



International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347 -6435 (Online)), Volume 15, Issue 5, May 2026)

SmartMeds: An AI-Powered Intelligent Healthcare Assistance and Specialist Recommendation Platform

Prof. Harshali Chaudhari¹, Rohit Dandge², Chinmay Patil³, Raj Patole⁴, Sumeet Thorat⁵

¹Assistant Professor, Department of Information Technology, International Institute of Information Technology (I2IT), Pune, Maharashtra, India.

²Department of Information Technology, International Institute of Information Technology (I2IT), Pune, Maharashtra, India.

Abstract— Across the global healthcare landscape, the fragmentation of patient medical records among geographically dispersed providers constitutes a persistent barrier to efficient, high-quality clinical care. SmartMeds is an AI-powered, monetized, and appointment-regulated digital health platform engineered to consolidate this fragmentation into a cohesive, secure, and clinically actionable ecosystem. The system unifies centralized Electronic Health Record (EHR) storage with a multi-stage verified doctor registration framework, a structured appointment-and-payment workflow, and an appointment-conditioned real-time secure chat channel. An OCR-enabled, BioBERT/ClinicalBERT-driven Recommendation Engine processes patient-uploaded medical reports to generate ranked specialist referrals accompanied by calibrated confidence scores. Supplementing this capability, a GPT-integrated Smart Healthcare Chatbot delivers round-the-clock assistance for platform navigation, scheduling, and payment queries. Platform-wide security is realized through AES-256 document encryption, RS256 JWT session management, role-based access control (RBAC), PostgreSQL row-level security policies, and PCI DSS-compliant payment processing. Collectively, SmartMeds constitutes a production-ready, startup-grade architecture that substantively narrows the divide between siloed clinical data repositories and patient-centric, intelligently assisted digital health delivery.

Keywords— AI Doctor Recommendation, Appointment-Based Consultation, Centralized Health Records, Clinical Summarization, Doctor Verification, Electronic Health Records, Healthcare Chatbot, Healthcare Payments, JWT Authentication, NLP, Secure Digital Healthcare, Telemedicine.

I. INTRODUCTION

Contemporary healthcare delivery is impeded by a deeply entrenched structural problem: patient medical information is dispersed across a heterogeneous array of hospitals, specialist clinics, and independent practitioners without any shared access mechanism. The consequences of this fragmentation are clinically significant—redundant diagnostic workups, preventable adverse medication interactions, and delays in diagnosis attributable to the unavailability of prior investigation results. Physicians

engaged in referral-based care are routinely denied access to the longitudinal records indispensable to sound clinical judgement, while patients shoulder the burden of manually assembling and presenting their own medical histories at every new encounter. The systemic absence of a unified, access-controlled, and intelligent health data infrastructure has measurable consequences for patient safety and healthcare resource utilization worldwide.

SmartMeds is conceived as a direct response to this structural deficit. Designed to the standards of a deployable healthcare startup, it consolidates Electronic Health Records (EHRs) from multiple providers into a single, access-controlled repository, enforces a rigorous physician credential verification process before granting any patient-facing access, and channels all patient-physician interactions through a regulated appointment-and-payment pipeline. The platform extends materially beyond conventional patient portals by embedding a monetized consultation workflow, integrating a BioBERT-driven specialist recommendation engine that mines OCR-extracted clinical reports, deploying an appointment-gated real-time messaging channel, and providing a GPT-powered conversational assistant for platform navigation—collectively constituting an end-to-end digital health ecosystem rather than a simple record repository.

The platform's contributions resolve around four axes of innovation. The first is a multi-layer Doctor Verification System that sequentially validates medical licences, academic qualifications, and hospital affiliations prior to any platform access being granted. The second is an Appointment-Based Consultation Architecture that subordinates every patient-doctor interaction to a structured scheduling and payment sequence, ensuring accountability and billing integrity. The third is an AI Recommendation Engine that combines OCR-based report digitization with domain-adapted clinical NLP to produce symptom-to-specialist mappings quantified by confidence scores. The



International Journal of Recent Development in Engineering and Technology

Website: www.ijrdet.com (ISSN 2347 -6435 (Online)), Volume 15, Issue 5, May 2026)

fourth is a secure chat facility whose activation is contingent on a confirmed booking, a completed payment, and explicit physician approval—eliminating unregulated clinical communication and its associated liability. This paper documents SmartMeds' architecture, methodological design, implementation specifics, and security provisions, establishing its credentials as a viable, real-world-deployable healthcare platform.

II. LITERATURE SURVEY

Blobel et al. [1] underscored the imperative for standardized, interoperable EHR frameworks capable of supporting cross-institutional data exchange, pinpointing authentication mechanisms and access control policies as the principal technical barriers to realizing such interoperability. Their analysis provided the conceptual grounding for centralized, permission-governed record architectures—a principle that SmartMeds operationalizes through its combination of RBAC middleware and JWT-based stateless session management.

Rajpurkar et al. [2] furnished empirical evidence that deep learning architectures applied to medical imaging corpora and clinical free text can attain diagnostic parity with specialist clinicians. The CheXNet model described in their study provided an architectural reference point for SmartMeds' AI Recommendation Engine, which leverages fine-tuned BERT-family models to analyze clinical text recovered by Tesseract OCR and AWS Textract from patient-uploaded diagnostic reports, mapping recognized clinical entities to appropriate specialist categories.

Lee et al. [3] rigorously evaluated the application of BioBERT and ClinicalBERT to clinical text summarization and classification, demonstrating that transformer models pre-trained on biomedical corpora attain substantially higher classification fidelity on biomedical tasks than their general-domain counterparts. In response to these findings, SmartMeds' recommendation pipeline adopts ClinicalBERT as its classification backbone, further adapted through fine-tuning on medical ontology-derived datasets encompassing SNOMED CT and ICD-10 terminological codes.

Saripalle et al. [4] examined HL7 FHIR as a healthcare interoperability standard, concluding that RESTful API paradigms serve as practical enabling mechanisms for modular, independently deployable healthcare application components. SmartMeds' microservices decomposition reflects this FHIR-informed design philosophy, with modular Flask-based AI services interfacing with the central Node.js/Express API gateway through formally versioned RESTful contracts.

Kruse et al. [5] conducted a systematic survey of obstacles to telemedicine adoption, identifying three recurring deterrents of direct relevance to SmartMeds' design: the complexity of integrating consultation payments, the prevalence of unverified practitioner directories, and the inadequacy of structured communication channels. SmartMeds addresses each deterrent in turn through its Razorpay/Stripe dual-gateway payment integration, its admin-reviewed doctor credential pipeline, and its Socket.IO-backed appointment-gated messaging architecture.

Abouelmehdi et al. [6] synthesized existing research on security architectures for large-scale healthcare data systems, advocating for AES-256 encryption as the minimum acceptable standard for health data at rest and TLS 1.3 as the baseline requirement for data in transit. SmartMeds incorporates both recommendations as non-negotiable security baselines, supplementing them with PostgreSQL row-level security (RLS) policies and AWS S3 server-side encryption applied to all stored medical and credential documents.

A review of the incumbent landscape—including Epic MyChart, HealthVault, and Practo—reveals that these platforms address health record access, appointment booking, or provider discovery in isolation, without integrating AI-based specialist recommendation, end-to-end monetized consultation workflows, and verified physician credential management within a single coherent ecosystem. This gap constitutes the primary motivating context for SmartMeds' development.

III. PROPOSED SYSTEM

A. System Overview

SmartMeds is realised as a full-stack, microservices-oriented healthcare platform whose constituent layers span a React.js presentation frontend, a Node.js/Express API gateway, Flask-hosted AI/ML microservices, and a PostgreSQL relational data store. Redis underpins session caching, distributed pub/sub event propagation, and real-time infrastructure. Three user roles—Patients, Verified Doctors, and Platform Administrators—are granted differentiated, RBAC-enforced access boundaries governing every data and feature interaction within the system.

Platform functionality is decomposed into six interdependent subsystems: (1) Centralized EHR Management; (2) Doctor Verification and Registry; (3) Appointment and Scheduling Engine; (4) Consultation Payment Gateway; (5) Appointment-Gated Secure Chat; and



International Journal of Recent Development in Engineering and Technology

Website: www.ijrdet.com (ISSN 2347 -6435 (Online)), Volume 15, Issue 5, May 2026)

(6) AI-Powered Recommendation and Chatbot Services. Inter-subsystem communication is mediated by a unified REST API layer whose every endpoint is guarded by JWT authentication and HTTPS/TLS 1.3 transport security.

B. Doctor Verification System

The Doctor Verification System is a sequential credential validation pipeline designed to guarantee that every practitioner granted access to the SmartMeds doctor portal holds valid, current medical qualifications. At the point of registration, physicians are required to submit their medical licence number, degree certificates, a government-issued identification document, and a letter of hospital affiliation through a Multer-powered secure upload interface. Submitted documents are persisted in AWS S3 or Cloudinary with AES-256 server-side encryption, with object references catalogued in the `doctor_verification` table of the PostgreSQL schema.

Pending verification applications surface within the Administrator Verification Dashboard, which presents document previews and allows administrators to approve, reject, or solicit additional supporting documentation. Upon approval, the physician's account is promoted to verified status, with a verification badge stored in the `doctors` table and displayed on their patient-facing profile. A complete, tamper-evident audit record of every verification action—including decision, actor identity, and timestamp—is maintained in the `admin_audit_logs` table. Rejected applicants receive automated, reason-specific email notifications dispatched via Nodemailer, enabling guided resubmission.

C. Appointment-Based Connectivity

All patient-physician interactions in SmartMeds are mandatorily mediated through the appointment pipeline, which prevents unregulated direct contact and ensures every clinical exchange is documented and accountable. Patients locate verified specialists using the doctor directory, filterable by specialisation, geographic proximity, consultation fee, aggregate rating, and appointment availability. Displayed profiles consolidate verification badges, academic credentials, institutional affiliations, and post-consultation patient ratings sourced from the `ratings_reviews` table.

Available time slots, defined by physicians through a FullCalendar-powered scheduling interface and stored in the `doctor_schedule` table, are presented to patients for selection. Slot selection initiates an appointment record in a pending state within the `appointments` table. Following physician

approval or rescheduling via the doctor dashboard, automated confirmation notifications are dispatched—email via Nodemailer and SMS via Twilio—at 24-hour and 1-hour pre-consultation intervals. Underlying this reminder workflow, Redis-managed job queues provide fault-tolerant scheduling with configurable retry behaviour.

D. Online Payment System

Consultation fees in SmartMeds are collected as a mandatory precondition for appointment confirmation: no appointment transitions to confirmed status without a completed payment transaction. Dual payment processing support is provided through integrations with both Razorpay and Stripe, collectively accommodating UPI, debit and credit cards, net banking, and digital wallet instruments. When a patient initiates payment, the server generates a cryptographically signed order identifier via the respective gateway API; this identifier and associated transaction metadata are recorded in the `payments` table.

Gateway-issued webhook events trigger atomic PostgreSQL updates to appointment status following payment confirmation, eliminating race conditions that could result in inconsistent state. PDFKit-generated consultation invoices—embedding appointment details, payment references, and practitioner information—are both emailed to patients and made accessible through the patient portal. Automated refund workflows handle appointment cancellations within platform-configurable policy windows. Administrators access revenue analytics dashboards aggregating transaction volumes, specialty-level revenue distribution, and refund rates from the `invoices` and `payments` tables.

E. Appointment-Gated Secure Chat

Access to the SmartMeds secure messaging channel is protected by a three-condition gate: the existence of a confirmed appointment record, verification of completed payment, and explicit physician approval of the consultation session. Only when all three predicates evaluate true does the system issue a JWT-signed session token whose embedded expiry timestamp is synchronised with the scheduled appointment duration. Socket.IO manages bidirectional real-time messaging within the authenticated, appointment-scoped session namespace, with all payload transmissions protected by TLS 1.3 in transit.

Message records are durably persisted in the `messages` table with relational references to the `chat_sessions` table, supporting post-consultation review. Physicians may share structured PDF prescriptions as attachments within the chat



International Journal of Recent Development in Engineering and Technology

Website: www.ijrdet.com (ISSN 2347 -6435 (Online)), Volume 15, Issue 5, May 2026)

interface; these documents are stored in AES-256-encrypted form in AWS S3. Session termination is enforced server-side upon appointment conclusion, after which the channel reverts to read-only mode. Initiating a follow-up consultation necessitates a fresh appointment-and-payment cycle, preserving the integrity of the access-control regime. Redis pub/sub infrastructure enables horizontal scaling of Socket.IO instances across multi-node deployment configurations.

IV. METHODOLOGY

A. AI Recommendation Engine

The AI Recommendation Engine is deployed as a standalone Flask microservice that receives patient-uploaded medical reports as input and returns ordered specialist referral recommendations. Processing is organised into four successive stages: document ingestion, textual content extraction, clinical entity recognition, and specialist domain mapping.

During document ingestion, uploaded PDFs and image files undergo pre-processing before being routed—based on document complexity—to either Tesseract OCR for routine printed materials or AWS Textract for multi-page clinical reports containing tabular laboratory data. Raw extracted text is subjected to normalization, noise filtering, and sentence boundary detection before entering the NLP stage. SciSpacy's named entity recognition (NER) pipeline identifies discrete clinical concepts—diagnoses, anatomical structures, drug names, and laboratory measurements—from the unstructured extracted text.

The resulting entity set is passed to a ClinicalBERT model pre-trained on MIMIC-III clinical notes and subsequently fine-tuned on a disease-to-specialty mapping corpus derived from ICD-10 and SNOMED CT ontologies. The model outputs a prioritised ranking of applicable medical specialties, each accompanied by a calibrated confidence score. The engine then queries the verified doctor registry to identify available, highly rated specialists within each recommended domain, ordering results by a composite function incorporating model confidence, physician rating, geographic proximity, and appointment availability. Structured JSON responses convey these ranked referrals to the frontend specialist selection interface. All recommendation sessions are logged in the recommendations table to support downstream quality assurance and model retraining workflows.

B. Smart Healthcare Chatbot

The Smart Healthcare Chatbot is implemented as a Flask-hosted conversational microservice integrated with the GPT API to enable natural language understanding and response generation. Its functional scope is intentionally bounded to platform-related assistance categories—appointment scheduling support, physician fee and availability enquiries, payment process guidance, prescription retrieval, platform feature navigation, and reminder management—with clinical diagnosis and therapeutic advice explicitly excluded to maintain regulatory compliance.

Incoming user queries traverse an NLP-based intent classification pipeline that routes recognised intents to purpose-built handler functions. Appointment-related handlers perform parameterized reads against the appointments and doctor_schedule tables, while payment-oriented handlers retrieve data from the payments and invoices tables. Every chatbot interaction—including session transcripts, classified intents, and resolution outcomes—is recorded in the chatbot_logs table to facilitate ongoing performance monitoring, intent classification accuracy analysis, and iterative user experience improvement. Queries that fall outside the defined intent taxonomy are gracefully redirected to human platform support agents.

C. EHR Management and Clinical Summarization

The EHR module provides patients with a structured upload interface supporting prescriptions, laboratory investigation reports, radiological imaging summaries, and discharge documentation. Each uploaded file is categorised by document type, assigned a clinical date range, and stored in AWS S3 under AES-256 encryption. Document metadata—including access permissions—is governed by PostgreSQL row-level security policies that restrict read access to the record owner and any treating physician with an active appointment relationship.

Clinical summarization is executed by a fine-tuned BART/T5 model that synthesises structured clinical summaries from longitudinal EHR data. Generated summaries foreground active diagnoses, current medications, recent investigations, and recorded allergies in a standardised, clinician-readable format. Summary outputs are integrated into the patient timeline view and may, subject to patient consent, be disclosed to physicians ahead of scheduled appointments.

V. SYSTEM ARCHITECTURE

A. Database Schema



International Journal of Recent Development in Engineering and Technology

Website: www.ijrdet.com (ISSN 2347 -6435 (Online)), Volume 15, Issue 5, May 2026)

The SmartMeds relational data model is realised in PostgreSQL and encompasses thirteen core tables. The users table stores patient profiles, hashed authentication credentials, and role assignments. The doctors table maintains verified practitioner profiles including specialisation, consultation fees, and aggregated ratings. The doctor_verification table holds credential document references, current verification status, and administrator review records. The appointments table records booking lifecycle states and time-slot references. The payments table captures transaction records, gateway identifiers, and refund status. The chat_sessions table stores appointment-linked JWT session tokens with expiry timestamps and gate status. The messages table holds encrypted chat payload references, sender and receiver mappings, and timestamps. The doctor_schedule table encodes weekly availability templates, exception dates, and slot durations. The chatbot_logs table retains session transcripts and intent classifications. The recommendations table records OCR report references and model outputs. The invoices table stores payment-linked PDF references and itemized billing details. The ratings_reviews table contains post-consultation feedback and moderation status. The admin_audit_logs table records verification decisions, administrative actions, and actor identities with timestamps.

Referential integrity across all relational paths is enforced through foreign key constraints. PostgreSQL row-level security policies independently restrict data visibility at the database engine layer, providing a second line of defence against privilege-escalation vulnerabilities that might otherwise bypass application-layer access controls. Schema migrations are managed through a version-controlled migration framework that guarantees reproducible schema deployments across development, staging, and production environments.

B. Microservices and API Architecture

The SmartMeds backend adopts a modular microservices decomposition spanning five discrete service units: (1) a Core API Gateway built on Node.js/Express, responsible for authentication, request routing, and business logic governing EHR operations, appointment management, payments, and user lifecycle; (2) an AI Recommendation Service implemented in Flask/Python, encapsulating the OCR ingestion, SciSpacy NER, and ClinicalBERT inference pipeline; (3) a Chatbot Service in Flask/Python managing GPT API integration and intent-based response routing; (4) a Notification Service in Node.js coordinating Nodemailer email and Twilio SMS delivery; and (5) a File Processing

Service handling Multer-mediated uploads, AWS S3 integration, and AES-256 document encryption.

Synchronous REST calls serve request-response interaction patterns between services, while Redis pub/sub channels propagate asynchronous events—appointment state transitions, payment confirmations, and notification triggers—without tight service coupling. Each microservice exposes versioned API endpoints, permitting independent deployment and horizontal scaling without cascading system disruption. The React.js frontend communicates exclusively with the Core API Gateway, which proxies AI and chatbot requests internally to the Flask microservices.

VI. IMPLEMENTATION

A. Technology Stack

The full SmartMeds technology stack is as follows. Frontend: React.js with FullCalendar for appointment scheduling, Tailwind CSS for responsive layout, and Socket.IO client for real-time messaging. Backend: Node.js 18 LTS with Express.js, Multer for file handling, Nodemailer for email dispatch, and the Twilio SDK for SMS. Database: PostgreSQL 15 with Sequelize ORM; Redis 7 for caching and pub/sub. AI/ML Services: Python 3.10 with Flask; HuggingFace Transformers (ClinicalBERT, BioBERT); SciSpacy; Tesseract OCR; AWS Textract SDK. Payment: Razorpay SDK and Stripe SDK with webhook signature verification. Cloud Storage: AWS S3 for primary document storage; Cloudinary as an alternative CDN. Security: AES-256 encryption; JWT with RS256 signing; bcrypt password hashing; Helmet.js; rate-limiting middleware. PDF Generation: PDFKit for invoice and prescription document production.

B. Key Implementation Details

Incoming credential documents traverse a server-side Multer pipeline that enforces MIME type validation against an allowlist of PDF, PNG, and JPEG formats, caps individual file sizes at 10 MB, and performs ClamAV malware scanning prior to S3 upload. Object keys for stored credentials incorporate cryptographic content hashes to defeat enumeration attacks. Administrator access to credential documents during the review process is mediated through pre-signed S3 URLs bearing a 15-minute validity window, ensuring that sensitive documents are never exposed via permanently accessible public endpoints.

Payment webhook integrity is verified through HMAC-SHA256 re-computation against the gateway-supplied signature header before any state transition is executed,



International Journal of Recent Development in Engineering and Technology

Website: www.ijrdet.com (ISSN 2347 -6435 (Online)), Volume 15, Issue 5, May 2026)

neutralising webhook spoofing attack vectors. Idempotency key enforcement prevents duplicate webhook deliveries from producing duplicate payment entries or redundant appointment confirmations. All payment processing logic is encapsulated within PostgreSQL serializable transactions, guaranteeing atomic joint updates to the appointments and payments tables.

Socket.IO sessions are isolated within per-appointment namespaces to prevent message cross-contamination between concurrent consultation sessions. JWT-based socket authentication middleware re-validates session tokens on every connection event, immediately severing connections presenting expired or modified tokens. Server-side expiry enforcement disconnects active socket clients upon appointment end time, with a five-minute grace window accommodating clinical note completion.

VII. SECURITY ANALYSIS

A. Data Security

SmartMeds implements security as a layered, defense-in-depth model operating across three distinct tiers. At the data tier, AES-256-CBC encryption safeguards all medical and credential documents at rest within AWS S3; envelope encryption via AWS KMS governs key management. PostgreSQL RLS policies impose per-user data boundaries directly at the database engine level, independent of application-layer access logic. High-sensitivity structured fields—diagnostic notes, prescription content—receive an additional application-layer encryption pass before database persistence.

At the transport tier, TLS 1.3 with forward secrecy encrypts all client-to-server and inter-service traffic. HTTP Strict Transport Security (HSTS) headers enforce HTTPS-only access from all browser clients. At the application tier, JWT access tokens are signed with RS256 asymmetric keys, distributing public verification keys to microservices to support stateless token validation. Access token lifetimes are bounded to 15 minutes, with refresh token rotation limiting the exposure window of any compromised credential.

B. Payment Security

SmartMeds' payment processing pipeline satisfies PCI DSS requirements by relying on gateway tokenization: raw cardholder data never transits platform servers. Server-side payment order creation eliminates any client-side opportunity to manipulate charged amounts. HMAC-SHA256 webhook signature verification prevents spoofed payment event injection. Refund authorisations for amounts

above configurable thresholds require multi-factor administrator confirmation, providing a supplementary fraud prevention barrier.

C. Access Control and Audit

Granular RBAC middleware partitions system capabilities across patient, doctor, and administrator roles. Resource-level ownership checks confirm that patients may only access their own records and that physicians may retrieve patient records solely when an active appointment exists for the patient-physician pair. All consequential administrative actions—credential verifications, treating physician record access, payment processing operations—are durably recorded in the `admin_audit_logs` table, capturing actor identity, timestamp, originating IP address, and action payload to support forensic investigation and regulatory compliance demonstration.

VIII. RESULTS AND DISCUSSION

Platform evaluation was conducted across four complementary dimensions: AI recommendation fidelity, response performance under concurrent load, security robustness, and user experience quality.

The AI Recommendation Engine was assessed against a withheld test corpus of 1,200 anonymized clinical reports. The system achieved a Top-1 specialist assignment accuracy of 84.3% and a Top-3 accuracy of 96.1%, benchmarked against ground-truth specialist allocations confirmed by licensed clinicians. The fine-tuned ClinicalBERT model exceeded baseline TF-IDF and general-purpose BERT configurations by 11.2 percentage points on the Top-1 metric, providing empirical validation of the domain-adaptation strategy.

Performance benchmarking revealed that the appointment booking workflow concludes within 320 ms at the 95th percentile under a 500-concurrent-user load. End-to-end payment processing latency, inclusive of gateway API round-trips, averaged 1.8 seconds. Socket.IO-mediated chat message delivery under equivalent concurrent load measured a median latency of 48 ms, confirming the suitability of the WebSocket infrastructure for real-time clinical consultation use cases.

Security assessment via OWASP ZAP automated scanning, complemented by manual penetration testing, revealed no critical vulnerabilities in the production-equivalent deployment environment. SQL injection, reflected and stored XSS, and CSRF attack surfaces were systematically neutralised through parameterized query construction, Content Security Policy header enforcement,



International Journal of Recent Development in Engineering and Technology

Website: www.ijrdet.com (ISSN 2347 -6435 (Online)), Volume 15, Issue 5, May 2026)

and CSRF token validation respectively. All synthetically crafted malicious webhook payloads were rejected by the HMAC-SHA256 spoofing defence in 100% of test cases.

Usability testing with a cohort of 45 participants—drawn from patient and simulated doctor user classes—produced a System Usability Scale (SUS) aggregate score of 82.4, corresponding to the Excellent adjective rating. This result demonstrates that the platform's full multi-module feature set does not impose unacceptable navigational complexity on end users representing the intended clinical population.

IX. FUTURE SCOPE

Six development vectors are envisioned for future SmartMeds iterations. First, the implementation of HL7 FHIR API compliance will enable interoperability with hospital information systems and national EHR infrastructure, enabling bidirectional health record synchronisation with participating institutions. Second, WebRTC-based video consultation integration will extend the appointment-regulated connectivity model to support synchronous audiovisual clinical encounters, supplementing the existing asynchronous text channel.

Third, IoT device connectivity will permit automated ingestion of continuous biometric measurements—heart rate, blood oxygen saturation, blood glucose—from wearable devices into the longitudinal patient EHR timeline, furnishing AI-driven chronic disease risk stratification with continuous real-world health signals. Fourth, federated learning architectures will support distributed model improvement across multiple institutional data sources without requiring centralisation of sensitive patient data, addressing the privacy governance constraints that presently limit collaborative AI training at scale.

Fifth, blockchain-anchored audit trails will generate cryptographically immutable, externally verifiable records of medical record access events, addressing compliance requirements in jurisdictions with particularly stringent health data governance legislation. Sixth, multi-modal input support for clinical summarization—combining laboratory report text, imaging observations, and clinical narrative through vision-language transformer models—will substantially extend the platform's clinical decision support value for complex, multi-system patient presentations.

X. CONCLUSION

SmartMeds constitutes a substantive architectural evolution relative to conventional health record portals, integrating within a single, securely engineered ecosystem:

credential-verified physician onboarding, appointment-regulated patient-physician connectivity, mandatory online consultation payments, appointment-gated secure clinical communication, AI-driven specialist recommendation, and intelligent conversational platform assistance. The system demonstrates that production-ready healthcare platforms can simultaneously achieve clinical-grade AI recommendation fidelity, sub-second interactive performance characteristics, and rigorous compliance with contemporary security standards.

The platform's six-subsystem decomposition, sustained by a thirteen-table PostgreSQL data model, a versioned microservices API, and domain-adapted clinical NLP, furnishes a reproducible architectural template for the next generation of digital health platforms. SmartMeds' appointment-gated access paradigm guarantees that every clinical interaction is accountable, appropriately monetized, and auditable through comprehensive logging infrastructure. The empirically demonstrated SUS score of 82.4 and Top-3 recommendation accuracy of 96.1% collectively validate SmartMeds' readiness for deployment in controlled real-world clinical environments, establishing its standing as a meaningful contribution to accessible, intelligent, and secure digital healthcare delivery.

REFERENCES

- [1] B. Blobel, P. Pharow, and M. Nerlich, Eds., *eHealth: Combining Health Telematics, Telemedicine, Biomedical Engineering and Bioinformatics to the Edge*. Berlin: IOS Press, 2008.
- [2] P. Rajpurkar et al., "CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning," *arXiv preprint arXiv:1711.05225*, 2017.
- [3] J. Lee et al., "BioBERT: A Pre-trained Biomedical Language Representation Model for Biomedical Text Mining," *Bioinformatics*, vol. 36, no. 4, pp. 1234–1240, Feb. 2020.
- [4] R. Saripalle, C. Runyan, and M. Russell, "Using HL7 FHIR to Achieve Interoperability in Patient Health Records," *Journal of Biomedical Informatics*, vol. 94, p. 103188, Jun. 2019.
- [5] C. S. Kruse et al., "Telehealth and Patient Satisfaction: A Systematic Review and Narrative Analysis," *BMJ Open*, vol. 7, no. 8, p. e016242, 2017.
- [6] K. Abouelmehdi, A. Beni-Hssane, H. Khaloufi, and M. Saadi, "Big Healthcare Data: Preserving Security and Privacy," *Journal of Big Data*, vol. 5, no. 1, p. 1, Dec. 2018.
- [7] E. Alsentzer et al., "Publicly Available Clinical BERT Embeddings," in *Proc. 2nd Clinical Natural Language Processing Workshop*, Minneapolis, MN, 2019, pp. 72–78.



International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347 -6435 (Online)), Volume 15, Issue 5, May 2026)

- [8] A. Johnson et al., "MIMIC-III, A Freely Accessible Critical Care Database," *Scientific Data*, vol. 3, p. 160035, May 2016.
- [9] G. Huang, Z. Liu, L. van der Maaten, and K. Q. Weinberger, "Densely Connected Convolutional Networks," in *Proc. CVPR*, Honolulu, HI, Jul. 2017, pp. 4700–4708.
- [10] A. Vaswani et al., "Attention Is All You Need," in *Advances in Neural Information Processing Systems*, vol. 30, 2017, pp. 5998–6008.
- [11] M. Abadi et al., "TensorFlow: A System for Large-Scale Machine Learning," in *Proc. 12th USENIX OSDI*, Savannah, GA, Nov. 2016, pp. 265–283.
- [12] OWASP Foundation, OWASP Top Ten 2021. Accessed: Jan. 2025. [Online]. Available: <https://owasp.org/www-project-top-ten/>