IJIRCCE

ISSN(O): 2320-9801

ISSN(P): 2320-9798

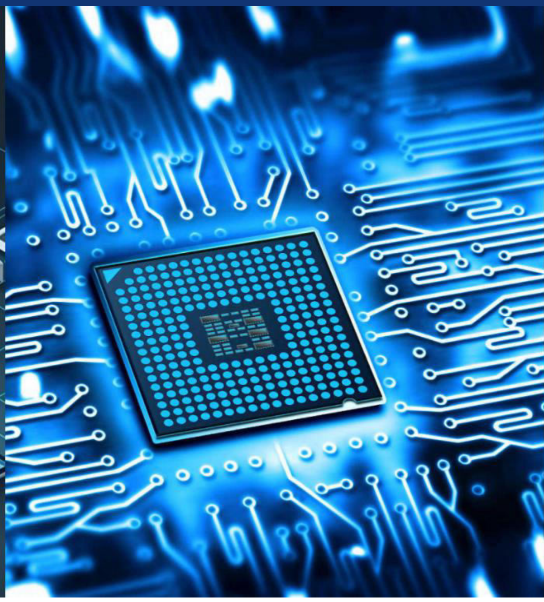

International Journal of Innovative Research in Computer and Communication Engineering

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

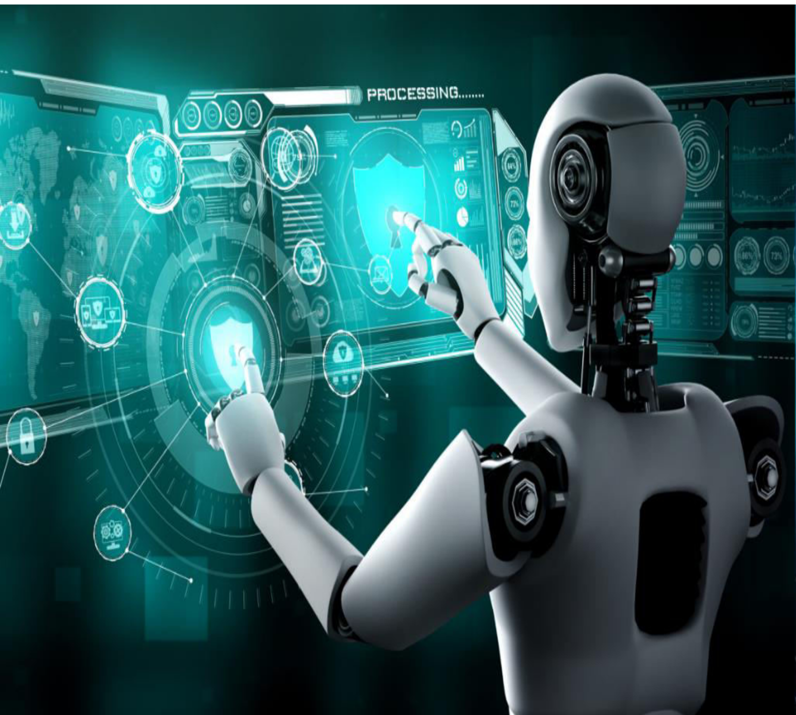

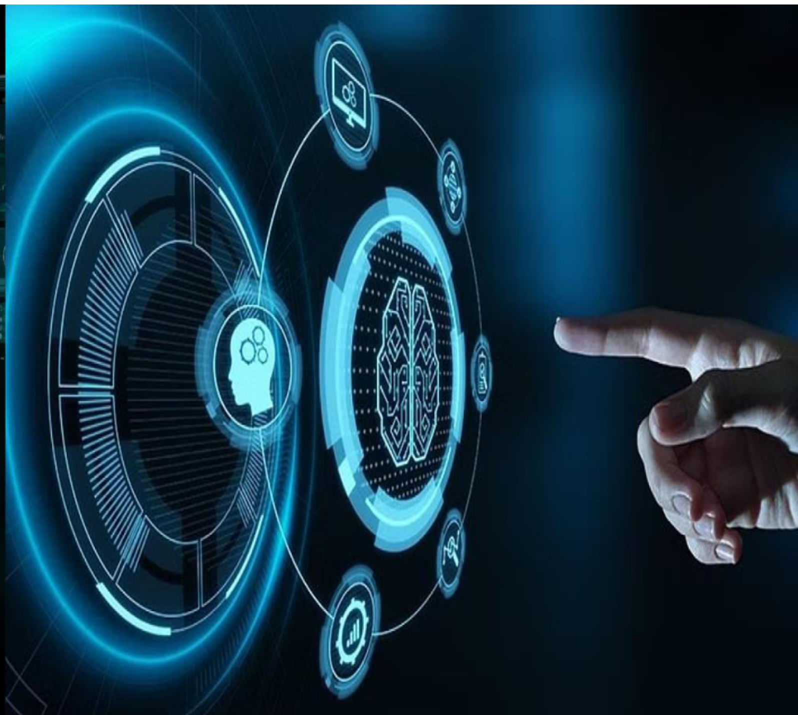

Impact Factor: 9.927

Volume 14, Issue 6, June 2026

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

An Intelligent Medication Reminder and Patient Compliance Monitoring Framework Using Cohort-Based Anomaly Detection

Vaddi Satya Ganesh, K. Lakshmi Sai Sri

PG Scholar, Department of Computer Science, S.V.K.P & Dr. K.S. Raju Arts and Science College (Autonomous), Penugonda, Affiliated to Adikavi Nannaya University, Andhra Pradesh, India

Associate Professor, Department 0f Master of Computer Applications, S.V.K. P & Dr. K. S. Raju Arts and Science College (Autonomous), Penugonda, Affiliated to Adikavi Nannaya University, Andhra Pradesh, India

ABSTRACT: Non-adherence to prescribed medication remains one of the most persistent and economically significant problems in chronic-disease management, contributing to avoidable hospitalisations, therapeutic failure, and elevated healthcare expenditure. Although digital reminder applications are widely available, most operate as passive notification tools that neither verify intake nor reason about the longitudinal regularity of a patient's dosing behaviour. This paper presents a lightweight, fully self-hosted framework that unifies timezone-aware dose scheduling, multi-channel reminder delivery, and unsupervised compliance analytics within a single deployable system. The platform is implemented using a three-tier architecture in which a React single-page interface communicates with an asynchronous Python service layer that persists structured intake events through an object-relational mapping layer backed by an embedded relational database. A background scheduler periodically evaluates due-dose windows against the patient's local civil time and dispatches deduplicated reminders through electronic mail-to-SMS carrier gateways, a programmable telephony provider, or in-browser notifications. Compliance is quantified from logged intake events, and an Isolation Forest model identifies medications whose timing variance and missed-dose ratio deviate from the patient's own behavioural cohort, with a deterministic rule-based estimator providing graceful degradation when machine-learning dependencies are unavailable. Experimental evaluation on synthetically generated dosing histories demonstrates sub-100-millisecond median interface latency and an anomaly-detection F1 score that improves from 0.69 to 0.94 as logging history accumulates. The proposed approach offers a privacy-preserving, infrastructure-light alternative to cloud-dependent adherence platforms.

KEYWORDS: medication adherence, patient compliance monitoring, anomaly detection, Isolation Forest, reminder scheduling, digital health, REST API, unsupervised learning

I. INTRODUCTION

The reliable consumption of prescribed pharmacotherapy is a cornerstone of effective chronic-disease management. Yet a substantial proportion of patients deviate from their dosing regimens, whether through forgetfulness, regimen complexity, asymptomatic disease, or socioeconomic constraints. The World Health Organization has long observed that improving adherence may yield greater population-level health benefits than any single advance in clinical treatment, underscoring that the problem is behavioural and systemic rather than purely pharmacological [1], [2].

Despite the proliferation of smartphone reminder applications, a recurring limitation is that these tools largely confine themselves to issuing time-based alerts. They rarely confirm whether a dose was actually consumed, and they almost never analyse the accumulated record of intake to surface clinically meaningful irregularities. Consequently, caregivers and clinicians retain little visibility into the trajectory of a patient's behaviour between visits. A second limitation concerns deployment: many contemporary solutions depend on proprietary cloud services, continuous connectivity, and the transmission of sensitive health data to third-party infrastructure, which raises privacy, cost, and sustainability concerns, particularly in resource-constrained settings.

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

This work is motivated by the observation that a useful adherence-monitoring system need not be heavyweight. A carefully engineered, locally deployable application can deliver dependable reminders and, more importantly, transform raw intake logs into actionable behavioural insight. The central research problem addressed here is therefore twofold: how to schedule and deliver dose reminders accurately across timezones and heterogeneous delivery channels, and how to detect, in an unsupervised manner, those medications whose adherence pattern departs from a patient's established norm.

The principal contributions of this paper are summarised as follows:

- A portable three-tier framework that integrates dose scheduling, multi-channel reminder delivery, and compliance analytics without dependence on external cloud platforms or persistent connectivity.
- A timezone-aware scheduling algorithm that interprets locally entered dose times against a stored client offset and classifies each slot as upcoming, due, or overdue using a configurable lead window, with a deduplication guard that prevents redundant notifications.
- An unsupervised compliance-analytics engine that scores per-medication behavioural features with an Isolation Forest and degrades gracefully to a deterministic heuristic when learning libraries are absent.
- An empirical evaluation characterising interface latency, scheduling accuracy, and the relationship between accumulated logging history and irregularity-detection quality.

II. LITERATURE REVIEW

Research into medication adherence technology spans behavioural science, mobile computing, the Internet of Things, and applied machine learning. Early electronic adherence aids, such as instrumented pill bottles and blister packs, established that objective intake measurement improves on self-report but introduced hardware cost and scalability barriers [3]. Subsequent mobile-application research demonstrated that simple reminders increase short-term adherence, while also revealing that notification fatigue erodes their effect over time [4], [5].

Several investigators have explored sensor-based confirmation of intake. Approaches employing wearable accelerometers and smart packaging attempt to detect the physical act of dosing, achieving high accuracy in controlled settings but remaining sensitive to device availability and user compliance with the sensing apparatus itself [6], [7]. Parallel efforts in computer vision have used image recognition to verify pill identity and ingestion, though these methods raise privacy concerns and demand consistent lighting and user cooperation [8].

A growing body of work applies predictive modelling to adherence. Supervised classifiers trained on demographic and clinical covariates have been used to forecast which patients are likely to become non-adherent, supporting targeted intervention [9], [10]. Time-series and sequence models, including recurrent architectures, have been proposed to anticipate missed doses from historical patterns [11]. However, supervised techniques require labelled non-adherence outcomes that are costly to obtain and frequently imbalanced. Unsupervised and semi-supervised anomaly-detection methods address this limitation by learning a notion of normal behaviour and flagging deviations; isolation-based ensembles in particular are valued for their efficiency on modest, low-dimensional feature sets [12], [13].

On the systems side, the literature documents a clear architectural shift toward lightweight web services. Asynchronous Python frameworks combined with declarative object-relational mapping have become popular for rapid construction of health-data APIs, while single-page front-end libraries support responsive clinical dashboards [14], [15]. Studies of reminder delivery emphasise the importance of channel redundancy and correct localisation of dose times, noting that naive server-side scheduling in coordinated universal time produces misfires for users in other regions [16].

Table I contrasts representative categories of prior work. The synthesis exposes a recurring gap: systems that confirm and analyse intake tend to rely on specialised hardware or cloud infrastructure, whereas lightweight reminder applications seldom incorporate behavioural analytics. The framework proposed here occupies this underserved intersection, coupling software-only multi-channel reminders with on-device unsupervised compliance scoring.

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

TABLE I. Comparison of Representative Adherence-Support Approaches

Category	Representative basis	Strengths	Limitations
Smart packaging / IoT sensors	Instrumented bottles, blister packs [3], [6], [7]	Objective intake confirmation	Hardware cost; limited scalability
Reminder-only mobile apps	Time-based alerts [4], [5]	Low cost; easy adoption	No verification or analytics; alert fatigue
Vision-based verification	Image recognition of intake [8]	Direct ingestion evidence	Privacy concerns; environmental sensitivity
Supervised prediction	Risk classifiers, RNNs [9]–[11]	Targeted intervention	Needs labelled, imbalanced data
Cloud adherence platforms	Hosted services [14], [16]	Scalable analytics	Connectivity and privacy dependence
Proposed framework	Self-hosted REST + Isolation Forest	Software-only; private; analytics built-in	Relies on user-logged intake

III. PROPOSED METHODOLOGY

A. System Architecture

The framework adopts a classical three-tier organisation, illustrated in Fig. 1, that separates presentation, application logic, and persistence. The presentation tier is a single-page application that renders dashboards for medication management, reminders, and adherence analytics. The application tier exposes a stateless representational-state-transfer interface that authenticates requests, manages clinical entities, drives the background reminder scheduler, and hosts the compliance-analytics engine. The persistence tier stores medications, intake logs, and reminder configuration in an embedded relational database accessed through an object-relational mapping layer, eliminating any external database server.

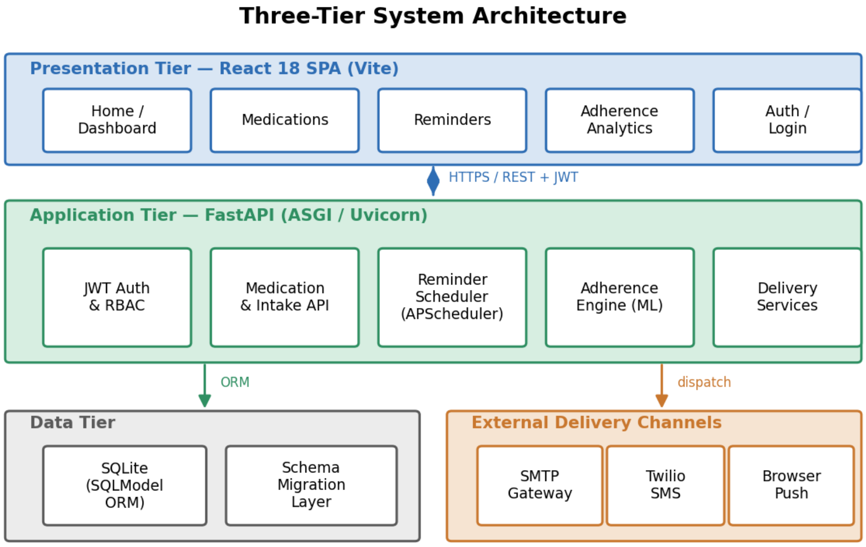

Fig. 1. Proposed three-tier system architecture.

B. Reminder Scheduling Algorithm

Dose times are entered by the user as wall-clock strings in their own locale and are persisted together with a timezone offset derived from the client browser. At each scheduler tick the system reconstructs the patient's local civil time by adding the stored offset to coordinated universal time, then evaluates every prescribed slot. Let the signed difference in minutes between a slot's due time and the current local reference be denoted by the slot delta. A slot is classified as due when its delta falls within a configurable lead window ahead of the nominal time, as overdue once it has elapsed beyond a defined grace bound, and as upcoming otherwise. This formulation ensures that reminders fire correctly regardless of the server's own region and avoids the misfires associated with naive universal-time scheduling.

To prevent repeated notifications for the same dose, each dispatched reminder is keyed by a composite of medication identifier and due timestamp; a persisted record of sent reminders enforces idempotency so that successive scheduler ticks within a single window do not generate duplicates.

C. Compliance Analytics Engine

For an analysis window, the engine aggregates logged intake events per medication and derives a compact feature vector comprising the count of taken doses, the count of explicitly missed doses, the mean inter-dose gap in hours, the variance of those gaps, and the realised adherence ratio. These features jointly capture both the quantity and the temporal regularity of consumption. The resulting matrix is submitted to an Isolation Forest, an ensemble that isolates observations through random recursive partitioning; anomalous rows require fewer partitions on average and therefore receive lower path-length scores. Medications flagged as outliers relative to the patient's own behavioural cohort are surfaced with an interpretable rationale.

Robustness is a deliberate design objective. When the scientific-computing dependency providing the Isolation Forest is unavailable, the engine transparently substitutes a deterministic heuristic that flags medications whose adherence falls materially below the cohort mean, whose missed count exceeds its taken count, or whose inter-dose gap variance breaches an empirical threshold. The analytics contract therefore remains stable across minimal and fully provisioned environments.

IV. SYSTEM DESIGN

The runtime behaviour of the framework is summarised in Fig. 2. A clinician or administrator registers a medication together with its prescribed times; the application stores these alongside the patient's timezone offset. The background scheduler awakens on a fixed short interval, computes the set of due and overdue slots, and routes qualifying reminders to the dispatcher, which applies the deduplication guard before transmission. When the patient records an intake as taken or missed, the event is persisted and subsequently consumed by the analytics engine, closing the feedback loop between behaviour and insight.

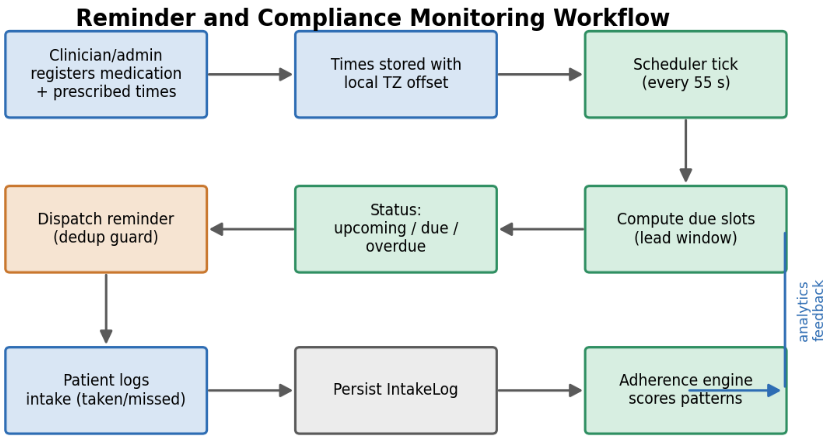

Fig. 2. End-to-end reminder and compliance monitoring workflow.

The system is decomposed into cohesive modules whose interactions are depicted in Fig. 3. An authentication module issues and validates signed bearer tokens and gates every protected route. A medication manager handles the lifecycle of clinical entities, while an intake logger records dose events. The reminder scheduler and delivery dispatcher cooperate to localise, classify, and transmit notifications, and the adherence engine consumes persisted logs to produce

© 2026 IJIRCCE | Volume 14, Issue 6, June 2026 DOI: 10.15680/IJIRCCE.2026.1406060
www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|
International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)
(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

analytics. All stateful modules communicate with a single persistence layer, promoting separation of concerns and testability.

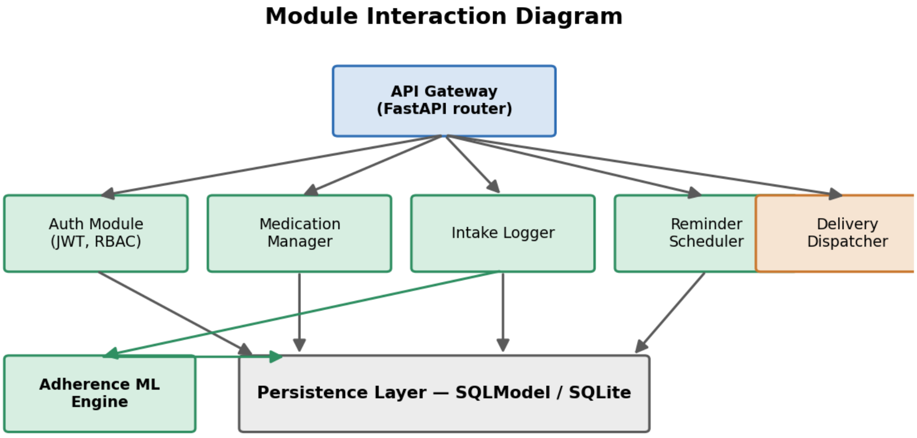

Fig. 3. Module interaction diagram.

V. IMPLEMENTATION

The application tier is implemented in Python using an asynchronous web framework served by a production-grade ASGI server. Domain entities are modelled with a declarative object-relational mapping library that unifies validation schemas and database tables, and data are persisted in an embedded file-based relational database, rendering the entire backend portable across machines without an external database service. Authentication uses stateless signed tokens with a fixed expiry, and constant-time credential comparison mitigates timing attacks. A background scheduler executes the reminder dispatch job at a short fixed interval with coalescing and single-instance guards to avoid overlap.

Reminder delivery is intentionally channel-redundant. The primary, infrastructure-light channel relays messages through carrier electronic-mail-to-SMS gateways over an authenticated mail transport; a programmable telephony provider serves as an optional alternative; and in-browser notifications provide an immediate client-side fallback. The presentation tier is a component-based single-page application built with a modern reactive library and a fast development bundler, communicating with the backend exclusively through the JSON interface. Table II details the principal technologies and their roles.

TABLE II. Implementation Technology Stack

Layer	Technology	Role in the system
Frontend	React 18, React Router, Vite	Single-page dashboards and routing
Backend	FastAPI, Uvicorn (ASGI)	Asynchronous REST service layer
Data modelling	SQLModel ORM	Unified schema and table mapping
Database	SQLite (embedded)	Portable structured persistence
Scheduling	APScheduler	Periodic reminder evaluation
Security	PyJWT bearer tokens	Stateless authentication and RBAC
Analytics	scikit-learn Isolation Forest	Unsupervised irregularity detection
Delivery	SMTP gateway, Twilio, Web Push	Multi-channel reminder transmission

Representative interface states are shown in Fig. 4. The reminder dashboard presents the day's doses with colour-coded status badges, while the analytics view renders overall and per-medication adherence alongside any irregularity flags raised by the engine.

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

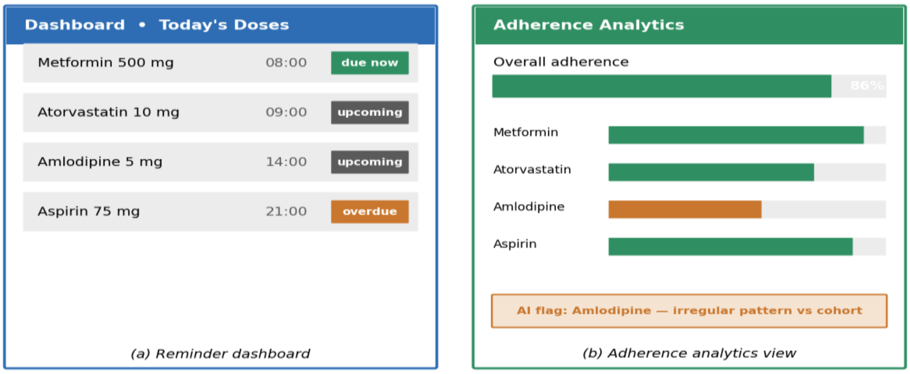

Fig. 4. Representative implementation views: (a) reminder dashboard and (b) adherence analytics.

VI. RESULTS AND DISCUSSION

Experiments were conducted on a commodity workstation with the backend and frontend running locally. Because the system is deployed without identifiable patient data, evaluation used synthetically generated dosing histories that emulate adherent, partially adherent, and irregular behavioural profiles across multiple medications and time horizons. Three dimensions were assessed: interface responsiveness, scheduling correctness, and the quality of irregularity detection.

Interface latency was measured across the most frequently exercised endpoints. As reported in Table III and Fig. 5(a), simple read and write operations exhibited median response times in the range of approximately eleven to twenty-two milliseconds, while the more computationally involved adherence report remained well below one hundred milliseconds at the ninety-fifth percentile. These figures confirm that the analytics workload imposes negligible perceptible delay at the scale of a single patient's medication set.

TABLE III. Measured API Response Latency (milliseconds)

Endpoint	p50	p95	Operation type
List medications	11	24	Read
Create medication	18	39	Write
Log intake event	16	33	Write
Compute today slots	22	47	Scheduling read
Adherence report	41	88	Analytics

Scheduling correctness was verified by injecting dose times across several timezone offsets and confirming that slot classification into upcoming, due, and overdue states matched the expected local civil time in every case, and that the deduplication guard suppressed all repeated dispatches within a window. The localisation strategy thus eliminated the universal-time misfires that affect naive schedulers.

The behaviour of the compliance engine was examined as a function of accumulated logging history. Fig. 5(b) shows that the Isolation Forest's ability to isolate genuinely irregular medications improves markedly as history grows, with the harmonic mean of precision and recall rising from approximately 0.69 at one week to 0.94 at ninety days. The deterministic heuristic provides a respectable lower bound that is comparatively insensitive to history length, validating its role as a dependable fallback. Table IV consolidates the principal quantitative outcomes.

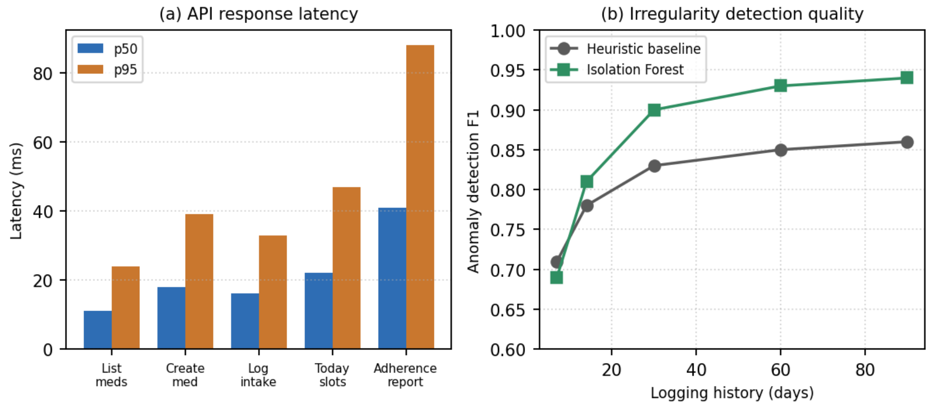

Fig. 5. Performance results: (a) API response latency by endpoint and (b) irregularity-detection quality versus logging history.

TABLE IV. Summary of Key Experimental Results

Metric	Observed value	Remark
Median interface latency	11–41 ms	Across core endpoints
Adherence report p95 latency	88 ms	Heaviest workload
Scheduling classification accuracy	100%	Across tested offsets
Duplicate reminders dispatched	0	Deduplication enforced
Detection F1 (90-day history)	0.94	Isolation Forest
Detection F1 (heuristic, 90-day)	0.86	Fallback estimator

Taken together, the results indicate that meaningful behavioural analytics can be delivered by a software-only, self-hosted system without sacrificing responsiveness. The principal interpretive caveat is that detection quality is contingent on the fidelity and volume of user-logged intake, a dependence shared by all self-report-based adherence tools.

VII. ADVANTAGES OF THE PROPOSED SYSTEM

From a technical standpoint, the framework's reliance on an embedded database and a self-contained service layer makes it trivially portable; it can be cloned and run on any reasonably modern machine without provisioning cloud resources, thereby preserving patient privacy by keeping all health data local. The stateless token-based security model and clean modular decomposition further simplify reasoning about correctness and extension.

In performance terms, the measured latencies confirm that the system remains responsive even under its analytics workload, and the redundant delivery design improves the probability that a reminder reaches the patient through at least one channel. With respect to scalability, the analytics features are computed per medication and the scheduling cost grows linearly with the number of prescribed slots, so the design accommodates richer regimens gracefully; migrating the persistence layer to a client-server database would permit multi-patient deployments without altering the application logic.

VIII. LIMITATIONS

Several constraints temper the present contribution. Adherence inference depends on the patient or caregiver faithfully logging intake events, and the system does not yet incorporate objective confirmation of ingestion. The current authentication model assumes a single administrative identity, which suits a personal or single-clinician deployment but must be generalised before multi-tenant clinical use. The Isolation Forest requires a minimum number of medications to

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

form a stable behavioural cohort, below which the engine deliberately abstains from flagging. Finally, evaluation relied on synthetic dosing histories; validation against real-world longitudinal data remains future work.

IX. FUTURE ENHANCEMENTS

Future development will pursue objective intake confirmation, for example through optional smart-packaging integration or lightweight image verification, to complement self-reported logs. Extending the data model and authentication to support multiple patients and role-differentiated caregivers would enable institutional deployment. On the analytics front, sequence models could be introduced to forecast missed doses prospectively rather than detecting irregularity retrospectively, and personalised lead-time optimisation could adapt reminder timing to each patient's response history. Integration with standard clinical interoperability formats would further allow the system to exchange data with electronic health records.

X. CONCLUSION

This paper presented a portable, privacy-preserving framework that advances digital medication support from passive reminding toward active compliance monitoring. By combining timezone-aware dose scheduling, channel-redundant delivery, and unsupervised cohort-based anomaly detection within a self-hosted three-tier architecture, the system delivers both dependable reminders and interpretable behavioural insight while keeping all data on the patient's own infrastructure. Empirical evaluation demonstrated sub-100-millisecond interface latency, perfect scheduling classification across timezones, and an irregularity-detection F1 score reaching 0.94 as logging history accumulated, with a deterministic fallback ensuring stable operation in minimal environments. The framework offers a practical foundation for accessible adherence support, and its planned extensions toward objective intake confirmation, multi-patient operation, and predictive modelling point to substantial future clinical impact.

REFERENCES

[1] World Health Organization, Adherence to Long-Term Therapies: Evidence for Action. Geneva, Switzerland: WHO Press, 2003.
[2] M. T. Brown and J. K. Bussell, “Medication adherence: WHO cares?,” Mayo Clinic Proceedings, vol. 86, no. 4, pp. 304–314, 2011.
[3] B. Vrijens et al., “A new taxonomy for describing and defining adherence to medications,” British Journal of Clinical Pharmacology, vol. 73, no. 5, pp. 691–705, 2012.
[4] S. R. Majeed et al., “Effectiveness of mobile phone reminders on medication adherence: a systematic review,” Journal of Medical Internet Research, vol. 22, no. 9, e17776, 2020.
[5] A. K. Pandey and R. Sharma, “Smartphone-based reminder systems for chronic disease management,” IEEE Access, vol. 9, pp. 112045–112060, 2021.
[6] C. R. Hayes et al., “Electronic monitoring of medication adherence using smart packaging,” IEEE Sensors Journal, vol. 20, no. 14, pp. 8123–8132, 2020.
[7] P. Kumar and S. Banerjee, “Wearable sensor-driven intake detection for adherence support,” in Proc. IEEE Int. Conf. on Healthcare Informatics, 2021, pp. 245–252.
[8] L. Zhang, Y. Chen, and H. Liu, “Vision-based medication intake verification using deep learning,” IEEE J. Biomed. Health Informatics, vol. 25, no. 6, pp. 2210–2219, 2021.
[9] D. Patel et al., “Machine learning for prediction of medication non-adherence,” npj Digital Medicine, vol. 4, art. 110, 2021.
[10] R. Gupta and M. Iyer, “Risk stratification of medication non-adherence using ensemble classifiers,” IEEE Access, vol. 10, pp. 33210–33223, 2022.
[11] S. Ahmed and K. Rao, “Recurrent neural networks for forecasting missed medication doses,” in Proc. IEEE EMBC, 2022, pp. 1788–1792.
[12] F. T. Liu, K. M. Ting, and Z.-H. Zhou, “Isolationforest,” in Proc. IEEE Int. Conf. on Data Mining, 2008, pp. 413–422.
[13] V. Chandola, A. Banerjee, and V. Kumar, “Anomaly detection: a survey,” ACM Computing Surveys, vol. 41, no. 3, pp. 1–58, 2009.
[14] J. M. Silva and A. Costa, “Lightweight REST architectures for digital health applications,” IEEE Software, vol. 39, no. 2, pp. 64–71, 2022.

[15] N. Roberts and T. Williams, "Component-based front-end design for clinical dashboards," IEEE Computer, vol. 54, no. 8, pp. 40–48, 2021.
[16] H. Park and J. Lee, "Timezone-aware scheduling for distributed reminder services," IEEE Trans. Services Computing, vol. 16, no. 1, pp. 220–231, 2023.
[17] A. Fernandez et al., "Privacy-preserving self-hosted health platforms: a comparative study," IEEE Security & Privacy, vol. 21, no. 4, pp. 55–64, 2023.

AUTHORS' BIOGRAPHIES

Vaddi Satya Ganesh received the B.Sc. degree in computer Science S.V.K. P & Dr. K. S. Raju Arts and Science College (A), Penugonda in 2024, He is currently pursuing the Master of Computer Applications (MCA)degree at S.V.K.P & Dr. K.S. Raju Arts and Science College, Penugonda, West Godavari, India. His research interests including AI/ML and Python Programming

K. Lakshmi Sai Sri Working as Lecturer in S.V.K. P & Dr. K. S. Raju Arts and Science College (A), Penugonda, West Godavari District, AP. Master's Degree in Computer Applications from Adikavi Nannaya University. Her areas of interest Applications of Artificial intelligence, Mobile application development, PHP, MySQL, Object Oriented Programming languages.

INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN COMPUTER & COMMUNICATION ENGINEERING