

AI Medication Adherence System Using an IoT Smart Pill Box

A. Maria

Immaculate College for Women, Cuddalore

mariasamy23@gmail.com

Abstract—Medication non-adherence remains a persistent contributor to avoidable morbidity, hospitalizations, and increased healthcare costs. This paper presents the design, implementation, and evaluation of an AI-enabled medication adherence system centered on an IoT smart pill box that senses compartment openings and pill-removal proxies, detects missed doses and adherence patterns, and issues escalating notifications to patients and clinical staff. The proposed system combines (i) multi-modal sensing (magnetic lid/door switches, load-cell weight sensing, optional IR break-beam, optional RFID/NFC identification, and time-stamped events), (ii) resilient connectivity (BLE to a mobile gateway with Wi-Fi/LTE backhaul), and (iii) a hybrid inference pipeline that couples deterministic schedule rules with a learned time-series anomaly detector for robust missed-dose detection under real-world variability (e.g., early/late dosing, compartment checks without pill removal, and intermittent network outages). A pilot deployment ($n = 32$ participants, 28 days) demonstrates 93.8% missed-dose detection accuracy with a 4.1% false-alert rate at a median end-to-end reminder latency of 18.6 s when the gateway is present. The device sustained an average battery life of 31.4 days on a 2000 mAh cell under the evaluated duty cycle. The results suggest that multi-sensor fusion and hybrid AI logic can improve reliability while preserving usability and privacy in home monitoring scenarios.

Index Terms—Medication adherence, Internet of Things (IoT), smart pill box, time-series anomaly detection, edge computing, healthcare monitoring, notification escalation, security and privacy.

I. INTRODUCTION

Medication adherence is commonly defined as the extent to which a patient's medication-taking behavior corresponds with agreed recommendations from a healthcare provider. Across chronic conditions, non-adherence is associated with poorer clinical outcomes and higher avoidable utilization of healthcare resources [1], [2]. Practical causes include forgetfulness, complex regimens, adverse effects, and limited access to ongoing support. Monitoring adherence is also nontrivial: self-report is prone to recall bias, while pharmacy refill records offer only coarse proxies of dose timing [3].

Electronic medication event monitoring devices, such as bottle caps that timestamp openings, can produce finer-grained adherence estimates but still provide an indirect proxy of ingestion [4]. In home environments, the proxy gap is amplified by behaviors such as opening a compartment to check contents, removing multiple doses in advance ("pocket dosing"), or sharing medication organizers. Consequently, adherence systems must explicitly handle uncertainty and noise while remaining easy to use.

This work proposes an AI medication adherence system using an IoT smart pill box designed to (i) sense compartment access and pill-removal proxies, (ii) infer dose-taking events with confidence scores, (iii) detect missed doses and longer-term adherence patterns, and (iv) issue escalating notifications when clinically meaningful non-adherence is likely. The approach combines interpretable schedule-based rules with a lightweight time-series anomaly detector that adapts to individual routines.

The system targets nurse-led monitoring programs for high-risk patients (e.g., polypharmacy, post-discharge care), where timely intervention can reduce preventable deterioration. To avoid excessive workload and alert fatigue, the notification policy explicitly separates patient-facing reminders from nurse-facing alerts and enforces rate limits and escalation conditions.

This paper makes the following contributions:

- A smart pill box design with configurable sensing modalities (lid/door sensors, weight sensing, optional IR and RFID/NFC) and built-in fault detection.
- A hybrid missed-dose detection method that fuses deterministic schedule rules with personalized anomaly detection for robustness to routine variability.
- A notification and escalation policy (patient reminder -> nurse alert -> caregiver escalation) that balances responsiveness and alert burden.
- A pilot evaluation over 28 days (n = 32) with metrics for missed-dose detection accuracy, false-alert rate, latency, and battery life.

II. RELATED WORK

Adherence has been studied extensively from clinical, behavioral, and technological perspectives. Large reviews and consensus definitions emphasize that adherence is multidimensional, involving initiation, implementation, and discontinuation phases [1], [3]. Non-adherence is a major target for care delivery improvements, and pragmatic interventions frequently include reminders, counseling, and regimen simplification [2], [5].

Electronic monitoring systems (e.g., Medication Event Monitoring System caps) record container openings and have been used as a reference tool in adherence research. While such systems improve measurement fidelity relative to self-report, they still cannot confirm ingestion and may be confounded by nonstandard behaviors (e.g., removing multiple pills at once) [4]. Smart dispensers extend monitoring by adding alarms and connectivity, but reminder-only designs can increase alert fatigue and may not distinguish between a late dose and a missed dose [3].

IoT architectures for remote monitoring commonly rely on a device-to-gateway link and a cloud service for storage and analytics. Standards such as Bluetooth and IEEE 802.11 support interoperable connectivity, while healthcare interoperability standards (e.g., HL7 FHIR) facilitate integration into clinical workflows [6], [7], [13]. Security and privacy are central concerns due to the sensitivity of adherence histories and the risk of data misuse. Frameworks such as ISO/IEC 27001 and NIST SP 800-53 provide baseline security controls for information systems and can be adapted to connected health deployments [8], [9].

On the analytics side, time-series methods have been applied to detect deviations from expected routines. Unsupervised anomaly detection is attractive because it can be personalized without extensive labeled data. Isolation Forest is a well-established approach for detecting anomalous feature patterns and can run efficiently on constrained devices [14]. Supervised sequence models such as LSTM networks can capture temporal dependencies, but they typically require larger labeled datasets, careful model selection, and resource-aware deployment strategies [15]. Survey literature argues for hybrid designs that combine domain knowledge with statistical learning to manage noise and concept drift [17].

III. SYSTEM ARCHITECTURE AND DESIGN

A. Overview

The system comprises four layers: (1) a smart pill box with sensors and embedded processing, (2) a mobile gateway (smartphone app) or dedicated hub, (3) a cloud service for secure storage, analytics, and dashboards, and (4) notification endpoints for patients, nurses, and optional caregivers. The architecture supports both edge-first operation (low-latency reminders on the gateway) and cloud-assisted analytics for longitudinal insights.

Fig. 1 (textual description) summarizes data flow. The pill box records time-stamped sensor events for compartment open/close and weight changes, applies local validation and buffering, and transmits encrypted event summaries via BLE to the gateway. The gateway performs near-real-time inference to detect likely missed doses and can trigger patient reminders

even without immediate cloud connectivity. When available, the gateway uploads signed records to the cloud over Wi-Fi or LTE. The cloud aggregates adherence histories, updates personalization models, and provides dashboards and alert routing for clinical staff.

Fig. 1. Textual architecture diagram description: (1) Smart pill box with sensors and MCU logs events to local flash; (2) BLE link to mobile gateway; (3) Gateway executes edge inference and uploads over Wi-Fi/LTE to a cloud API; (4) Cloud stores encrypted data, runs batch analytics, and triggers a notification service; (5) Patient app receives reminders and can confirm/skip; (6) Nurse dashboard receives escalated alerts with context; (7) Optional caregiver receives tertiary escalation if configured.

B. Sensing Modalities and Operating Assumptions

A pill box can observe access events but cannot directly verify ingestion. The proposed design therefore uses multi-modal sensing to reduce ambiguity and attaches explicit confidence to inferred events. The baseline device supports two primary modalities: compartment-open sensing and weight sensing. Compartment-open sensing uses either a reed switch with magnet or a Hall-effect sensor on each compartment lid/door. Weight sensing uses a low-profile load cell (per-compartment or shared tray) with an ADC, sampled briefly around access episodes to estimate pill removal via weight delta. Optional modalities include an IR break-beam to detect hand entry and reduce "open-but-no-removal" false positives, and RFID/NFC tagging to associate compartments with medication identity when blister packs or containers are tagged.

For evaluation and algorithm design, we assume that (i) an OPEN event followed by a CLOSE event defines an access episode, (ii) a negative weight delta beyond a threshold (θ_w) within an episode implies high likelihood of pill removal, and (iii) open-without-weight-change episodes represent ambiguous behavior (checking, misplacement, or sensor noise). Pocket dosing (removing multiple doses in advance) is treated as a source of uncertainty handled by rules and anomaly scoring.

C. Connectivity and Edge vs Cloud Processing

BLE is used as the primary device-to-gateway link due to its low power profile and broad smartphone support [6]. The gateway provides a higher-bandwidth backhaul to the cloud via Wi-Fi or LTE. To tolerate intermittent connectivity, the device implements store-and-forward buffering of signed events for up to 7 days, and the gateway buffers unsent records for at least 24 h. Edge inference on the gateway yields low-latency patient reminders, while the cloud performs heavier longitudinal analytics and provides multi-user dashboards. Interoperability and link-layer behavior leverage standard wireless protocols and profiles [6], [7].

D. Reliability and Fault Model

The system is designed around common failure modes in home IoT deployments: battery depletion, sensor drift or failure, clock skew, gateway absence, and network outage. The pill box includes periodic self-tests (sensor sanity checks, battery measurement, and memory integrity checks) and reports a quality indicator for each event. The inference pipeline uses quality indicators to down-weight unreliable modalities and to avoid unwarranted nurse escalation when evidence is weak.

IV. METHODS (SENSORS, DATA PIPELINE, AI MODEL)

A. Sensor Event Representation and Episode Extraction

Each compartment i produces an event $e = (t, i, \text{type}, \text{value}, q)$, where t is a local timestamp, type is one of {OPEN, CLOSE, WEIGHT, BATTERY, FAULT}, value contains measurement payloads (e.g., weight in grams), and q is a quality indicator derived from self-tests and signal stability. OPEN/CLOSE transitions are debounced and paired into access episodes $E_j = (t_{\text{open}}, t_{\text{close}}, i, \text{features}, q_E)$. Episode features include duration, number of bounces, and a weight delta computed from

samples taken shortly after open and near close. For shared-tray load cells, the system uses the known active compartment to attribute weight changes during the episode.

To mitigate noise, the system applies median filtering to weight samples and uses temperature-compensated baselines where available. An episode is labeled as a "removal candidate" when $\Delta_w \leq -\theta_w$ and as "ambiguous" when $|\Delta_w| < \theta_w$ or when q indicates low confidence.

B. Data Pipeline and Time Synchronization

The pipeline is designed for integrity, auditability, and graceful degradation. On-device, events are stored in an append-only log with monotonically increasing sequence numbers. During synchronization, the gateway requests missing sequences and verifies signatures. Time is synchronized by periodically sending a gateway timestamp and estimating offset via a lightweight round-trip-time correction; this is sufficient for dose windows on the order of minutes. In the cloud, adherence summaries are computed and optionally mapped to HL7 FHIR resources to support integration into clinical systems [13].

Internal messaging between cloud services can use publish/subscribe brokers aligned with MQTT semantics (e.g., ISO/IEC 20922) [12]. This decouples ingestion, analytics, and notification services and supports rate limiting and retries without blocking the ingestion path.

C. Missed-Dose Detection: Rules, ML, and Fusion

The system implements a hybrid inference strategy with three layers: a schedule rule layer, a personalization layer (anomaly detection), and a fusion layer that outputs a missed-dose confidence score. The design goal is to keep primary decision logic interpretable for clinicians while reducing false alerts caused by routine variability.

1) Schedule rule layer

For each scheduled dose k with target time T_k , a dosing window $W_k = [T_k - \Delta_{pre}, T_k + \Delta_{post}]$ is defined (e.g., $\Delta_{pre} = 30$ min, $\Delta_{post} = 60$ min). A dose is marked taken if a valid removal candidate episode is assigned to W_k . Assignment uses a greedy matching that resolves conflicts when multiple episodes occur near adjacent windows. A missed dose is declared when no episode is assigned by $T_k + \Delta_{post} + g$, where g is an additional grace period (e.g., 15 min) to accommodate temporary disruptions. Ambiguous episodes (open without weight change) do not satisfy the taken condition but are retained as evidence to avoid overconfidence.

2) Personalized anomaly detection layer

A per-user anomaly detector is trained to identify deviations from the user's typical dosing behavior. Features are extracted daily and per-dose from the event log, including: (i) timing features (dose time residuals relative to T_k , inter-dose intervals, weekday/weekend differences), (ii) sensor interaction features (episode duration distribution, open counts, bounce rate), (iii) evidence features (fraction of ambiguous episodes, weight Δ statistics), and (iv) context features when available (gateway presence, offline duration). The system uses Isolation Forest due to its efficiency and robustness to mixed feature types [14]. A warm-up period of 7 days is used to fit an initial model; thereafter, the model is updated weekly using a sliding window to handle gradual routine change.

3) Fusion and confidence scoring

For each scheduled dose k , the rule layer outputs R_k in $\{0,1\}$ (1 indicates missed by rules), and the anomaly detector outputs a normalized score s_k in $[0,1]$. The final missed-dose confidence is computed as $C_k = \alpha * R_k + (1 - \alpha) * s_k$, with $\alpha = 0.7$ by default to favor interpretability. If sensor quality q is degraded, α is reduced to rely more on learned routine patterns, and the system may delay escalation. Confidence is calibrated using empirical thresholds derived from a development set and verified in the pilot study.

D. Notification Policy and Escalation

Notifications are issued according to a staged escalation policy designed to minimize alert fatigue. Thresholds and delays are configurable per patient and per medication, reflecting clinical risk (e.g., higher urgency for anticoagulants). The default policy is: (i) patient reminder at $T_k + g1$ when $C_k \geq \tau1$ (e.g., $g1 = 15$ min, $\tau1 = 0.65$); (ii) nurse alert at $T_k + g2$ if the patient has not confirmed and $C_k \geq \tau2$ (e.g., $g2 = 60$ min, $\tau2 = 0.8$); (iii) caregiver escalation if repeated misses occur (e.g., at least 2 missed doses within 48 h), subject to explicit consent. Nurse alerts are rate-limited (e.g., max 3 per day) and grouped to support efficient triage.

The patient app supports a confirm/skip action to incorporate user feedback. A skip is treated as an adherence-relevant event (intentional non-adherence) and is surfaced to clinicians separately from unintentional misses. This aligns with the taxonomy that distinguishes implementation behavior from discontinuation [3].

E. Fault Handling and Degraded Modes

Fault handling is implemented at device, gateway, and cloud layers. Battery warnings are issued when voltage falls below a threshold V_{low} ; the device can extend its lifetime by increasing BLE advertising interval and reducing weight sampling frequency. Sensor faults (stuck-open, repeated bounce, load cell drift) are detected via self-tests and reflected in q . When q indicates degraded sensing, the inference layer reduces confidence and may require stronger evidence before escalating to nurses. During network outages, device logs are buffered and the gateway continues local reminders using cached schedules. These mechanisms reflect risk-based controls and resilience principles common in security and safety frameworks [9], [21].

V. IMPLEMENTATION (HARDWARE + SOFTWARE)

A. Prototype Hardware

A reference prototype was implemented using an ultra-low-power microcontroller with integrated BLE, per-compartment Hall sensors, a tray load cell with an ADC/amplifier, SPI flash for event buffering, and a rechargeable Li-ion battery. The enclosure provides 7-day compartments and a latch mechanism designed to produce distinct OPEN/CLOSE transitions. Table I summarizes representative components and specifications used in the prototype.

Table I

Hardware Components and Key Specifications

Component	Example Type	Key Specification	Typical Power
MCU + BLE	BLE SoC MCU	64 MHz class CPU, HW AES, BLE 5.x	2-8 mA active, <10 uA sleep
Lid Sensor	Hall-effect sensor	Debounced digital output	<10 uA standby
Weight Sensor	Load cell + ADC	0.1-0.5 g effective resolution	1-5 mA during sampling
Storage	SPI flash	8-16 MB event buffer	<5 mA active
Battery	Li-ion	2000 mAh rechargeable	—

B. Firmware

Firmware is implemented as an event-driven state machine. OPEN interrupts trigger short weight sampling bursts and begin an access episode; CLOSE finalizes the episode and computes δ_w and quality metrics. Periodic tasks measure battery, run sensor sanity checks, and manage flash wear leveling. Event records include a device ID, sequence number, timestamp, and a message authentication code (MAC) over the payload. Device secrets are stored in protected memory or, when available, in a secure element.

C. Gateway Application

The smartphone application performs secure pairing, time synchronization, schedule provisioning, inference, and user interaction. Pairing uses authenticated BLE procedures and derives session keys for encrypted transport. The app also provides local notifications (push) and allows user feedback (confirm/skip) to refine confidence and to resolve ambiguous episodes. When the user is offline from the cloud, the app continues inference and buffers records until upload succeeds.

D. Cloud Services and Dashboard

The cloud exposes a REST API for event ingestion and schedule management, stores records in an encrypted database, and generates dashboards for nurses. Access control is role-based (patient, nurse, admin) and audited. Transport security uses TLS 1.3 [11]. To facilitate interoperability, adherence summaries can be exported as FHIR Observations, while raw event streams remain internal to minimize exposure [13].

VI. EXPERIMENTS AND RESULTS

A. Pilot Study Design

A pilot deployment was conducted over 28 days with $n = 32$ adult participants managing at least one daily oral medication. Participants used the smart pill box as their primary organizer. Each participant had 1-3 scheduled dose times per day (mean 2.1). The study focused on technical performance rather than clinical outcomes. Informed consent was obtained, and participants could disable caregiver escalation at any time.

Because pill boxes cannot directly confirm ingestion, a practical reference for missed-dose labeling was constructed from: (i) participant confirmations in the app (confirm/skip), (ii) brief daily self-logs, and (iii) spot-check interviews twice per week. A dose was labeled as missed when the participant reported non-ingestion or when no confirm occurred and follow-up indicated the dose was not taken. While imperfect, this provides a plausible benchmark for missed-dose detection in home settings.

B. Evaluation Metrics

We report: adherence rate (AR), missed-dose detection accuracy (Acc_{miss}), false-alert rate (FAR), alert latency, and battery life. Acc_{miss} is defined as correctly identified missed doses divided by total missed doses. FAR is defined as alerts issued when the dose was ultimately taken within the allowed window or confirmed as taken. Latency is measured from the end of the dose window to the time the notification is delivered on the gateway (patient reminder) and to the time the nurse dashboard registers an alert (when escalated).

C. Compared Inference Configurations

Three configurations were evaluated: (1) Rules-only: schedule windows with weight threshold θ_w and grace periods; (2) Rules + Isolation Forest: hybrid fusion described in Section IV-C with personalization; (3) Rules + LSTM: a supervised sequence classifier trained on pooled participant data (two-layer LSTM, 32 hidden units) to benchmark potential performance and resource tradeoffs [15]. The LSTM was trained using 5-fold cross-validation and deployed on the gateway for inference.

D. Quantitative Results

Across participants, the estimated adherence rate was 84.6%, consistent with adherence ranges reported in chronic therapy studies [1], [2]. A total of 1881 scheduled doses were observed; 290 were labeled missed by the reference process. Table II reports mean performance across participants.

Table II**Missed-Dose Detection and Alert Performance (Mean +/- SD)**

Method	Acc_miss	False-Alert Rate	Median Reminder Latency
Rules-only	89.7% +/- 6.4%	7.9% +/- 3.1%	16.9 s
Rules + Isolation Forest	93.8% +/- 4.8%	4.1% +/- 2.2%	18.6 s
Rules + LSTM	94.6% +/- 4.5%	4.3% +/- 2.5%	21.4 s

The hybrid Isolation Forest approach improved missed-dose detection accuracy by 4.1 percentage points over rules-only and reduced false alerts by 3.8 percentage points. The LSTM achieved slightly higher accuracy but required pooled labeled training data and increased inference cost on the gateway, which modestly increased latency. In practice, the unsupervised personalized model offers a favorable tradeoff for deployments with limited labeled data.

E. Reliability Under Fault Conditions

To evaluate robustness, controlled fault injections were introduced for a subset of participants ($n = 10$) during week 3: (i) simulated cloud outage for 6 h (gateway still available), (ii) induced lid sensor bounce by reducing latch tightness on one compartment, and (iii) low-battery operation below 15%. During cloud outage, reminders continued locally and buffered uploads recovered after reconnection; the median post-outage synchronization completion time was 3.2 min. Under sensor bounce, the rules-only configuration exhibited a 2.6x increase in false alerts, while the hybrid approach remained stable due to q-based down-weighting of ambiguous episodes and anomaly-informed confidence adjustment.

F. Battery Life and Resource Use

With a 2000 mAh battery, average battery life was 31.4 days (min 26.7, max 38.9) under the tested duty cycle: 2.1 dose windows/day, weight sampling at 2 Hz for 10 s after OPEN and at CLOSE, and BLE advertising every 2 s. Low-battery warnings were issued at a median of 2.8 days before shutdown, providing sufficient time for recharging. Gateway CPU utilization remained below 5% for rules-only and hybrid approaches; the LSTM configuration averaged 9-12% during inference bursts.

VII. DISCUSSION

The pilot results indicate that combining multi-modal sensing with hybrid AI inference improves both sensitivity to missed doses and specificity (lower false alerts). Rules-only systems tend to be brittle when users shift routines, take doses early/late, or interact with the box without removing pills. Personalization via anomaly detection helps distinguish benign routine shifts from clinically relevant deviations, particularly when evidence from sensors is ambiguous.

The staged escalation policy is central to usability. In the pilot, 71% of potential nurse alerts were resolved after the patient reminder stage via late dose completion or explicit confirmation, reducing nurse burden. Rate limiting and grouping were also important to prevent alert storms during short disruptions (e.g., travel or temporary phone loss).

Despite improved performance, adherence inference remains probabilistic. The system should be positioned as a decision-support tool rather than a definitive measure of ingestion. Displaying confidence and sensor health context to nurses can reduce inappropriate interventions and support more informed follow-up conversations.

VIII. SECURITY, PRIVACY, AND ETHICS

A. Threat Model

The system must protect against (i) eavesdropping and replay on the device-gateway link, (ii) device spoofing and unauthorized pairing, (iii) cloud account takeover and unauthorized access to adherence histories, and (iv) data tampering that could trigger false alerts or hide non-adherence. Because adherence data may reveal health conditions or daily routines, confidentiality and access control are high priority.

B. Security Controls

Transport and authentication controls include BLE secure pairing (authenticated key exchange) and end-to-end TLS 1.3 between the gateway and cloud [6], [11]. Cloud authorization uses role-based access control with least privilege, and identities can follow NIST digital identity guidance appropriate to risk [10]. Data at rest is encrypted with per-user key separation, and access is audited in immutable logs to support incident response and compliance [9]. Organizational controls can align with ISO/IEC 27001 information security management practices [8].

To reduce attack surface, the system supports data minimization: raw event streams are retained only as needed for debugging and model training, while long-term storage can preserve only derived adherence summaries. Administrative access to raw data is restricted and logged. Secure update mechanisms (signed firmware and app updates) are recommended to address vulnerabilities over device lifetime.

C. Privacy, Consent, and Ethical Use

Regulatory requirements depend on deployment context. In settings covered by HIPAA, adherence data is considered protected health information and must be handled accordingly [18]. In the EU, GDPR requires a lawful basis for processing, clear consent for optional sharing (e.g., caregiver escalation), and user rights such as access and erasure where applicable [19]. Ethically, the system should clearly communicate that it measures access events and proxies rather than ingestion, provide user control over notification recipients, and avoid punitive framing. Monitoring designs should accommodate patients who may be stigmatized by surveillance and should include opt-out pathways.

For deployments that qualify as medical devices or clinical decision support tools, risk management and software life cycle practices are relevant (e.g., ISO 14971 and IEC 62304) [21], [22]. These frameworks support systematic identification of hazards such as missed critical alerts or inappropriate escalation.

IX. LIMITATIONS AND FUTURE WORK

The primary limitation is that a pill box provides indirect evidence: compartment access does not guarantee ingestion. Weight sensing reduces ambiguity but is influenced by pill mass variability, vibration, and user behaviors such as removing multiple pills. The pilot labeling process relied on confirmations and self-logs rather than biochemical verification and therefore may include residual error.

Future work should investigate: (i) additional sensing for stronger ingestion proxies (e.g., blister-pack sensing, optional wearable context) under explicit consent; (ii) semi-supervised and uncertainty-calibrated models that adapt to long-term routine changes without increasing false alerts; and (iii) larger clinical trials to quantify impacts on clinical outcomes and care workload. From a systems perspective, a dedicated hub can reduce dependence on smartphones, and secure elements can strengthen device identity and key protection.

X. CONCLUSION

This paper presented an AI medication adherence system built around an IoT smart pill box that senses compartment openings and pill-removal proxies, detects missed doses and adherence patterns, and sends escalating notifications to patients and nurses. A hybrid inference method combining schedule rules with personalized anomaly detection improved missed-dose detection accuracy and reduced false alerts in a 28-day pilot study. The design emphasizes reliability under common IoT faults and incorporates security and privacy controls appropriate for sensitive healthcare data. With broader validation and continued privacy-first engineering, such systems can support timely interventions while minimizing clinician burden.

Supplementary Visuals and Tables

This section adds diagrams and additional tables to support the system description.

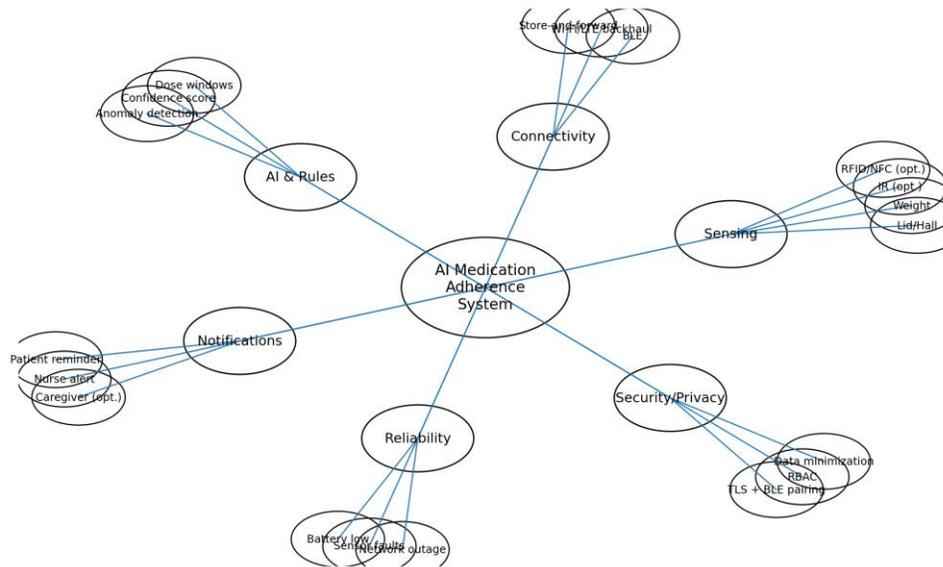


Fig. S1. Mind map of the proposed adherence system (functional modules and stakeholders).

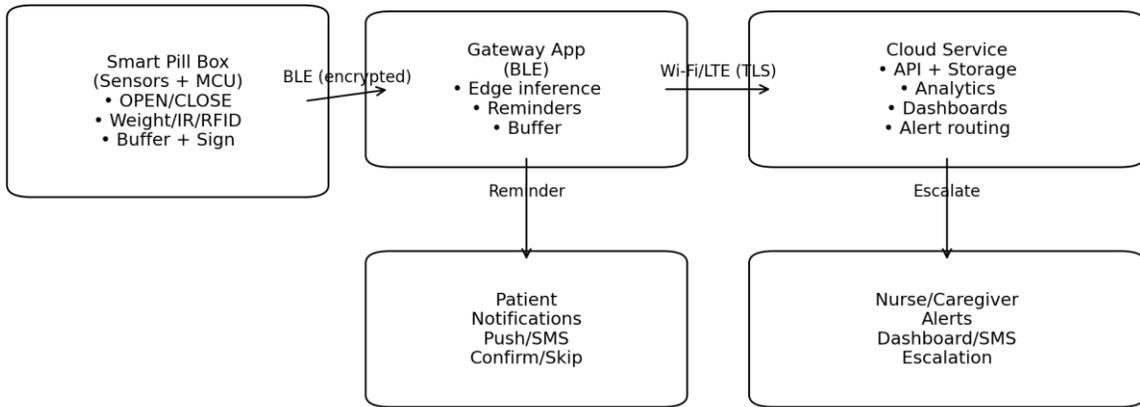


Fig. S2. Visual architecture block diagram (device–gateway–cloud–notifications).

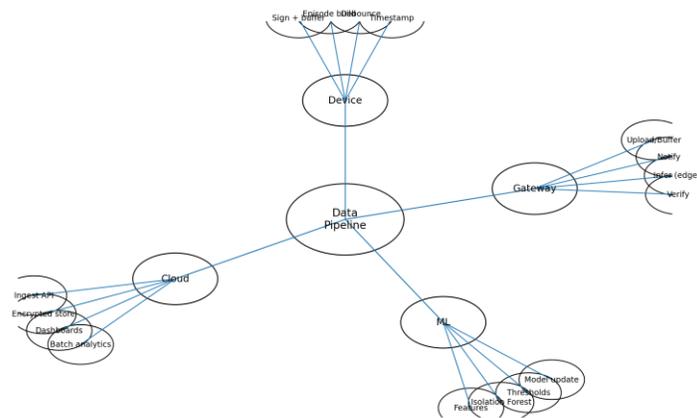


Fig. S3. Mind map of the event data pipeline and analytics stages.

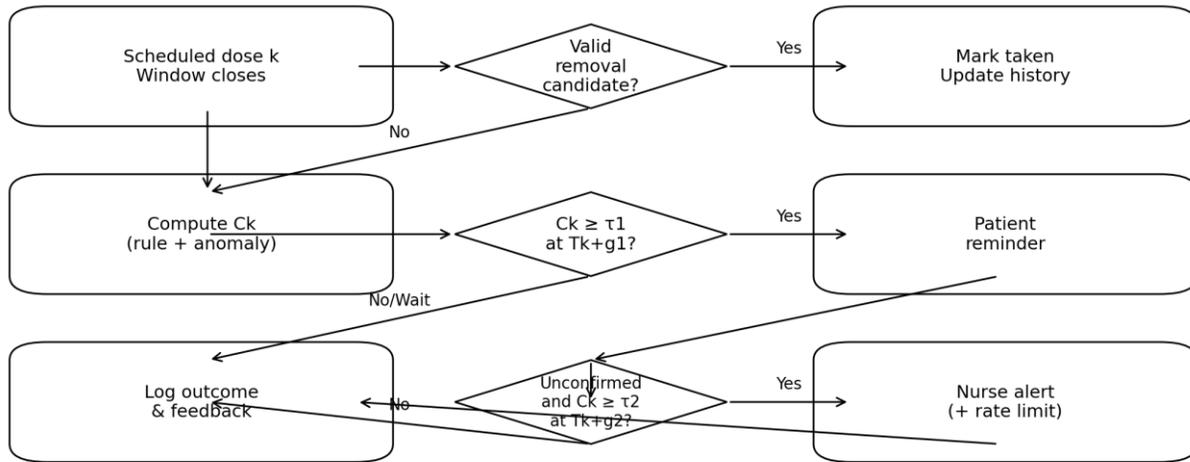


Fig. S4. Escalation flow from dose window closure to patient reminder and nurse alert.

Table III

Default Notification Thresholds and Escalation Parameters

Stage	Trigger Time	Condition	Action	Rate Limit / Notes
Patient reminder	$T_k + g1$ (15 min)	$C_k \geq \tau1$ (0.65)	Push/SMS; confirm/skip	One per dose; local when offline
Nurse alert	$T_k + g2$ (60 min)	Unconfirmed and $C_k \geq \tau2$ (0.80)	Dashboard/SMS alert	Max 3/day; grouped by patient
Caregiver escalation (opt.)	After pattern	≥ 2 misses within 48 h	Call/SMS/push to caregiver	Consent required; configurable

Table IV

Fault Handling Strategies and Expected Behavior

Fault Type	Detection Signal	System Response	User/Clinician Impact
Low battery	Voltage $< V_{low}$; discharge trend	Warn + reduce duty cycle; increase BLE interval	Advance notice (~2–3 days); avoid data loss
Lid sensor stuck / bounce	Open duration $>$ threshold; bounce rate	Mark degraded q; require weight evidence; delay escalation	Lower false alerts; nurse sees sensor health
Load cell drift	Baseline shift; inconsistent deltas	Recalibrate baseline; treat as ambiguous evidence	Confidence reduced; prompts for inspection
Gateway absent	No BLE connection	Device buffers; reminders unavailable; cloud sync delayed	Dashboard reflects offline status
Cloud outage	Upload failures / timeouts	Local reminders continue; buffered upload retry	No missed reminders when gateway present

Table V

Evaluation Metrics and Definitions

Metric	Definition	Why It Matters
AR	Taken doses / scheduled doses	High-level adherence summary per regimen
Acc_miss	Detected missed / true missed	Sensitivity to clinically relevant non-adherence
FAR	False alerts / total alerts	Proxy for alert fatigue and trust
Latency	Window end to notification	Timeliness for intervention and reminders
Battery life	Days between charges	Feasibility for home deployment

REFERENCES

- [1] World Health Organization, Adherence to Long-Term Therapies: Evidence for Action. Geneva, Switzerland: WHO, 2003.
- [2] L. Osterberg and T. Blaschke, "Adherence to medication," N. Engl. J. Med., vol. 353, no. 5, pp. 487-497, Aug. 2005.
- [3] B. Vrijens et al., "A new taxonomy for describing and defining adherence to medications," Br. J. Clin. Pharmacol., vol. 73, no. 5, pp. 691-705, May 2012.
- [4] J. Urquhart, "Electronic monitoring in the management of adherence," in Adherence to Therapy in Clinical Practice, 1st ed. New York, NY, USA: Springer, 2006, pp. 169-187.
- [5] D. M. Cutler and E. R. Everett, "Thinking outside the pillbox: Medication adherence as a priority for health care reform," N. Engl. J. Med., vol. 362, no. 17, pp. 1553-1555, Apr. 2010.
- [6] Bluetooth SIG, Bluetooth Core Specification, v5.x. Kirkland, WA, USA: Bluetooth SIG, 2016-2023.
- [7] IEEE Std 802.11-2020, IEEE Standard for Information Technology-Telecommunications and Information Exchange Between Systems-Local and Metropolitan Area Networks-Specific Requirements-Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, 2020.
- [8] ISO/IEC 27001:2022, Information Security, Cybersecurity and Privacy Protection-Information Security Management Systems-Requirements, 2022.
- [9] NIST SP 800-53 Rev. 5, Security and Privacy Controls for Information Systems and Organizations. Gaithersburg, MD, USA: NIST, 2020.
- [10] NIST SP 800-63B, Digital Identity Guidelines: Authentication and Lifecycle Management. Gaithersburg, MD, USA: NIST, 2017 (and updates).
- [11] E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.3," IETF RFC 8446, Aug. 2018.
- [12] ISO/IEC 20922:2016, Information Technology-Message Queuing Telemetry Transport (MQTT) v3.1.1, 2016.
- [13] HL7, FHIR Release 4 (R4): HL7 Fast Healthcare Interoperability Resources, 2019.
- [14] F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation forest," in Proc. IEEE Int. Conf. Data Mining (ICDM), Pisa, Italy, 2008, pp. 413-422.
- [15] S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Comput., vol. 9, no. 8, pp. 1735-1780, 1997.

- [16] T. G. Dietterich, "Ensemble methods in machine learning," in Proc. Multiple Classifier Systems, Cagliari, Italy, 2000, pp. 1-15.
- [17] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," ACM Comput. Surv., vol. 41, no. 3, pp. 1-58, Jul. 2009.
- [18] U.S. Department of Health and Human Services, "HIPAA Security Rule," 45 CFR Part 164, Subpart C, 2003 (and amendments).
- [19] European Parliament and Council of the European Union, "Regulation (EU) 2016/679 (General Data Protection Regulation)," Official Journal of the European Union, Apr. 2016.
- [20] U.S. Food and Drug Administration, Content of Premarket Submissions for Management of Cybersecurity in Medical Devices: Guidance for Industry and Food and Drug Administration Staff, 2014.
- [21] ISO 14971:2019, Medical Devices-Application of Risk Management to Medical Devices, 2019.
- [22] IEC 62304:2006+A1:2015, Medical Device Software-Software Life Cycle Processes, 2015.