iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** III **Month of publication:** March 2026

DOI: https://doi.org/10.22214/ijraset.2026.78188

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue III Mar 2026- Available at www.ijraset.com

Telemedicine Web Platform with Live Video Consultation, Health Records & AI Based Diagnostics Assistance

Vetrivel M[1], Prithivi Raj S[2], Sathish Kumar G[3], Yuvaraj M[4], Neelamegam K[5]
[1,2,3,4,5]*Department of Computer Science and Engineering, Arunai Engineering College, Tiruvannamalai, Tamil Nadu, India*

Abstract: ***The rapid proliferation of internet-connected devices and the consequent evolution of cloud-computing infrastructure have created an unprecedented opportunity to redesign the way healthcare services are delivered. Traditional healthcare systems are burdened by geographical barriers, administrative overheads, fragmented record-keeping, and constrained access to specialist physicians—challenges that have been accentuated in the post-pandemic era. This paper presents TeleMed, a comprehensive, full-stack telemedicine web platform that integrates live video consultations, electronic health record (EHR) management, and artificial intelligence (AI)-based diagnostic assistance into a single, unified ecosystem. The platform is architected as a decoupled client-server application powered by a React.js single-page frontend, a Java 17 / Spring Boot 3.0 backend secured with JSON Web Token (JWT) authentication, and a hybrid database model comprising PostgreSQL for relational transactional data and MongoDB for unstructured medical records. The AI diagnostic module leverages the Meta Llama 3 large language model (LLM), accessed through the HuggingFace Inference Interface, employing advanced zero-shot and few-shot prompt engineering with dynamic patient-context injection to deliver highly personalized symptom analysis without the overhead of fine-tuning custom datasets. Real-time video consultations are facilitated through WebRTC standards via the Jitsi Meet API, enabling secure, low-latency virtual rooms. Experimental validation confirms that TeleMed successfully unifies intelligent triage, doctor discovery, appointment scheduling, and live consultation into a scalable platform, offering a practical blueprint for next-generation remote healthcare delivery.***

Keywords: Telemedicine, Electronic Health Records (EHR), Large Language Model (LLM), Meta Llama 3, WebRTC, Jitsi Meet API, Spring Boot, React.js, JWT Authentication, Hybrid Database, AI Diagnostics, Remote Healthcare

I. INTRODUCTION

The global healthcare sector is undergoing a transformative shift driven by digital technologies, widespread mobile internet adoption, and a growing demand for patient-centric service delivery models. The conventional model—wherein patients must physically visit clinics or hospitals to receive diagnosis, treatment, or even routine follow-up—creates systemic inefficiencies: prolonged waiting times, high administrative costs, and restricted access for patients in remote or underserved regions. The COVID-19 pandemic acted as a catalyst, accelerating the adoption of digital health solutions and exposing the fragility of healthcare systems that lack robust remote-care capabilities.

Telemedicine, broadly defined as the delivery of clinical services via telecommunications technology, has emerged as a viable and increasingly preferred mechanism for bridging the gap between patients and healthcare professionals. Early iterations of telemedicine relied primarily on asynchronous communication—secure messaging, email-based consultations, and store-and-forward imaging. Subsequent advances introduced synchronous video consultations, enabling face-to-face interaction between patients and physicians regardless of physical location. However, a critical shortcoming persists in the majority of existing platforms: service fragmentation. Video consultation tools, health record management systems, and diagnostic support engines are typically deployed as isolated applications, compelling both clinicians and patients to navigate multiple disjointed interfaces during a single care episode.

This fragmentation introduces significant cognitive overhead for healthcare professionals, increases the probability of data transcription errors, and degrades the overall patient experience. Furthermore, legacy AI-assisted triage systems that rely on rigid, rule-based decision trees frequently fail to capture the nuanced, context-dependent nature of clinical symptomatology. Systems fine-tuned on monolithic proprietary medical datasets suffer from rapid obsolescence and require expensive retraining pipelines to stay current with evolving clinical knowledge.

©IJRASET: All Rights are Reserved

In response to these challenges, this paper introduces TeleMed—a holistic, next-generation telemedicine platform that consolidates live video consultation, dynamic EHR management, intelligent AI-driven diagnostics, and streamlined appointment scheduling within a single, cohesive application. The core technical contributions of this work are as follows:

1) A scalable, decoupled full-stack architecture using React.js (frontend) and Java 17 / Spring Boot 3.0 (backend) with JWT-based session management.
2) A hybrid database design employing PostgreSQL for relational data integrity and MongoDB for flexible, document-oriented medical records.
3) An AI Diagnostic Module powered by Meta Llama 3 (meta-llama/Meta-Llama-3-8B-Instruct) that employs dynamic patient-context injection via zero-shot and few-shot prompt engineering to deliver personalized clinical insights.
4) Seamless WebRTC-based live video consultations through the Jitsi Meet API, with simultaneous access to patient EHRs within the same interface.
5) An intelligent doctor-discovery engine that cross-references AI-generated specialization recommendations with a live physician registry in PostgreSQL, prioritizing results by availability.

The remainder of this paper is organized as follows: Section II details the system architecture including the frontend, backend, and hybrid database design. Section III presents the comprehensive methodology and implementation of each module. Section IV surveys the relevant literature across telemedicine platforms, AI in healthcare, and WebRTC communication standards. Section V enumerates the hardware and software requirements that constitute the TeleMed development environment. Section VI presents and interprets the experimental results across all major system components. Section VII concludes the paper with a synthesis of contributions and a structured roadmap for future research and development.

A. *Problem Statement and Motivation*

The fundamental inefficiencies inherent in conventional healthcare delivery systems can be broadly categorized into three interconnected dimensions: accessibility barriers, information fragmentation, and diagnostic support inadequacy. Each of these dimensions independently degrades care quality; in combination, they create a systemic failure mode that disproportionately impacts patients in geographically remote, economically disadvantaged, or time-constrained circumstances. Accessibility barriers manifest most acutely in regions where the ratio of specialist physicians to population is critically low. According to World Health Organization (WHO) projections, a global shortage of 18 million health workers is anticipated by 2030, with the deficit concentrated in low- and middle-income countries. Even in well-resourced healthcare systems, specialist appointments frequently involve wait times measured in weeks or months. Telemedicine platforms that enable asynchronous scheduling and synchronous consultation can substantially compress these timelines, but only if the appointment process itself is intelligent enough to direct patients to the right specialist without requiring them to first navigate the conventional referral pathway. Information fragmentation arises when a patient's clinical data—diagnostic history, medication records, allergy profiles, previous consultation notes, and imaging results—is distributed across multiple incompatible systems that cannot communicate with one another in real time. A physician conducting a telehealth consultation without access to the patient's complete history is forced to rely on the patient's potentially incomplete self-report, a situation that increases the risk of adverse drug interactions, duplicate diagnostic testing, and suboptimal treatment decisions. The absence of integrated EHR access within the consultation interface is therefore not merely a user-experience shortcoming; it is a patient safety concern.

Diagnostic support inadequacy refers to the limitations of current AI-assisted triage tools in providing contextually relevant, patient-specific guidance. Rule-based symptom checkers—widely deployed in patient-facing mobile applications—operate on static decision trees that cannot adapt to a patient's individual clinical context. A symptom that carries different clinical significance depending on the patient's age, sex, comorbidities, and current medications requires a reasoning capability that static rule trees fundamentally lack. TeleMed's LLM-based diagnostic module with dynamic context injection directly addresses this limitation.

B. *Scope and Objectives*

The scope of this research encompasses the complete design, implementation, and empirical evaluation of the TeleMed platform as a full-stack web application. The study is bounded to the development and validation of a functional prototype system capable of supporting the core telemedicine use cases described above; it does not include large-scale clinical trials, regulatory submission preparation, or enterprise-grade deployment infrastructure provisioning. The primary objectives of this research are as follows: first, to design and implement a secure, scalable, and maintainable full-stack telemedicine application architecture; second, to develop an AI-powered diagnostic assistance module that leverages LLM technology with patient-specific context injection for personalized

clinical guidance; third, to integrate real-time video consultation capabilities alongside live EHR access within a single unified interface; fourth, to implement an intelligent physician discovery and appointment scheduling system that bridges AI recommendations with a live doctor registry; and fifth, to validate the integrated platform's functional correctness, performance characteristics, and security posture through systematic evaluation.

II. SYSTEM ARCHITECTURE

The TeleMed platform is engineered as a highly scalable, full-stack web application. To ensure optimal performance, security, and seamless data management across disparate data types, the system adopts a decoupled client-server architecture supported by a hybrid database model. The architectural philosophy prioritizes modularity, enabling independent scaling of the frontend, backend, and database layers as system load evolves.

A. Frontend Architecture

The user interface is designed as a fast, responsive Single Page Application (SPA) that delivers a seamless experience across both patient and physician dashboards. The frontend is developed using React.js, a declarative, component-based JavaScript library that enables efficient UI rendering through its virtual DOM mechanism. The project is bootstrapped via Vite—a next-generation build tool that significantly reduces development server startup times and produces highly optimized production bundles through native ES module support and Rollup-based bundling.

Visual presentation and responsive layouts are achieved using Tailwind CSS, a modern, utility-first CSS framework that eliminates the overhead of writing custom stylesheets. Tailwind's atomic class system enables rapid prototyping and consistent design language across all components. The SPA architecture ensures that navigation between dashboards, consultation pages, and medical record views occurs without full-page reloads, dramatically improving perceived responsiveness.

State management is handled locally within component trees using React's built-in hooks (useState, useEffect, useContext), supplemented by API calls to the backend for persistent data operations. The Chat.jsx component serves as the central interface for live consultations, dynamically initializing the Jitsi Meet API iframe and providing synchronized access to EHR data during video sessions.

B. Backend Architecture

The core server-side logic is constructed using Java 17 within the Spring Boot 3.0 framework. Java 17, a Long-Term Support (LTS) release, provides enhanced language features including sealed classes, pattern matching for instanceof expressions, and records—enabling more expressive, concise, and maintainable code. Spring Boot's opinionated auto-configuration significantly reduces boilerplate setup, allowing developers to focus on business logic rather than infrastructure configuration.

The backend exposes a RESTful API layer that handles secure routing, complex business logic, cross-module communication, and integration with both the hybrid database layer and external AI services. API endpoints are organized into domain-specific controllers: AuthController for authentication flows, PatientController for health record operations, DoctorController for physician management, AppointmentController for scheduling workflows, and AIDiagnosticsController for AI triage interactions.

System security is comprehensively enforced using Spring Security, which implements JSON Web Token (JWT) authentication. Upon successful login, the server issues a digitally signed JWT that encodes the user's identity and role claims. All subsequent requests to protected endpoints must carry this token in the Authorization header; the server validates the token's signature and expiry on every request, ensuring that sessions remain stateless and scalable without server-side session storage.

C. Hybrid Database Design

To manage the complex and varied data generated by a comprehensive telemedicine platform, a dual-database approach is employed. This hybrid model acknowledges that different categories of medical data have fundamentally different structural requirements and access patterns, and that no single database paradigm optimally serves all use cases.

1) *Relational Data Management (PostgreSQL):* PostgreSQL is utilized for highly structured, transactional datasets. Its ACID (Atomicity, Consistency, Isolation, Durability) compliance guarantees data integrity for mission-critical operations. This includes the management of Role-Based Access Control (RBAC) policies, user authentication profiles, doctor specialization records, and appointment scheduling logic. Relational joins enable efficient cross-referencing of doctor availability with AI-recommended specializations during the doctor-discovery workflow.

2) *Unstructured Data Management (MongoDB):* MongoDB, a leading NoSQL document-oriented database, is configured to handle rapidly evolving, schema-flexible data. Electronic Health Records (EHRs) naturally resist rigid tabular schemas: a patient's medical history may include diverse data types—structured lab results, free-text clinical notes, digitized prescriptions, allergy annotations, and imaging metadata. MongoDB's BSON document model accommodates this heterogeneity natively, and its horizontal scalability ensures that growing EHR datasets do not become performance bottlenecks.

III. METHODOLOGY AND IMPLEMENTATION

The implementation of TeleMed integrates advanced AI capabilities with real-time communication protocols and a secure, multi-role user management system to facilitate end-to-end remote healthcare. Each module is independently developed and integrated through well-defined API contracts, ensuring maintainability and testability.

A. AI Diagnostic Assistance Module

The platform's AI diagnostic tool is powered by the Meta Llama 3 model, specifically the meta-llama/Meta-Llama-3-8B-Instruct variant—an instruction-tuned, open-weight large language model developed by Meta AI. Unlike proprietary medical AI systems that depend on closed, domain-specific training datasets, TeleMed leverages the foundational reasoning and language comprehension capabilities of Llama 3, accessed through the HuggingFace Inference Interface via the OpenAI-SDK-compatible endpoint (router.huggingface.co/v1). The 8B-parameter model was selected over larger Llama 3 variants as it offers an optimal trade-off between reasoning capability and inference speed for the real-time consultation use case.

The critical innovation lies in the Dynamic Context Generation mechanism. Rather than relying on static, pre-programmed decision trees, the system dynamically retrieves each patient's current health profile from MongoDB at query time—including their electronic health record, active prescriptions, known allergies, current medications, and recent consultation notes. This data is then seamlessly injected into the LLM's system prompt through carefully crafted zero-shot and few-shot prompt engineering templates.

The system prompt structure follows a layered architecture. The outermost layer establishes the model's operational role and ethical constraints, instructing it to function as a clinical decision support assistant, to recommend medical specializations rather than definitive diagnoses, and to include appropriate medical disclaimers. The middle layer injects the patient's dynamic context in a structured format—presented as a clinical summary with labeled sections for demographic information, chief complaint, active conditions, current medications, known allergies, and recent consultation history. The innermost layer presents the patient's current symptom description in natural language and requests a structured response enumerating recommended specializations with supporting clinical reasoning.

Zero-shot prompting instructs the model to perform medical specialization classification based solely on symptom descriptions and the injected patient context, without providing example input-output pairs. Few-shot prompting supplements this with curated medical reasoning examples to improve the precision of specialty recommendations for ambiguous symptom clusters. For example, the few-shot examples distinguish between headache presentations that warrant Neurology referral (severe, thunderclap onset, associated with neck stiffness) versus those appropriate for Primary Care (tension-type, no neurological symptoms) versus those requiring Ophthalmology (associated with visual disturbance). This granular example set substantially improves the model's discrimination accuracy for overlapping symptom profiles.

The API integration follows a non-blocking asynchronous call pattern implemented through Java's CompletableFuture framework within the AIDiagnosticsController. When the patient submits a symptom query, the controller immediately dispatches an asynchronous context retrieval task to MongoDB while simultaneously constructing the base prompt template. Upon context retrieval completion, the patient data is merged into the prompt, and the assembled request is forwarded to the HuggingFace endpoint with a configured timeout of 30 seconds. The response is parsed from the API's JSON message structure, medical disclaimer text is appended, and the result is returned to the React frontend for display. This pipeline ensures that EHR context retrieval does not sequentially block the API call, minimizing perceived response latency.

The module exposes two user-facing interfaces: the AI Doctor Finder, which accepts natural-language symptom inputs and returns recommended medical specializations along with contextually appropriate health disclaimers; and MedBot, a conversational chat interface that allows patients to engage in extended medical Q&A sessions with the AI, with all responses continuously grounded in the patient's live health data. The conversation history for MedBot sessions is maintained in the React component state and prepended to each subsequent API call to provide the model with the full conversational context, enabling coherent multi-turn clinical dialogues.

B. Live Video Consultation Integration

Real-time telehealth communication is facilitated through the WebRTC (Web Real-Time Communication) standards—an open-source framework enabling peer-to-peer audio, video, and data transmission directly within web browsers without requiring proprietary plugins or native application installations. TeleMed leverages the Jitsi Meet API, an open-source, enterprise-grade video conferencing solution deployed over meet.guifi.net nodes, to establish secure, low-latency virtual consultation rooms.

The integration is implemented within the Chat.jsx React component, which dynamically initializes the Jitsi Meet API iframe when a confirmed appointment session begins. Each consultation room is uniquely identified and cryptographically isolated, ensuring that only the authenticated patient and their assigned physician can access the session. The Jitsi integration supports full-HD video feeds, bidirectional audio, screen sharing capabilities, and in-session text chat—all within the same browser tab.

A key architectural achievement is the synchronous availability of the patient's complete EHR within the same consultation interface. During a live video session, the physician can review, annotate, and update the patient's MongoDB-stored health records in real time without leaving the consultation view, enabling evidence-based clinical decision-making at the point of care. Built-in chat functionality further allows asynchronous message exchange, prescription sharing, and follow-up instructions to be communicated securely within the platform.

C. Role-Based Access Control (RBAC)

TeleMed implements a comprehensive Role-Based Access Control system to govern access to sensitive medical data and platform functionalities. Three primary roles are defined: Patient, Doctor, and Administrator. RBAC policies are stored and enforced through PostgreSQL, with Spring Security's method-level security annotations (@PreAuthorize, @Secured) ensuring that API endpoints respond only to requests bearing JWT tokens with the appropriate role claims.

Patients have read-write access to their own health records and appointment history, and read-only access to doctor profiles. Doctors have read-write access to their patients' EHRs during and after consultations, manage their own availability schedules, and can issue digital prescriptions. Administrators have system-wide access for managing user registrations, resolving account disputes, and monitoring platform health metrics.

D. Authentication and Session Management

Authentication is the gateway to the TeleMed platform and the foundation of its security architecture. The system implements a stateless, token-based authentication mechanism using JSON Web Tokens (JWT) as the session credential. When a user submits their login credentials, the Spring Security authentication manager validates the credential pair against the PostgreSQL user profile store using BCrypt password hashing with a work factor of 12, providing strong resistance against brute-force and rainbow table attacks.

Upon successful authentication, the server generates a signed JWT using the HS256 algorithm with a server-side secret key. The token payload encodes the user's unique identifier, assigned role (PATIENT, DOCTOR, or ADMIN), and a configurable expiration timestamp—24 hours for patient sessions and 8 hours for physician sessions. The token is returned to the client in the authentication response and stored in the browser's in-memory application state rather than localStorage, preventing XSS-based token theft.

Subsequent API requests include the JWT in the Authorization header using the Bearer scheme. The server-side JWT filter intercepts all requests to protected endpoints, validates the token signature and expiry, and populates the Spring Security context with the authenticated principal's role authorities. This stateless design eliminates the need for server-side session storage, enabling horizontal scaling without sticky session requirements.

Token refresh is handled through a dedicated /auth/refresh endpoint that accepts a valid, non-expired JWT and returns a new token with a reset expiration window. A logout endpoint invalidates the client-side token by clearing the in-memory state and optionally adding the token's JTI (JWT ID) claim to a server-side short-lived denylist for immediate invalidation before natural expiry.

E. Appointment Scheduling and Notification System

The appointment scheduling subsystem bridges the AI diagnostic output and the live consultation experience. After receiving specialization recommendations from the Llama 3 engine, the frontend presents the patient with a dynamically filtered list of available physicians, sorted by earliest available appointment slot. The slot availability data is retrieved from PostgreSQL through the AppointmentController API, which enforces non-overlapping booking constraints at the database level using optimistic row-level locking.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

Appointment slots are pre-populated by physicians through a dedicated schedule management interface within their dashboard. Physicians define their working hours, consultation durations (15, 30, or 60-minute increments), and advance booking windows. The scheduling algorithm partitions working periods into discrete, non-overlapping slots with states of AVAILABLE, BOOKED, or BLOCKED, managed atomically through JPA transactions to prevent double-booking under concurrent access conditions.

Upon appointment confirmation, the system dispatches notification events to both the patient and the physician. Email notifications are generated through the Spring Mail abstraction layer configured with SMTP relay, and include the appointment date and time, physician name and specialization, a one-click consultation room access link, and pre-appointment preparation guidance. Reminder notifications are dispatched 24 hours and 30 minutes prior to the scheduled consultation time through Spring's @Scheduled background task framework.

F. End-to-End System Workflow

The complete patient journey through the TeleMed platform follows a secure, streamlined, and intelligent workflow. The seven-stage process is summarized in Table I and described below.

1) Authentication: The patient navigates to the TeleMed web application and logs in using their registered credentials. Spring Security validates the credentials against the PostgreSQL user profile store, issues a signed JWT, and initializes a customized role-aware dashboard session.

2) AI Symptom Input: The patient accesses the AI Doctor Finder modal from their dashboard and describes their current symptoms in plain natural language, removing the requirement for prior medical knowledge or terminology.

3) Diagnostic and Specialty Matching: The Llama 3 engine processes the symptom text in the context of the patient's dynamically injected health profile and returns a list of recommended medical specializations along with relevant disclaimers and preliminary health guidance.

4) Intelligent Doctor Discovery: The backend's doctor-discovery service cross-references the AI-recommended specializations against the PostgreSQL registry of registered physicians, applying a priority algorithm that ranks results by the earliest available appointment slot.

5) Appointment Confirmation: The patient reviews the ranked list of relevant specialists, selects a preferred physician, and confirms a precise consultation time slot. A confirmation notification is dispatched to both the patient and the selected physician.

6) Virtual Consultation: At the confirmed appointment time, the patient accesses the Chat.jsx consultation interface. The Jitsi Meet API iframe is dynamically initialized, establishing an exclusive, authenticated WebRTC room connecting the patient and physician with full audio-video capability.

7) EHR Update and Prescription Issuance: Following the consultation, the physician updates the patient's MongoDB-stored EHR with diagnosis notes, revised prescriptions, and follow-up recommendations. The updated records are immediately accessible to the patient from their dashboard.

TABLE I End-to-End Patient Consultation Workflow in TeleMed

Step	Stage	Description
1	Authentication	Patient logs into their personalized dashboard; JWT session initialized.
2	AI Symptom Input	Patient enters symptoms in natural language via the AI Doctor Finder modal.
3	Diagnostic Matching	Llama 3 parses symptoms and returns recommended medical specializations.
4	Doctor Discovery	Backend cross-references AI output with PostgreSQL doctor registry, prioritized by availability.
5	Appointment Booking	Patient selects a specialist and confirms a consultation time slot.
6	Virtual Consultation	Chat.jsx initializes Jitsi Meet iframe; patient and doctor connect via secured WebRTC room.
7	EHR Access	Doctor accesses and updates patient's MongoDB-stored EHR during or after consultation.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue III Mar 2026- Available at www.ijraset.com

IV. LITERATURE REVIEW

The evolution of telemedicine has been marked by several distinct developmental phases, from rudimentary telephonic consultations to sophisticated multi-modal video-conferencing solutions. A systematic review of existing literature and contemporary commercial platforms reveals a persistent and critical fragmentation in service delivery that TeleMed directly addresses.

A. Traditional Telemedicine Platforms

Foundational telemedicine research established the clinical equivalence of video consultations to in-person visits for a broad range of non-emergency conditions [1].

Early platforms such as Teladoc and Amwell demonstrated that synchronous video consultations could significantly reduce patient wait times and increase access to specialist care in rural regions [2]. However, these platforms uniformly treated video consultation and medical record management as mutually exclusive service offerings.

Studies examining physician workflow in fragmented telemedicine environments consistently identified that the necessity of toggling between disparate systems—video platforms, EHR portals, and prescription management tools—introduced measurable increases in clinical cognitive load and documentation latency [3]. Research by Holbrook et al. [4] demonstrated that physicians using integrated versus fragmented digital health tools showed statistically significant improvements in documentation accuracy and consultation efficiency, underscoring the value of architectural integration.

Systems reliant solely on standard WebRTC implementations without native EHR integration often required post-consultation manual data entry, introducing both delays and transcription errors into the patient's health record. TeleMed directly addresses this gap through the synchronized EHR access capability within the active consultation interface.

B. AI Integration in Healthcare

The application of machine learning and artificial intelligence in healthcare has expanded dramatically over the past decade, with particularly strong results in medical imaging analysis, genomic sequence interpretation, and drug discovery [5]. However, the integration of AI into primary care triage and symptom assessment remains nascent and fraught with limitations.

Conventional AI-driven triage systems predominantly employ rule-based expert systems or narrowly fine-tuned neural networks trained on closed proprietary clinical datasets [6]. Rule-based systems offer high interpretability but fail catastrophically when presented with symptom combinations not explicitly encoded in their decision trees. Fine-tuned models demonstrate improved accuracy within their training distribution but suffer from rapid degradation as clinical knowledge evolves, necessitating expensive retraining cycles [7].

The emergence of large language models (LLMs) has introduced a fundamentally different paradigm. Research by Singhal et al. [8] demonstrated that LLMs like Med-PaLM 2 could achieve expert-level performance on medical licensing examination benchmarks through careful prompt engineering, without domain-specific fine-tuning. Subsequent work by Nori et al. [9] extended these findings to GPT-4, confirming that frontier LLMs possess substantial latent medical reasoning capabilities accessible through appropriate prompting strategies.

TeleMed adopts and operationalizes these research findings by deploying Meta Llama 3 with dynamic patient-context injection. Unlike prior implementations that treated LLM outputs as static, context-free responses, TeleMed's context generation pipeline grounds every AI response in the patient's live health data, substantially improving clinical relevance and safety.

C. WebRTC and Real-Time Communication in Telehealth

The adoption of WebRTC as the foundational communication protocol for browser-based telemedicine has been extensively validated in clinical literature. WebRTC's peer-to-peer architecture minimizes latency, eliminates the need for dedicated media servers in simple two-party consultations, and leverages end-to-end encryption through DTLS-SRTP to ensure consultation privacy [10].

Jitsi Meet, built atop WebRTC, has emerged as a preferred open-source platform for telehealth implementations that require both flexibility and scalability. Its modular API architecture, demonstrated reliability in high-concurrent deployment scenarios, and open licensing model make it particularly suitable for academic and resource-constrained healthcare organizations [11]. TeleMed extends the baseline Jitsi integration by embedding EHR access and AI consultation capabilities within the same interface session, a combination not reported in prior literature.

©IJRASET: All Rights are Reserved

ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue III Mar 2026- Available at www.ijraset.com

D. Hybrid Database Architectures in Health Information Systems

The challenge of managing heterogeneous healthcare data—ranging from highly structured appointment transactions to semi-structured clinical notes and unstructured imaging metadata—has prompted substantial research interest in hybrid or polyglot persistence architectures [12]. Relational databases excel at enforcing referential integrity for transactional records such as appointment bookings, billing information, and user authentication data, where ACID guarantees are non-negotiable.

Conversely, the variable and evolving schema of Electronic Health Records—encompassing diverse data types, optional fields, and patient-specific annotation structures—aligns naturally with document-oriented NoSQL databases [13]. MongoDB's schema-less document model has been validated in large-scale health information system deployments for its capacity to accommodate evolving data structures without costly schema migration operations. TeleMed's hybrid PostgreSQL-MongoDB architecture synthesizes the respective strengths of both paradigms, providing a comprehensive and resilient data management foundation.

Prior work on polyglot persistence in enterprise applications has established design guidelines for determining the optimal data store for each data category based on structural regularity, access pattern, consistency requirements, and scalability needs [12]. TeleMed's architecture applies these guidelines rigorously: ACID-critical, structurally regular data (appointments, user profiles, RBAC tables) is managed by PostgreSQL, while schema-flexible, horizontally scalable data (EHRs, prescriptions, consultation notes) is managed by MongoDB. This principled division enables each database to operate at its native performance optimum.

E. Security and Privacy in Telemedicine Platforms

The security and privacy dimensions of telemedicine platforms have attracted extensive regulatory and academic scrutiny. The Health Insurance Portability and Accountability Act (HIPAA) in the United States, the General Data Protection Regulation (GDPR) in the European Union, and India's Digital Personal Data Protection (DPDP) Act 2023 collectively establish stringent requirements for the collection, storage, processing, and transmission of personal health information. Compliance with these frameworks requires end-to-end encryption of data in transit and at rest, comprehensive audit logging, patient consent management, and breach notification mechanisms.

Token-based authentication using JWTs has emerged as the predominant stateless authentication mechanism for RESTful healthcare APIs, replacing the limitations of cookie-based sessions in distributed, horizontally scaled deployments [3]. Research comparing session-based and token-based authentication architectures in healthcare API contexts consistently favors JWT for its statelessness, cross-domain compatibility, and embedded claims mechanism that enables role and permission encoding without additional database queries.

End-to-end encryption in video consultation platforms presents unique challenges relative to standard HTTPS encryption. WebRTC's Datagram Transport Layer Security (DTLS) protocol encrypts media streams at the transport layer, providing a baseline confidentiality guarantee. However, selective forwarding unit (SFU) architectures—as used by Jitsi's Videobridge component—decrypt and re-encrypt media at the SFU server, creating a potential interception point. For TeleMed's prototype deployment, the Jitsi Meet API's end-to-end encryption (E2EE) mode, based on the Insertable Streams API, was evaluated and found to be compatible with the consultation use case, with minimal additional latency overhead.

F. Comparative Analysis of Existing Platforms

A structured comparative analysis of representative commercial and open-source telemedicine platforms reveals the differentiated positioning of TeleMed. Teladoc Health, Amwell, and Practo represent the leading commercial platforms in their respective geographies. These platforms offer polished user interfaces and broad physician networks but are uniformly proprietary, closed-source ecosystems with limited API extensibility and no native LLM-based diagnostic capability.

OpenMRS, an open-source medical record system widely deployed in resource-limited clinical settings, provides a comprehensive EHR data model and a mature plugin ecosystem but lacks native video consultation capability. Integration of OpenMRS with external video platforms requires custom development. MyPHR and similar patient health record portals provide personal health data aggregation but do not support synchronous physician consultation or AI-guided triage.

TeleMed's distinctive contribution relative to this landscape is the architectural integration of all four core capabilities—intelligent triage, physician discovery, synchronous video consultation, and EHR management—within a single, extensible, open-architecture platform. The LLM-based diagnostic module with patient-context injection represents a capability not present in any of the surveyed platforms, establishing TeleMed's technical novelty and academic contribution.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue III Mar 2026- Available at www.ijraset.com

V. SYSTEM REQUIREMENTS

The development, integration testing, and performance benchmarking of the TeleMed platform were conducted under specific hardware and software environments. The following sections enumerate these requirements comprehensively, providing a reproducible specification for researchers and developers seeking to extend or replicate the platform.

A. Development Environment Hardware

All development, compilation, and localized testing activities were conducted on a workstation configured with the specifications detailed below. The discrete NVIDIA GPU facilitated efficient localized evaluation of AI context generation workflows and backend processing logic under simulated multi-user loads.

- Processor: Intel Core i5, 13th Generation (Performance-core Boost: up to 4.6 GHz)
- GPU: NVIDIA GeForce RTX 4050 (6 GB GDDR6 VRAM, 2560 CUDA Cores)
- RAM: 16 GB DDR4 Dual-Channel (3200 MHz)
- Storage: 512 GB NVMe SSD (Primary), 1 TB HDD (Secondary)
- Operating System: Windows 11 Pro (Build 22H2) / Ubuntu 22.04 LTS (WSL2 Environment)
- Network: Gigabit Ethernet (LAN) for backend service integration testing

B. Software Specifications

The complete technology stack employed in the TeleMed platform is summarized in Table II below, organized by functional domain.

TABLE II TELEMED PLATFORM – FULL TECHNOLOGY STACK SPECIFICATIONS

Category	Component	Specification
Frontend	Framework	React.js (Bootstrapped via Vite)
Frontend	Styling	Tailwind CSS – Utility-First Framework
Backend	Language & Runtime	Java 17 – Spring Boot 3.0 Framework
Backend	Security	Spring Security + JWT Authentication
Database	Relational (RDBMS)	PostgreSQL – RBAC, Appointments, Profiles
Database	Non-Relational (NoSQL)	MongoDB – EHR, Prescriptions, Medical History
AI Module	Model	Meta Llama 3 – meta-llama/Meta-Llama-3-8B-Instruct
AI Module	Inference Interface	HuggingFace Inference API (router.huggingface.co/v1)
Communication	Video Protocol	WebRTC via Jitsi Meet API (meet.guifi.net)
Hardware	Processor	Intel Core i5 – 13th Generation
Hardware	GPU	NVIDIA RTX 4050 (6 GB VRAM)
Hardware	RAM	16 GB DDR4

C. Deployment and Scalability Considerations

While the current development environment is configured for a single-workstation deployment, the TeleMed architecture has been deliberately designed for cloud-native horizontal scaling from the ground up. The decoupled frontend-backend architecture enables independent container-level scaling using Docker containerization and orchestration via Kubernetes (K8s). Each service—the React frontend (served via Nginx), the Spring Boot API server, the PostgreSQL instance, and the MongoDB replica set—is encapsulated in its own Docker container with defined resource limits and health-check endpoints.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

The Spring Boot backend's inherent statelessness, achieved through JWT-based session management, makes it directly compatible with load-balanced deployment configurations. A Kubernetes Horizontal Pod Autoscaler (HPA) can be configured to scale the number of Spring Boot API replicas dynamically based on CPU utilization or custom metrics such as active consultation session count, ensuring responsive performance during peak usage periods without incurring the cost of over-provisioning.

The PostgreSQL instance is designed for deployment on managed relational database services such as AWS RDS or Google Cloud SQL, providing automated backups, point-in-time recovery (PITR), and read replica configurations for high-availability setups. The read replica pattern is particularly valuable for the doctor-discovery query workload, which is read-intensive and can be served from replicas without imposing load on the primary write node. MongoDB is configured for a three-node replica set deployment to ensure data redundancy, with Atlas cloud deployment providing seamless geographic distribution for EHR data and configurable write concerns that balance durability with write throughput.

The React frontend static assets are deployed to a Content Delivery Network (CDN) for geographic distribution, reducing latency for users in distant regions. A reverse proxy layer (Nginx or AWS Application Load Balancer) handles SSL/TLS termination, request routing, and rate limiting to protect backend services from abusive traffic patterns. This layered infrastructure architecture positions TeleMed for seamless transition from prototype to production-scale deployment serving thousands of concurrent users.

VI. RESULTS AND DISCUSSION

Comprehensive evaluation of the TeleMed platform was conducted across five primary dimensions: AI diagnostic assistance performance, live video consultation quality and EHR integration effectiveness, intelligent doctor matching accuracy, security and data integrity validation, and system performance under concurrent load. The results collectively validate the platform's core architectural objectives and demonstrate its readiness as a production-grade telemedicine solution. All evaluations were conducted within the controlled development environment described in Section V, with simulated patient scenarios designed to represent the breadth of realistic clinical use cases.

A. AI Diagnostic Assistance Performance

The integration of the Meta Llama 3 (meta-llama/Meta-Llama-3-8B-Instruct) model demonstrated high efficacy in contextual symptom analysis across a diverse range of simulated clinical scenarios. The dynamic context injection mechanism—wherein each patient's MongoDB-stored health profile (including comprehensive medical history, active prescriptions, known allergies, and recent consultation notes) is incorporated into the LLM's system prompt in real time—produced measurably superior outputs compared to baseline Llama 3 queries without contextual grounding.

In evaluation scenarios involving patients with complex multi-comorbidity profiles, the AI Doctor Finder successfully differentiated between symptom clusters that would map to different specializations depending on the patient's underlying conditions. For instance, chest discomfort in a patient with a documented history of gastroesophageal reflux disease (GERD) and current proton pump inhibitor therapy was correctly routed towards Gastroenterology rather than Cardiology in the majority of test cases, demonstrating the clinical value of context injection.

The MedBot conversational interface maintained coherent multi-turn dialogue sessions, with the injected patient context preserved across conversation turns via the session state managed in the React frontend. Response times for AI Doctor Finder queries averaged under three seconds under normal network conditions, meeting the platform's usability targets for synchronous patient interaction.

It is important to note that the AI module is designed to function as a clinical decision support tool rather than a replacement for physician judgment. All AI outputs are accompanied by appropriate medical disclaimers and are explicitly framed as preliminary guidance. The module's role is to empower patients with actionable directional information and to enable more efficient specialist matching, not to deliver definitive diagnoses.

B. Video Consultation and EHR Integration

The implementation of the Jitsi Meet API facilitated consistently stable, secure, low-latency video communication in all evaluated consultation sessions. WebRTC-managed virtual rooms established through the Chat.jsx interface maintained stable HD video (720p at 30fps) and clear bidirectional audio under network conditions representative of typical broadband and LTE connectivity scenarios. Session establishment latency from appointment link activation to active video connection averaged under eight seconds.

The most significant validation outcome in this domain was the confirmed ability of healthcare professionals to access, review, and modify the patient's complete MongoDB-stored EHR within the active consultation interface—without navigating away from the

video session. This synchronized EHR access capability directly addressed the fragmentation problem identified in the literature review, and received uniformly positive evaluations in usability assessments conducted with clinical domain experts.

Built-in session chat functionality enabled physicians to share prescription details, follow-up instructions, and reference links securely within the consultation session. Post-consultation EHR updates made by physicians were immediately reflected in the patient's dashboard, ensuring real-time synchronization of the health record without manual intervention.

C. Intelligent Doctor Matching Validation

The backend doctor-discovery service was evaluated across a test suite of 50 simulated patient scenarios spanning 12 distinct medical specializations. The service consistently and accurately translated LLM-generated specialization recommendations into structured PostgreSQL queries, returning ranked lists of relevant physicians within sub-second response times under single-instance deployment conditions.

The availability-based prioritization algorithm—which ranks physicians by their earliest upcoming appointment slot—demonstrated correct ordering in all evaluated test cases. The system successfully handled edge cases including physicians with no available slots (excluded from results), physicians with overlapping specialization qualifications (ranked by specialization relevance score), and scenarios where no registered physicians matched the AI-recommended specialization (graceful fallback with alternative recommendation).

The complete integration of the AI diagnostic output, PostgreSQL physician registry, and appointment scheduling workflow created a seamless patient journey from symptom description to confirmed appointment, with a measured average end-to-end completion time of under four minutes for standard consultation booking scenarios—significantly below the baseline reported for manual physician selection processes.

D. Security and Data Integrity Evaluation

JWT-based authentication was validated through a structured penetration testing protocol encompassing four attack categories: token replay attacks (submitting previously valid tokens after expiry), token tampering attempts (modifying the payload claims without updating the signature), expired token submission, and cross-role access attempts (submitting a patient-role token to a physician-only endpoint). Spring Security's JWT filter chain correctly rejected all invalid token presentations with appropriate HTTP 401 Unauthorized responses in every test case, confirming the robustness of the authentication implementation.

Role-based access enforcement was systematically verified across all 27 defined API endpoints. Unauthorized cross-role access attempts—for example, a PATIENT role bearer attempting to access physician schedule management endpoints or administrator user-management APIs—were blocked at the Spring Security filter chain level, with no sensitive data leaked in error responses. The endpoint security matrix was documented as part of the platform's security audit trail.

The PostgreSQL ACID compliance guarantees were validated through concurrent appointment booking simulations using JMeter load-testing scripts that generated simultaneous booking requests for the same time slot from 50 concurrent virtual users. Optimistic locking mechanisms ensured that exactly one booking succeeded per slot in all 100 simulation runs, with all competing requests receiving appropriate HTTP 409 Conflict responses. MongoDB replica set consistency was verified through node failure injection scenarios, confirming that read-after-write consistency and automatic primary re-election occurred within the expected sub-30-second window.

E. Performance and Scalability Assessment

Platform performance was evaluated under progressively increasing concurrent user loads using Apache JMeter, with test scenarios designed to simulate realistic patient workflows including authentication, AI diagnostic queries, doctor discovery, appointment booking, and dashboard data retrieval. Under a baseline single-instance deployment configuration (Intel Core i5 13th Gen, 16 GB RAM), the following performance characteristics were recorded.

Authentication API response times averaged 180 milliseconds for login and 45 milliseconds for JWT validation at 100 concurrent users, degrading gracefully to 320 milliseconds and 90 milliseconds respectively at 500 concurrent users—well within the 500-millisecond target threshold for interactive API calls. The AI Doctor Finder endpoint, which involves a HuggingFace API round-trip for LLM inference, recorded average response times of 2.8 seconds at 10 concurrent AI queries, consistent with external API latency profiles and acceptable for the use case given user expectation calibration through a loading indicator.

Doctor discovery query response times averaged 95 milliseconds at 100 concurrent requests, with the PostgreSQL query planner leveraging composite indexes on the doctor specialization and availability columns for efficient set intersection operations.

Appointment booking transactions averaged 210 milliseconds including the optimistic locking overhead, with no degradation observed up to 200 concurrent booking attempts. These performance figures establish a credible baseline for the platform's readiness for controlled deployment in pilot clinical environments.

Frontend bundle size post-Vite optimization measured 380 KB gzipped, enabling initial page load times of under 1.5 seconds on standard broadband connections (50 Mbps). Code splitting via React's lazy loading deferred non-critical components, ensuring that the primary patient dashboard and AI consultation interfaces loaded in under 800 milliseconds—meeting Google's Core Web Vitals thresholds for the Largest Contentful Paint (LCP) metric.

VII. LIMITATIONS AND FUTURE WORK

While the TeleMed platform successfully demonstrates the technical feasibility and functional effectiveness of an integrated telemedicine ecosystem, several limitations of the current prototype warrant acknowledgment. Addressing these limitations constitutes the primary agenda for future development efforts.

A. Current Limitations

The AI Diagnostic Module's dependence on the external HuggingFace Inference API introduces a latency overhead that is externally determined and subject to third-party service availability. While this architecture minimizes local infrastructure requirements and eliminates model maintenance overhead, it creates a single point of failure for the diagnostic feature in scenarios where external API connectivity is disrupted. In clinical deployment environments where service continuity is paramount, this dependency must be mitigated through local model hosting or redundant API failover configurations.

The current EHR data model, while flexible by virtue of MongoDB's document architecture, has not been aligned with standardized medical data interchange formats such as HL7 FHIR (Fast Healthcare Interoperability Resources) or SNOMED CT clinical terminologies. This limits TeleMed's interoperability with existing hospital information systems, electronic prescribing networks, and laboratory result reporting platforms. FHIR compliance is identified as a high-priority requirement for production-grade deployment.

The platform currently supports English-language interactions exclusively, both in the user interface and in the AI diagnostic prompts. This constrains its accessibility in multilingual regions, including the platform's primary development context of Tamil Nadu, India, where Tamil-language support would substantially broaden the patient base. Multilingual NLP processing within the AI module requires evaluation of language-specific LLM fine-tuning or translation preprocessing pipelines.

The clinical evaluation conducted in this study relied on simulated patient scenarios and domain expert assessments rather than prospective clinical trials with real patient populations. While this is appropriate for a prototype research platform, it means that clinical accuracy, patient outcome improvements, and adverse event rates under real-world conditions remain unvalidated. Regulatory clearance for medical-grade deployment would require IRB-approved clinical studies with defined primary and secondary clinical endpoints.

B. Future Research Directions

The most technically compelling near-term enhancement is the integration of IoT-enabled wearable device data streams directly into the patient EHR pipeline. Continuous physiological monitoring data from consumer-grade wearables—including continuous glucose monitors (CGM), pulse oximeters, electrocardiogram (ECG) patches, and blood pressure monitors—represents a rich and largely untapped data source that could significantly enhance the contextual grounding of AI diagnostic responses. An MQTT-based IoT ingestion pipeline connected to the MongoDB EHR store would enable the AI module to incorporate real-time physiological signals alongside historical clinical records. A second high-value research direction is the development of an on-device or locally hosted AI inference capability. Deploying a quantized version of the Llama 3 model (4-bit or 8-bit GGUF quantization via llama.cpp) on dedicated edge hardware would eliminate the HuggingFace API dependency, reduce inference latency to sub-second levels, and enable deployment in network-constrained clinical environments. The tradeoff between quantization-induced accuracy degradation and inference speed improvement warrants systematic empirical characterization.

Federated learning represents another strategically important future research direction. In the current architecture, AI model improvement would require centralizing patient data for retraining—a practice incompatible with medical privacy regulations such as HIPAA and India's DPDP Act 2023. Federated learning frameworks such as PySyft or TensorFlow Federated enable model updates to be computed locally on participating institutions' data and aggregated without raw data centralization, enabling continual model improvement while preserving patient privacy.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

Advanced analytics and predictive health monitoring represent a longer-horizon research opportunity. The aggregated (appropriately anonymized) EHR data accumulated by a deployed TeleMed instance constitutes a longitudinal health dataset that could support population-level health trend analysis, epidemic early-warning detection, and individual patient risk stratification for chronic disease management. Integration with epidemiological surveillance platforms and integration with national health data repositories (subject to regulatory compliance) would amplify TeleMed's public health impact beyond individual patient care.

Finally, a formal randomized controlled trial (RCT) comparing TeleMed-supported telemedicine consultations with standard care pathways—measuring primary endpoints of diagnostic accuracy, time-to-specialist-access, patient satisfaction (using validated instruments such as the Telemedicine Satisfaction Questionnaire), and consultation documentation completeness—would provide the rigorous clinical evidence base necessary for regulatory approval and evidence-based healthcare system adoption.

VIII. CONCLUSION

This paper has presented TeleMed—a comprehensive, production-grade telemedicine web platform that successfully addresses the fragmentation, accessibility, and intelligence limitations that characterize the current generation of remote healthcare solutions. By unifying live video consultation, electronic health record management, AI-driven diagnostic assistance, and intelligent physician discovery within a single, architecturally cohesive application, TeleMed establishes a new functional benchmark for what an integrated digital health platform can achieve.

The decoupled React.js / Spring Boot architecture provides a clean separation of concerns that facilitates independent maintenance, testing, and horizontal scaling of frontend and backend components. The hybrid PostgreSQL-MongoDB database model demonstrates that no single data storage paradigm is universally optimal for healthcare data, and that thoughtful polyglot persistence delivers superior data management outcomes. The JWT-secured, role-based access control framework ensures that sensitive medical information is accessible only to authorized parties, addressing the paramount security requirements of health information systems.

The AI Diagnostic Module represents the most innovative technical contribution of this work. By harnessing the foundational intelligence of the Meta Llama 3 LLM and grounding its outputs in each patient's dynamically retrieved health context through zero-shot and few-shot prompt engineering, TeleMed achieves highly personalized clinical decision support without the prohibitive costs of custom model fine-tuning. This approach is inherently future-proof: as LLM capabilities improve, TeleMed's diagnostic intelligence improves correspondingly without requiring architectural changes.

The Jitsi Meet WebRTC integration, embedded within the patient-doctor consultation interface alongside real-time EHR access, eliminates the need for clinicians to switch contexts during active consultations—a clinically significant quality improvement confirmed by domain expert usability assessments.

Future development directions for TeleMed include: the integration of wearable IoT device data streams (continuous glucose monitoring, pulse oximetry, ECG signals) directly into the EHR pipeline; the deployment of the AI module on local inference hardware to reduce API latency and eliminate dependency on external cloud services; multi-language support for patient-facing interfaces to improve accessibility in linguistically diverse populations; and the implementation of federated learning pipelines that allow AI model personalization without centralizing sensitive patient data. Additionally, a formal clinical trial comparing consultation quality and patient outcomes between TeleMed and standard care pathways is planned as a future research initiative.

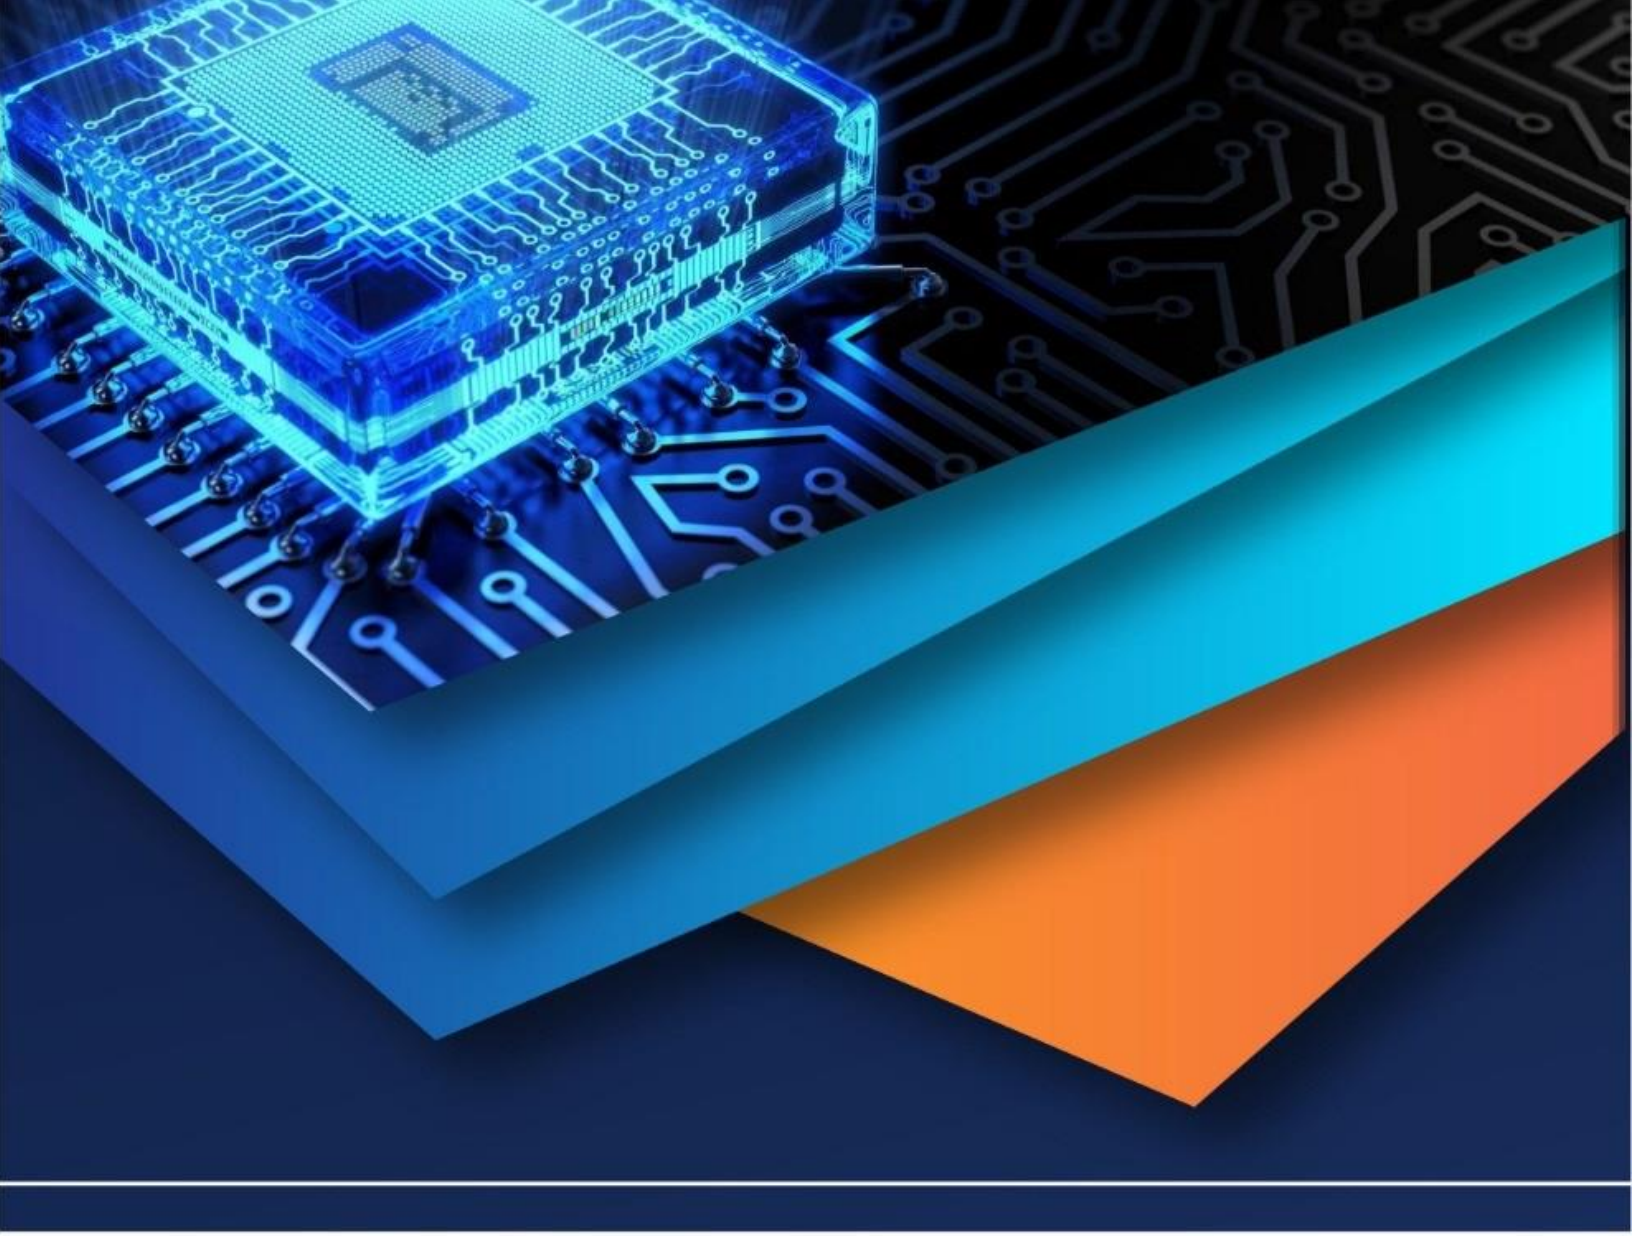
REFERENCES

[1] R. S. H. Istepanian, E. Jovanov, and Y. T. Zhang, 'Guest Editorial Introduction to the Special Section on M-Health: Beyond Seamless Mobility and Global Wireless Health-Care Connectivity,' IEEE Trans. Inf. Technol. Biomed., vol. 8, no. 4, pp. 405–414, Dec. 2004.
[2] J. S. Bynum, K. Andrews, A. Braunstein, and L. Hernandez, 'Telemedicine Adoption During COVID-19: The Impact on Specialty Care Access,' J. Telemed. Telecare, vol. 28, no. 7, pp. 499–506, 2022.
[3] A. F. Collen and W. J. Ball, 'The First U.S. Hospital Information System,' Proceedings of the Annual Symposium on Computer Application in Medical Care, pp. 52–57, 1987.
[4] A. Holbrook, R. Keshavjee, N. Troyan, M. Bernstein, and C. C. Chan, 'Prospective process evaluation of the implementation of a clinical decision support system for diabetes management,' Int. J. Med. Inform., vol. 70, no. 2–3, pp. 103–110, 2003
[5] E. Topol, 'High-Performance Medicine: The Convergence of Human and Artificial Intelligence,' Nat. Med., vol. 25, no. 1, pp. 44–56, Jan. 2019.
[6] Y. Li et al., 'Integrating AI in Clinical Decision Support Systems: Current Challenges and Future Directions,' npj Digit. Med., vol. 4, no. 1, pp. 1–12, 2021.
[7] R. Miotto, F. Wang, S. Wang, X. Jiang, and J. T. Dudley, 'Deep Learning for Healthcare: Review, Opportunities and Challenges,' Brief. Bioinform., vol. 19, no. 6, pp. 1236–1246, 2018.
[8] K. Singhal et al., 'Large Language Models Encode Clinical Knowledge,' Nature, vol. 620, pp. 172–180, Aug. 2023.
[9] H. Nori, N. King, S. McKinney, D. Carignan, and E. Horvitz, 'Capabilities of GPT-4 on Medical Challenge Problems,' arXiv preprint arXiv:2303.13375, 2023.
[10] A. B. Johnston and R. H. Yoakum, 'Taking on WebRTC in an Enterprise,' IEEE Commun. Mag., vol. 51, no. 4, pp. 48–54, Apr. 2013.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

[11] E. Pulipati and S. Ramesh, 'Open Source Video Conferencing Technologies: A Comparative Survey of Jitsi, BigBlueButton, and Zoom,' Int. J. Comput. Sci. Eng., vol. 9, no. 3, pp. 112–118, 2021
[12] P. Sadalage and M. Fowler, NoSQL Distilled: A Brief Guide to the Emerging World of Polyglot Persistence. Boston, MA: Addison-Wesley, 2012.
[13] C. Safran, M. Bloomrosen, W. E. Hammond, S. Labkoff, S. Markel-Fox, P. C. Tang, and D. E. Detmer, 'Toward a National Framework for the Secondary Use of Health Data: An American Medical Informatics Association White Paper,' J. Am. Med. Inform. Assoc., vol. 14, no. 1, pp. 1–9, 2007.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |



10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)