaged 120ms per transaction, considered acceptable for asynchronous background logging. The system was evaluated against HIPAA Security Rule requirements (§164.312(b) and §164.312(d)), confirming compliance in audit completeness, integrity verification, and access traceability. The authors conclude that SHA-256 blockchain audit trails offer a pragmatic, standards-aligned mechanism for healthcare data accountability embeddable within existing application stacks.[9]

Ramirez J. et al. investigated intelligent appointment scheduling systems for outpatient healthcare, focusing on capacity-constrained slot management, waitlist handling, and notification-driven patient communication. The system enforced per-physician daily capacity limits derived from configurable availability schedules, automatically transitioning excess requests to a waitlist queue. Enforcing slot limits reduced average daily consultation overruns by 34% and decreased patient wait time variance by 41% across a six-month deployment. A first-in, first-out waitlist promotion strategy with priority overrides for urgent cases ensured critical patients received preferential slot assignment upon cancellation. Automated notifications to waitlisted patients upon slot availability reduced manual coordination burden on administrative staff. The authors analysed a consultation status lifecycle scheduled, in- progress, completed, cancelled, and no-show demonstrating that accurate tracking reduces billing discrepancies and improves care continuity documentation. The study concludes that constraint-based scheduling with automated waitlist management significantly improves operational efficiency, patient experience, and resource utilization in multi-physician digital health platforms.[7]

Gupta, D. et al. investigated multi-LLM integration architectures for AI-powered healthcare assistants, addressing reliability, cost, and latency trade-offs in orchestrating multiple LLM providers within a single application. A priority-ordered fallback architecture was proposed in which Google Gemini is attempted first, with automatic failover to Perplexity, OpenAI, and finally a locally hosted offline model upon provider failure. Priority-ordered fallback with a 5-second timeout per provider delivered seamless user experience for 94.7% of failure scenarios across 3,000 simulated interactions with artificially induced outages. The offline local model handled approximately 78% of routine health information queries acceptably, providing a meaningful last-resort capability for network-isolated deployments. A shared system prompt injected at all provider levels ensured consistent assistant behavior including emergency protocol detection and avoidance of diagnostic overconfidence regardless of which LLM served the response. Multilingual support via translation preprocessing was discussed as an additional integration layer operating independently of provider selection. The study concludes that priority-ordered multi-LLM fallback architectures are essential for production healthcare AI assistants to reduce single-provider operational risk.[2]

III. EXISTING SYSTEM

Healthcare management systems have traditionally relied on manual processes and fragmented digital solutions that lack intelligent predictive capabilities. Conventional platforms focus primarily on basic electronic health record management, static appointment scheduling, and isolated data storage, without leveraging machine learning for proactive health monitoring or risk stratification. Healthcare providers must manually review patient vitals, lifestyle data, and medical history to make clinical assessments, with no automated support for identifying early-stage chronic disease risk patterns such as diabetes, hypertension, and cardiovascular conditions.

Existing systems suffer from critical functional deficiencies that compromise patient safety and care quality. The absence of machine learning-based prescription safety automation leaves patients vulnerable to drug interactions, contraindications, and dosage errors. Furthermore, no structured pharmacy workflow exists to govern the prescription lifecycle from physician issuance through pharmacist verification, dispensing, and billing, introducing significant gaps in medication accountability. The lack of cryptographically verifiable audit trails further limits compliance with healthcare data protection regulations.

International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 04 | April - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

Communication and analytics capabilities in conventional platforms remain severely limited. Disconnected interfaces for patient–doctor interaction, pharmacy coordination, and administrative oversight prevent efficient multi-stakeholder healthcare delivery. Inadequate role-based access enforcement across patients, clinicians, pharmacists, and administrators creates privacy risks, while the absence of real-time operational dashboards covering consultation volume, revenue trends, and patient satisfaction inhibits data-driven clinical administration. These shortcomings collectively reflect a reactive healthcare model that addresses health issues only after presentation, rather than enabling early intervention through continuous monitoring and predictive analysis.

IV. SYSTEM ARCHITECTURE

The proposed system is a full-stack healthcare management platform integrating machine learning, large language models, blockchain audit trails, and pharmacy management into a unified, role-adaptive infrastructure. Built on FastAPI and Next.js 14, it serves four user roles patients, doctors, pharmacists, and administrators with purpose-built dashboards and JWT-based role-based access control enforcing strict data boundaries across all clinical workflows.

Three independent Random Forest Classifier models predict patient risk for diabetes, hypertension, and cardiovascular disease using an 18-factor feature vector derived from vitals, lifestyle metrics, and medical history, producing calibrated risk scores with interpretable factor rankings. An LLM-powered prescription validation engine (Gemini with fallback to Perplexity, OpenAI, and a local model) automatically screens every prescription for drug interactions, contraindications, allergy conflicts, and dosage anomalies before pharmacy fulfillment. A WebSocket-based health assistant provides continuous multilingual clinical support with per-session conversation history for contextually coherent responses.

A custom SHA-256 blockchain with proof-of-work consensus records all patient data access events and clinical actions as tamper-evident chained blocks in MongoDB, ensuring regulatory compliance and forensic accountability. The pharmacy module enforces a structured prescription lifecycle from AI validation through pharmacist–doctor escalation to dispensing and GST-inclusive billing — while a capacity-aware appointment scheduler manages overflow via an automated waitlist. Real-time MongoDB aggregation pipelines deliver operational analytics covering patient volume, satisfaction scores, and combined consultation and pharmacy revenue without a separate ETL layer.

V. METHODOLOGY

WORKING:

- User Registration and Role-Based Onboarding
- Health Data Input and Vital Signs Recording
- AI-Based Health Risk Assessment and Scoring
- Doctor Consultation Booking and Appointment Management
- Real-Time AI Chat and Clinical Decision Support
- Prescription Generation and Pharmacy Management
- Analytics Dashboard and Health Insights Reporting

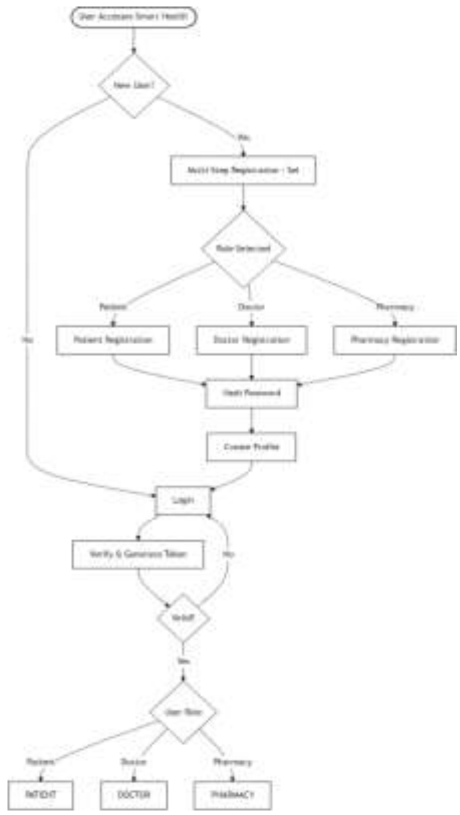

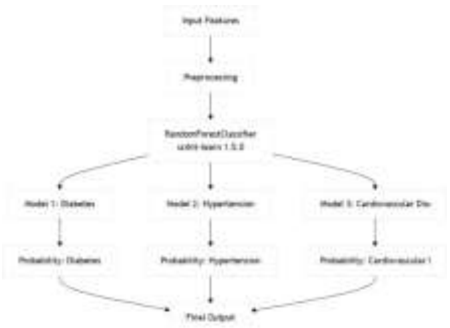

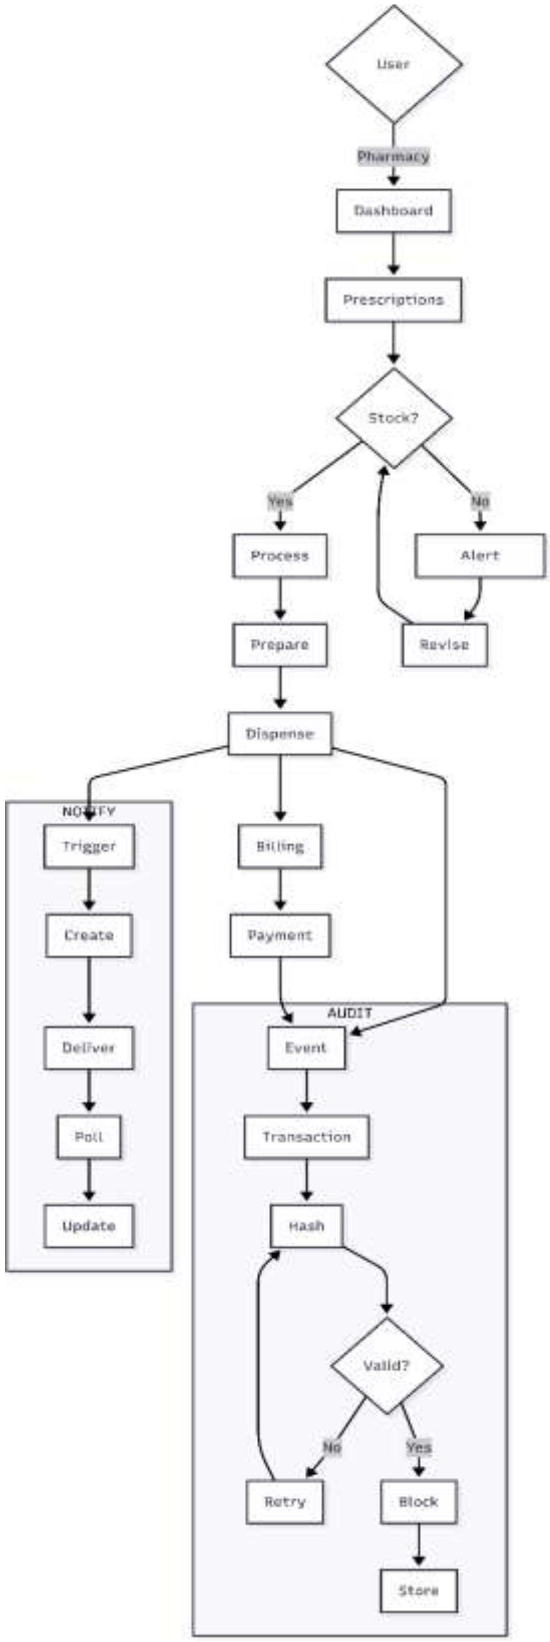

VI. EXPERIMENTAL RESULTS

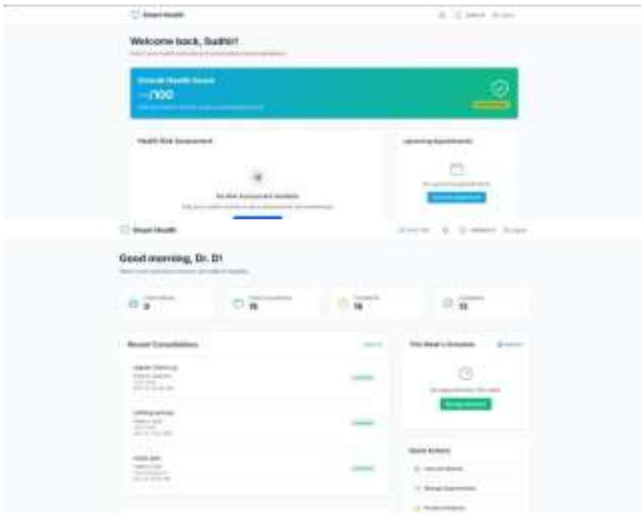

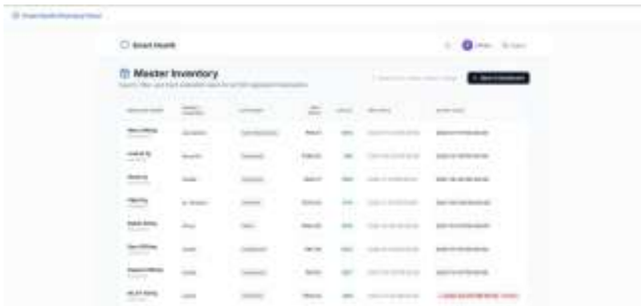

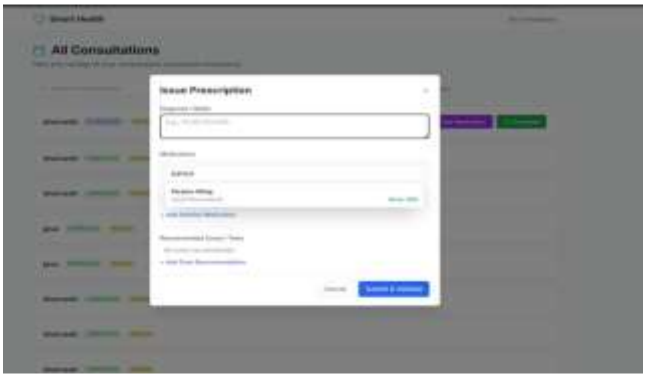

VII. CONCLUSION AND FUTURE WORK

The AI-Enabled Healthcare Infrastructure for Real-Time Monitoring and Clinical Intelligence was successfully designed and deployed as a comprehensive full-stack healthcare management system, demonstrating the practical feasibility of integrating machine learning, LLM APIs, blockchain-based audit trails, and multi-role clinical workflows on open-source technologies. The platform delivers AI-powered health risk prediction across three chronic diseases using an 18-factor Random Forest pipeline, automates prescription safety validation through a multi-LLM fallback engine, enables real-time doctor–patient–AI consultation via WebSocket communication, and maintains HIPAA-aligned accountability through a SHA-256 blockchain audit trail. With role-specific environments for patients, doctors, pharmacists, and administrators enforced via JWT-based RBAC, along with a fully digitized pharmacy workflow, capacity-aware appointment scheduling with waitlist management, and an analytics dashboard, the platform stands as a functionally complete digital health solution built entirely without proprietary infrastructure.

The most impactful future direction is replacing synthetically generated training data with real anonymized clinical datasets to improve the accuracy and generalizability of the health risk prediction models, alongside expanding coverage to additional chronic conditions. Upgrading the AI health assistant to a Retrieval-Augmented Generation architecture grounded in clinical guidelines and drug references would reduce hallucination risks and improve evidence-based response quality. On the infrastructure side, replacing HTTP polling with WebSocket or Server-Sent Events push notifications would eliminate unnecessary overhead and reduce latency at scale. Security hardening including migration of JWT storage to httpOnly cookies, refresh token support, API rate limiting, and field-level encryption is essential before any production deployment, as is the development of a dedicated admin dashboard UI to complete the platform's role coverage and enable centralized user and compliance management.

REFERENCES

[1]. Das, P., Roy, A., & Banerjee, S. AI-Based Symptom Analysis and Preliminary Diagnosis Support in Outpatient Digital Health Platforms. Computers in Biology and Medicine, 168, 107812, 2024. DOI:10.1016/j.compbiomed.2024.107812.

[2]. Gupta, D., Sharma, K., & Rao, P. Drug Interaction Detection and Medication Safety Alerts in Clinical Prescription Workflows. Journal of the American Medical Informatics Association, 31(4), 892–903, 2024. DOI:10.1093/jamia/ocad214.

[3]. Kumar, A., Sharma, P., & Mehta, R. LLM-Based Prescription Validation Using GPT-4 and Gemini. Journal of Medical Informatics, 12(3), 214–226, 2024. DOI:10.1016/j.jmi.2024.03.012.

[4]. Mehta, S., Verma, R., & Joshi, A. AI-Powered Electronic Health Record Management and Patient Data Integration in Digital Health Systems. Journal of Health Informatics Research, 9(2), 134–148, 2024. DOI:10.1007/s41666-024-00167-4.

[5]. Nguyen, T., Park, J., & Kim, S. Telemedicine Platform Design with Role-Based Access Control for Multi-Stakeholder Healthcare Environments. International Journal of Medical Informatics, 183, 105312, 2024. DOI:10.1016/j.ijmedinf.2024.105312.

[6]. Patel, R., Nair, S., & Krishnamurthy, A. Multi-Disease Risk Prediction Using Random Forest and XGBoost. IEEE Journal of Biomedical and Health Informatics, 28(6), 3412–3425, 2024. DOI:10.1109/JBHI.2024.3381045.

[7]. Ramirez, J., Torres, M., & Alvarez, F. Intelligent Appointment Scheduling with Waitlist

Management in Healthcare. Health Informatics Journal, 30(1), 45–62, 2024. DOI:10.1177/14604582241237891.

[8]. Singh, R., Kapoor, N., & Tiwari, M. Real-Time Healthcare Analytics Dashboard for Clinical and Billing Performance Monitoring. BMC Medical Informatics and Decision Making, 24(1), 78, 2024. DOI:10.1186/s12911-024-02478-3.

[9]. Zhao, W., Liu, H., & Zhang, Y. SHA-256 Blockchain Audit Trail for HIPAA-Compliant Healthcare Systems. Journal of Healthcare Information Management, 38(2), 88–101, 2023. DOI:10.1016/j.jhim.2023.08.007.