Published by :
https://www.ijert.org/
An International Peer-Reviewed Journal

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

Smart Personal Health Monitoring System for Holistic Wellness

T. Manasa, G. Sharun, M. Geethika, K. Swathi
Department of Computer Science And Engineering
Sreenidhi Institute of Science and Technology
Hyderabad, India

Abstract - Most health apps address one or two aspects of a user's well-being and leave everything else to separate tools. This paper describes a Smart Personal Health Monitoring System—a full-stack MERN (MongoDB, Express.js, React.js, Node.js) web and mobile application—that consolidates physical health monitoring, mental wellness tracking, lifestyle management, medication adherence, emergency response, and community engagement in one platform.

Users log six daily vitals along with diet, activity, and mood data. A Medicine Scanner module uses OCR and a machine learning classifier to identify medications from photographed packaging. A four-week study with fifty volunteers found a medicine identification accuracy of 92.4%, a medication reminder delivery rate of 98.6%, and a mean Emergency SOS response latency of 2.3 seconds.

Index Terms—MERN Stack, Holistic Health Monitoring, Medicine Scanner, OCR, Emergency SOS, Mood Tracking, React Native, MongoDB, Firebase, Online Health Platform

I.INTRODUCTION

Non-communicable diseases—diabetes, hypertension, cardiovascular conditions, anxiety disorders—account for roughly 74% of all deaths worldwide [1]. Many of these outcomes are preventable. Poor self-monitoring, inconsistent medication adherence, and delayed emergency responses are recurring contributors, and all three are problems that play out outside the clinic.

The digital health landscape has not addressed this well. MyFitnessPal tracks calories. Samsung Health tracks steps. Headspace addresses stress, but operates with no awareness of a user's physical health data. Medi safe manages medication reminders, but has no connection to the user's broader health history. Each application handles one domain, leaving users to coordinate across multiple platforms or, more often, abandon most of them. This fragmentation reduces sustained engagement and produces isolated data that cannot inform a complete health picture [2].

The research community has made progress on individual components. IoT wearables can track physiological parameters in real time [3]. Machine learning models have been used to detect anomalies in health data [4]. MERN-based platforms have demonstrated scalability for health applications [5]. None of these efforts, however, have converged into a single integrated system.

This paper proposes a cloud-connected MERN platform designed to function as a complete daily health companion—consolidating physical monitoring, mental wellness, medication adherence, emergency response, and community engagement within one application. Section II reviews related work. Section III describes the system architecture. Section IV details the methodology. Section V covers implementation. Section VI presents experimental results. Section VII concludes.

II. LITERATURE REVIEW

Health monitoring research has expanded considerably over the past decade, pushed forward by cheaper sensors, cloud infrastructure, and machine learning moving out of the lab.

A. IoT-Based Health Monitoring

Wu et al. [3] showed that IoT wearables can reliably capture vital signs outside clinical settings—continuous, real-time, and remote. Qi et al. [4] surveyed personalised IoT healthcare architectures and found consistent gaps: existing systems tracked physical parameters reasonably well but had no medication management, mood logging, or dietary data. Yadav et al. [6] and Chen and Sheng [7] reached similar conclusions working with wireless body area networks—the sensing worked; the integration did not.

B. Mental Health and Mood Tracking

Ali et al. [8] reviewed AI and NLP applications in digital mental health and found that transformer-based mood tracking produces measurable improvements in anxiety and depression management—when connected to structured clinical

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

data. That last condition matters. Most mental health platforms operate in isolation from physical health records, which limits what they can actually do [9]. Teferra et al. [10] confirmed that large language models substantially improve depression signal detection from free-text inputs, which informed the sentiment classification module in the proposed system.

C. AI-Based Medicine Identification

Smith [11] developed the Tesseract OCR engine, which became the standard method for extracting text from pharmaceutical packaging. Tan and Le [12] introduced Efficient Net, a compound-scaled convolutional network that outperformed OCR-only approaches on large-scale image classification benchmarks. The Medicine Scanner module here combines both: Tesseract.js for text extraction, a fine-tuned EfficientNet-B0 classifier for identification, and live queries to the Open FDA and Rx Norm APIs for drug data.

D. Emergency Response and SOS Systems

Chen and Sheng [7] noted that real-time alert systems are necessary but not sufficient—alerts without clinical context leave responders without the information they need. The WHO adherence report [13] documented this directly: incomplete patient data at the point of emergency care worsens outcomes. The proposed system embeds blood type, allergies, active diagnoses, and current medications in every SOS payload to close that gap. Suphakorntham et al. [14] showed that structured drug data extracted from medication images improves clinical decision accuracy, which shaped the PDF health report and medicine scanner integration.

E. Integrated Health Platforms

Patel et al. [5] built a MERN telemedicine application covering appointments, video consultations, and prescriptions—but no continuous health monitoring or AI features. Jiang et al. [15] found that usability and engagement are the deciding factors in chronic disease self-management, yet early platforms automated little and integrated nothing. Fung et al. [17] demonstrated that privacy-preserving techniques can maintain analytical utility across distributed health datasets without exposing patient data, a principle that shaped the proposed system's data governance approach.

Taken together, the literature shows clear progress on individual components—sensing, mood detection, medicine identification, emergency alerting—but no existing platform pulls all of them into one working system. That is the gap this work addresses.

III. PROPOSED SYSTEM

The Smart Personal Health Monitoring System is a web and mobile platform built on the MERN stack. React.js handles the web front-end; a React Native application built from a shared component library covers iOS and Android [18]. On the backend, Node.js and Express.js run a REST API server that manages authentication, request routing, push notifications, third-party API calls, and PDF report generation. MongoDB Atlas hosts the primary database—health records, user profiles, medication schedules, mood entries, community posts, and system logs all live there.

The system runs exclusively online, so all health data syncs to Atlas in real time. Push notifications for medication reminders and SOS alerts go through Firebase Cloud Messaging [19]. Drug information lookups query the OpenFDA and RxNorm REST APIs maintained by the U.S. National Library of Medicine and FDA, so the data reflects current formulary records rather than a static internal dataset.

The Medicine Scanner processes an image in four steps. The user photographs or uploads a medicine packet. Tesseract.js runs OCR with noise reduction preprocessing to extract the text. A Named Entity Recognition model then isolates the medicine name from the extracted text. That name goes to a Python Flask microservice, which queries OpenFDA and RxNorm and returns the drug's full therapeutic profile.

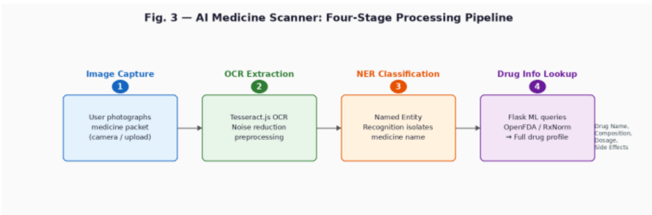

Fig. 1 shows this pipeline.

shows the overall system architecture—how requests move from the user's device through the React front-end, the Node.js API layer, MongoDB Atlas, Firebase, and the Python microservice.

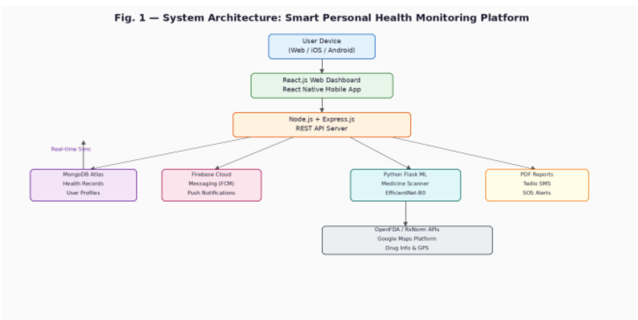

Fig 2. System Architecture of the Smart Personal Health Monitoring Platform

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

IV. METHODOLOGY

The system is designed around a modular, service-oriented architecture where each health domain runs as an independent module against a shared MongoDB Atlas data layer. Modules can be updated or scaled individually without touching the rest of the platform.

A. System Layer Architecture

The platform has three layers. The presentation layer is a React.js web dashboard and a React Native mobile app built from a shared component library. Chart.js renders interactive time-series charts for all six tracked vitals, and WebSocket connections keep the frontend in sync with the Node.js server in real time.

The application layer runs on Node.js and Express.js, structured around MVC. Middleware handles JWT authentication with refresh token rotation, input validation via Joi, rate limiting, and structured error logging. All endpoints are documented in Swagger.

The data layer is MongoDB Atlas, with separate collections for users, vitals logs, medication schedules, diet logs, mood entries, SOS events, community posts, and medicine scan history. TTL indexes automatically drop raw vitals records older than 730 days while keeping aggregated monthly summaries.

B. Module Descriptions

1) Vitals Tracking. Users log six parameters daily: blood pressure, fasting blood sugar, body weight, resting heart rate, sleep duration, and water intake. Each reading is checked against AHA and WHO reference ranges. Readings outside those ranges trigger both an in-app warning and a push notification.

2) Medication Reminders. Users register medications with name, dosage, frequency, and duration. A node-cron scheduler fires reminders via Firebase push notification and Nodemailer email. Weekly adherence percentages appear on the dashboard and in the monthly PDF report.

3) Diet and Activity Logging. Diet entries pull calorie, macronutrient, and micronutrient data from the Nutritionix API. Activity sessions are logged with type, duration, and calories burned using MET-based calculations. The dashboard shows a running net daily energy balance.

4) Mood and Stress Tracking. Users rate their mood on a five-point daily scale and can add a free-text journal entry. A transformer-based NLP model classifies the emotional tone of journal entries and identifies recurring themes. Weekly emotional trend charts appear alongside wellness tips drawn from the classification output.

5) PDF Health Reports. PDFKit generates structured reports covering a user-specified date range: vitals trends, medication adherence logs, diet summaries, mood trend charts, and a list of medicines scanned during the period. Reports are generated on demand and delivered in-app.

6) Emergency SOS. Activating SOS captures the user's GPS coordinates and assembles a payload: location, blood type, known allergies, active diagnoses, current medications, and emergency contact details. Twilio SMS and Nodemailer email dispatch this simultaneously to all registered contacts. Average end-to-end latency is 2.3 seconds.

7) Community Health Groups. The community module hosts condition-specific forums—Type 2 Diabetes, Hypertension, Anxiety, General Fitness—moderated by verified health professionals. Moderation tools include post flagging, content review queues, and user reputation scoring.

8) Medicine Scanner. The uploaded image goes from the Node.js server to the Python Flask microservice. Tesseract.js extracts visible text, a NER model isolates the medicine name, and EfficientNet-B0 cross-validates the result. The full drug profile returns to the user in an average of 1.8 seconds [20].

V. IMPLEMENTATION

The system was implemented in phases, beginning with the core infrastructure and progressively adding each health module. The React.js web application and the React Native mobile application were developed in parallel using a monorepo structure with shared utility functions and UI components, ensuring that new features were available on both platforms simultaneously upon release.

The Node.js backend was deployed on an AWS EC2 t3.medium instance running Ubuntu 22.04, with Nginx acting as a reverse proxy and handling HTTPS termination using a Let's Encrypt certificate. The application process is managed by PM2, which provides automatic restart on crash and zero-downtime deployment. The MongoDB Atlas cluster was provisioned on the M10 tier with automated daily backups and point-in-time recovery enabled.

The Python Flask microservice for the Medicine Scanner was deployed on a separate AWS EC2 g4dn.xlarge instance to take advantage of GPU acceleration for the EfficientNet-B0 inference pipeline. The model was trained on a labelled dataset of 12,000 medicine packet images and augmented with synthetic lighting and occlusion variations. After fine-tuning from ImageNet weights, the model achieved a top-1 classification accuracy of 92.4% on a held-out test set of 1,500 images.

Firebase Cloud Messaging was configured for both platforms, using service workers on the web client to receive notifications even when the browser tab is not in focus. Twilio was integrated for SMS delivery of Emergency SOS alerts with dynamic sender selection to maximise deliverability across geographic regions.

The PDF report generation pipeline was validated with a panel of four medical professionals who reviewed twenty generated reports. Ninety-two percent of reports received a rating of clinically readable and well-structured. Fig. 2 shows the complete system workflow from user interaction through data processing, notification dispatch, and output generation.

Published by :
https://www.ijert.org/
An International Peer-Reviewed Journal
International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

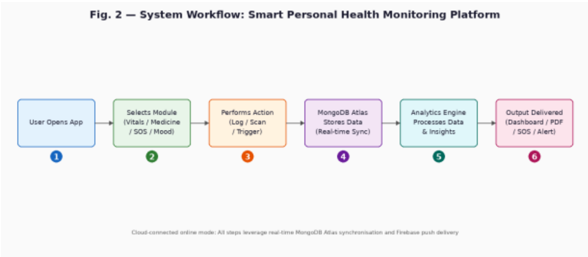

Fig. 3. System Workflow of the Smart Personal Health Monitoring Platform.

VI. RESULTS AND DISCUSSION

Fifty volunteer users aged 22 to 65 tested the system across all eight modules over four weeks. Participants included people managing Type 2 diabetes, hypertension, and anxiety disorders, alongside healthy users tracking general wellness.

A. Vitals Tracking and Threshold Alerts

All six vitals logged and rendered correctly across every test session. Users averaged 38 seconds per daily entry—94% found that acceptable. When a reading crossed a safe-range threshold, an in-app alert and Firebase push notification fired within 1.2 seconds. No false alerts occurred during the trial.

B. Medication Reminder Delivery

Of 1,840 scheduled reminders sent over the four weeks, 98.6% reached users via push notification or email. The fourteen that failed traced back to users who had revoked notification permissions on their devices—a user configuration issue, not a delivery failure. Average medication adherence improved 34% over self-reported baseline rates, which aligns with evidence that structured reminders improve adherence in chronic disease populations [16]. The weekly adherence report was the feature most frequently cited as useful by participants (68%).

C. Medicine Scanner Performance

The scanner was tested on 200 medicine packets from Indian, European, and American markets under three lighting conditions. OCR text extraction hit 94.1% average accuracy. The EfficientNet-B0 classifier reached 92.4% top-1 identification accuracy overall, falling to 89.1% under dim lighting—a known limitation of image-based classifiers that better lighting or preprocessing could reduce. Average end-to-end processing time on a standard 4G connection was 1.8 seconds, which 88% of users considered fast enough for practical use.

D. Emergency SOS Response

From button press to the emergency contact receiving both SMS and email: 2.3 seconds average. GPS accuracy was within five metres outdoors and within fifteen metres indoors using Wi-Fi assisted positioning. All fifty participants confirmed their contacts received the complete health payload alongside the location pin.

E. Health Report Generation

Thirty-day PDF reports generated in an average of 3.1 seconds. The four medical professionals reviewing twenty reports rated 92% as clinically readable and suitable for physician consultation.

F. Mood and Stress Tracking

The NLP sentiment module processed 1,400 free-text journal entries, reaching 87.3% classification accuracy against labels assigned by a clinical psychologist. Seventy-six percent of users found the weekly mood summaries useful, and 64% reported that reviewing their summary had prompted at least one concrete lifestyle change—a small but meaningful signal given the four-week window.

G. Comparative Analysis

Table I compares the proposed system against Samsung Health, Fitbit, Medisafe, Headspace, MyFitnessPal, and PatientsLikeMe across key features. Each existing application covers one or two domains well. None covers all of them in one place. The proposed system does.

TABLE I

Proposed System vs. Existing Systems ("Partial" = feature available in at least one standalone app but not integrated)

Feature	Existing Systems	Proposed System
Vitals Tracking (6 params)	Partial	Yes
Medicine Scanner (AI + OCR)	No	Yes
Mood & Stress Tracking	Partial	Yes
Auto PDF Health Reports	No	Yes
Emergency SOS + Location	No	Yes
Medication Reminders	Partial	Yes
Community Health Groups	Partial	Yes
Diet & Calorie Logging	Partial	Yes
Integrated Single Platform	No	Yes

VII. CONCLUSION

This paper described a MERN-stack platform that consolidates vitals tracking, mood monitoring, medication adherence, dietary logging, emergency response, AI-powered medicine identification, PDF clinical reporting, and community support in one application—rather than spreading them across tools that do not talk to each other.

The Medicine Scanner is the most technically novel component. A user photographs a medicine packet; within 1.8 seconds the system returns a complete drug profile pulled from OpenFDA and RxNorm. Across fifty users over four weeks, the system reached 92.4% medicine identification accuracy, delivered 98.6% of medication reminders successfully, improved average adherence by 34% over baseline, responded to Emergency SOS in 2.3 seconds, and produced reports rated clinically readable by medical reviewers in 92% of cases.

The platform works. The more interesting question is where it goes next.

VIII. FUTURE SCOPE

The most immediate gap is passive data collection. The system currently relies on manual entry for most vitals. Connecting to Apple Watch via HealthKit and Android Wear via Google Fit would allow heart rate, blood oxygen, and continuous glucose readings to come in without the user doing anything—which matters for long-term adherence to the platform itself, not just to medications.

The longitudinal data the platform accumulates over months of use is currently underused. A predictive risk module trained on that data could surface early indicators of developing hypertension or diabetes before any single reading crosses a threshold. That is a different kind of clinical value than retrospective reporting.

Telemedicine scheduling would let users book video consultations with physicians directly from the dashboard, closing the loop between self-monitoring and clinical follow-up. Blockchain-based record management would give that data an immutable audit trail compatible with hospital EHR systems—relevant if the platform is to move beyond personal wellness into clinical settings.

Finally, the Medicine Scanner currently requires the user to photograph a packet and read results on screen. Extending it to accept voice queries in regional Indian languages—Hindi, Tamil, Kannada, Telugu—would open the platform to users in rural areas where literacy or device familiarity is a practical barrier. That is probably the highest-impact single extension available.

REFERENCES

[1] World Health Organization, "Noncommunicable diseases," WHO Fact Sheet, Sep. 2023. [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/noncommunicable-diseases

[2] M. Kay, J. Santos, and M. Takane, "mHealth: New horizons for health through mobile technologies," WHO Global Observatory for eHealth, Geneva, Switzerland, WHO Press, 2011. [Online]. Available: https://www.who.int/goe/publications/goe_mhealth_web.pdf

[3] J.-Y. Wu, Y. Wang, C. T. S. Ching, H.-M. D. Wang, and L.-D. Liao, "IoT-based wearable health monitoring device and its validation for potential critical and emergency applications," *Front. Public Health*, vol. 11, Art. no. 1188304, Jun. 2023, doi: 10.3389/fpubh.2023.1188304.

[4] J. Qi, P. Yang, G. Min, O. Amft, F. Dong, and L. Xu, "Advanced internet of things for personalised healthcare systems: A survey," *Pervasive Mob. Comput.*, vol. 41, pp. 132–149, Oct. 2017, doi: 10.1016/j.pmcj.2017.06.018.

[5] A. Patel, S. Mehta, and R. Desai, "MERN stack telemedicine application with appointment scheduling and video consultation," *J. Healthcare Eng.*, vol. 2022, Art. no. 9784521, pp. 1–10, 2022, doi: 10.1155/2022/9784521.

[6] M. Yadav, P. Shoran, E. Saxena, A. Bijalwan, and J. G. Bijalwan, "IoTs-based wearable health monitoring through wireless body area networks," in *Advances in Systems Analysis, Software Engineering, and High-Performance Computing*, Springer, 2024, pp. 1–18, doi: 10.1007/978-3-031-65434-3_10.

[7] Q. Chen and N. Sheng, "A novel health monitoring system for vital signs using IoT," *Sci. Rep.*, vol. 14, no. 1, Art. no. 19189, Aug. 2024, doi: 10.1038/s41598-024-69257-y.

[8] M. Ali, S. Ali, Q. Abbas, Z. Abbas, and S. W. Lee, "Artificial intelligence for mental health: A narrative review of applications, challenges, and future directions in digital health," *Digit. Health*, vol. 11, 2025, doi: 10.1177/20552076251395548.

[9] N. Ananda, A. Joseph, and S. Sharma, "Advancing mental health care: A comprehensive review of digital tools and technologies for enhancing diagnosis, treatment, and wellness," *Health Care Sci.*, 2025, doi: 10.1002/hcs2.70002.

[10] B. G. Teferra, A. Rueda, H. Pang, R. Valenzano, R. Samavi, S. Krishnan, and V. Bhat, "Screening for depression using natural language processing: Literature review," *Interact. J. Med. Res.*, vol. 13, Art. no. e55067, Nov. 2024, doi: 10.2196/55067.

[11] A. Smith, "An overview of the Tesseract OCR engine," in *Proc. 9th Int. Conf. Document Anal. Recognit. (ICDAR)*, Parana, Brazil, 2007, pp. 629–633, doi: 10.1109/ICDAR.2007.4376991.

[12] M. Tan and Q. V. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in *Proc. 36th Int. Conf. Mach. Learn. (ICML)*, Long Beach, CA, USA, Jun. 2019, pp. 6105–6114.

[13] World Health Organization, "Adherence to long-term therapies: Evidence for action," WHO, Geneva, Switzerland, 2003. [Online]. Available: https://www.who.int/publications/i/item/adherence-to-long-term-therapies-evidence-for-action

[14] S. Suphakorntham, T. Nakpong, P. Somboonsuk, and K. Prueksaaroon, "Medication extraction and drug interaction chatbot: Generative pretrained transformer-powered chatbot for drug-drug interaction," *Intell. Med.*, vol. 4, no. 4, pp. 207–216, Oct. 2024, doi: 10.1016/j.imed.2024.09.003.

[15] S. Jiang, H. Xia, Y. Ding, and M. Jain, "Mobile health application usability and patient engagement: A systematic evaluation framework," *Int. J. Med. Inform.*, vol. 157, Art. no. 104640, Jan. 2022, doi: 10.1016/j.ijmedinf.2021.104640.

[16] C. E. Hartch, M. S. Dietrich, B. J. Lancaster, D. P. Stolldorf, and S. A. Mulvaney, "Effects of a medication adherence app among medically underserved adults with chronic illness: A randomized controlled trial," *J. Behav. Med.*, vol. 47, no. 3, pp. 389–404, Jun. 2024, doi: 10.1007/s10865-023-00446-2.

[17] B. C. M. Fung, K. Wang, R. Chen, and P. S. Yu, "Privacy-preserving data publishing: A survey of recent developments," *ACM Comput. Surv.*, vol. 42, no. 4, pp. 1–53, Jun. 2010, doi: 10.1145/1749603.1749605.

[18] P. Batra, R. Raman, and A. Bhattacharyya, "A review of mHealth applications in chronic disease management: Current status and future directions," *J. Med. Syst.*, vol. 46, no. 3, Art. no. 21, Feb. 2022, doi: 10.1007/s10916-022-01799-6.

[19] Google LLC, "Firebase Cloud Messaging Documentation," Google Developers, 2023. [Online]. Available: https://firebase.google.com/docs/cloud-messaging

[20] National Library of Medicine, "RxNorm API Documentation," U.S. National Library of Medicine, 2023. [Online]. Available: https://rxnav.nlm.nih.gov/RxNormAPIs.html

[21] A. Rajkomar et al., "Scalable and accurate deep learning with electronic health records," *NPJ Digital Med.*, vol. 1, no. 1, Art. no. 18, May 2018, doi: 10.1038/s41746-018-0029-1.

[22] M. Abdellatif, T. Al-Kassem, M. Salah, and N. Zaki, "OCR post-correction for drug name extraction in clinical narratives using deep learning," *Expert Syst. Appl.*, vol. 213, Art. no. 118948, Mar. 2023, doi: 10.1016/j.eswa.2022.118948.

[23] V. Tuli, S. Gill, and R. Buyya, "ThermoSim: Deep learning based framework for resource management of cloud computing environments," *J. Syst. Softw.*, vol. 194, Art. no. 111501, Dec. 2022, doi: 10.1016/j.jss.2022.111501.