International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 10 | Oct - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

Decentralized Epidemic Tracking using Federated Learning and Edge AI on Consumer Wearables

Prof. S. Abikayil Aarthi
Assistant Professor,
Department of Computer Science and Engineering,
Kings College of Engineering,
Punalkulam, TN, India,
aarthi.cse@kingsengg.edu.in

Hariharan. N
Department of Computer Science and Engineering,
Kings College of Engineering,
Punalkulam, TN, India,
hariharanhari102004@gmail.com

Ajay. V
Department of Computer Science and Engineering,
Kings College of Engineering,
Punalkulam, TN, India,
ajayvijayakumar710@gmail.com

Gurudev N
Department of Computer Science and Engineering,
Kings College of Engineering,
Punalkulam, TN, India,
gurudevneps@gmail.com

Deerga Tharsan. A. V
Department of Computer Science and Engineering,
Kings College of Engineering,
Punalkulam, TN, India,
deerga0607@gmail.com

Chandra Mukilan V
Department of Computer Science and Engineering,
Kings College of Engineering,
Punalkulam, TN, India,
chandramukilanv@gmail.com

Aravinth R
Department of Computer Science and Engineering,
Kings College of Engineering,
Punalkulam, TN, India,
aravinthramamoorthi57@gmail.com

ABSTRACT

Early detection of infectious-disease outbreaks depends on how quickly weak, distributed signals can be turned into a public-health decision, yet conventional surveillance relies on clinic visits, laboratory confirmation, and syndromic reporting that lag infection by days. Consumer wearables such as smartwatches, fitness bands, and smart rings continuously record resting heart rate, heart-rate variability, skin temperature, blood-oxygen saturation, respiration, sleep, and activity, and prior studies show that deviations from a person's own baseline can appear before symptoms. Centralizing this data, however, creates serious privacy, security, and regulatory exposure. This paper proposes a decentralized epidemic-tracking framework that combines Federated Learning (FL) and Edge AI on consumer wearables. Personalized anomaly detectors run on the device; only privacy-protected model updates and coarse, thresholded anomaly flags leave the user's phone; and a hierarchical wearable, smartphone, regional-edge, and health-authority architecture converts local signals into population-level outbreak indicators. Building on a recent wearable-centric review of federated AI, the framework adds eight design innovations: personalized baseline-deviation modelling, hierarchical hybrid aggregation, outbreak-risk-adaptive differential privacy, energy-aware client participation, federated analytics with secure aggregation and k-anonymous geo-cells, a permissioned-ledger audit trail kept off the critical path, federated unlearning for erasure rights, and a pathogen-agnostic cold-start strategy. Six learning approaches are compared conceptually, and detection, timeliness, communication, energy, and privacy metrics are defined. As no field deployment or labelled dataset was collected for this study, the work is presented as a proposed framework and illustrative experimental design rather than a report of empirical results, and all performance figures shown are clearly marked as illustrative placeholders. The discussion covers practical significance for public-health agencies, limitations such as sensor noise, non-IID data, client dropout, dataset shift, and regulatory classification, and the future directions needed to move from prototype to field-ready deployment.

Keywords: Federated Learning, Edge AI, Consumer Wearables, Epidemic Surveillance, Privacy Preservation, Anomaly Detection, Differential Privacy, Digital Public Health

I. INTRODUCTION

A. Background of Epidemic Surveillance

The speed of an outbreak response is set by how early the first signals are noticed. Traditional surveillance combines clinical case reports, laboratory confirmation, sentinel clinics, and syndromic indicators such as pharmacy sales or school absenteeism. These sources are valuable but lag infection, depend on people seeking care, and under-represent mild or asymptomatic cases, which are precisely the cases that drive early transmission. Epidemic dynamics also vary continuously in space and time, so a useful tracker needs frequent, geographically resolved, population-wide observation rather than periodic reports.

B. Limitations of Conventional and Centralized Approaches

Facility-based reporting is slow, uneven across regions, and difficult to standardize. A natural alternative is to collect wearable data in a central cloud and train a single model, but this introduces three problems. First, physiological and location data are highly sensitive, and centralized collection increases breach risk and conflicts with regulations such as HIPAA and the GDPR. Second, continuous raw streaming drains battery and bandwidth. Third, centralized pipelines create a single point of failure and weaken public trust, which directly limits participation. Federated Learning addresses part of this by keeping raw data on local devices and sharing only model updates [1], [11].

C. Role of Consumer Wearables

Smartwatches, fitness trackers, smart rings, and earbuds are already worn by large populations and measure heart rate, blood oxygen, sleep, temperature, respiration, and activity [1]. Studies using such devices reported that physiological deviations can precede symptom onset for COVID-19 and for influenza and the common cold [2], [3], [4]. Wearables therefore provide a dense, time-continuous, passively collected signal that also reaches people who never visit a clinic.

D. Role of Federated Learning and Edge AI

FL trains a shared model across many devices without moving the data, and Edge AI runs inference and light training close to where the data is produced, reducing latency and dependence on connectivity [11]. Together they allow personalized, privacy-preserving detection on the device and population-level learning across devices. The Smart Health review [1] shows that in wearable settings the real difficulty is system-level: strict energy limits, noisy high-frequency data, intermittent connectivity, and the need for real-time response force trade-offs among privacy, communication, convergence, and device sustainability.

E. Research Problem

Existing work either studies single-person infection detection from wearables, treats federated learning as a generic privacy tool, or surveys healthcare FL without an epidemic-surveillance objective. There remains a need for a clearly specified, end-to-end framework that (i) detects infection-related anomalies on the device, (ii) aggregates them into population-level outbreak signals without exposing individuals, and (iii) respects wearable energy, connectivity, and regulatory constraints.

F. Research Objectives

1. To review the role of wearables, federated learning, and edge intelligence in infectious-disease surveillance and identify gaps in the literature.
2. To design a hierarchical, privacy-preserving architecture for decentralized epidemic tracking.
3. To define an on-device preprocessing and feature-engineering scheme for noisy wearable streams.
4. To describe and conceptually compare six learning approaches for infection-anomaly detection.
5. To introduce innovative mechanisms for privacy, efficiency, trust, and compliance specific to wearable FL.
6. To specify evaluation metrics covering detection, timeliness, communication, energy, and privacy.
7. To identify practical limitations, regulatory considerations, and future research directions.

G. Contributions of the Study

- A consolidated framework for decentralized epidemic tracking on consumer wearables, covering the full pipeline from sensing to public-health decision support.
- Eight proposed design innovations (Section III-G) addressing gaps the wearable FL literature leaves open.
- A structured comparison of six learning approaches for noisy, non-IID wearable data.
- A literature comparison table of wearable-based infection detection and federated healthcare studies.
- An explicit, clearly labelled evaluation design that can be populated with real data, avoiding unsupported claims of accuracy.

II. LITERATURE REVIEW

This section reviews prior work on wearable-based infection detection, federated learning for wearables, edge intelligence, and trust and compliance mechanisms, and identifies the gap that the proposed framework addresses.

A. Wearable-Based Infection Detection

Mishra et al. [2] showed that smartwatch data can reveal COVID-19 before symptom onset through resting-heart-rate and related deviations. Quer et al. [3] combined wearable data with self-reported symptoms to improve COVID-19 detection. Grzesiak et al. [4] studied noninvasive wearable sensors for detecting influenza and the common cold before symptom onset. Cacheda et al. [5] proposed early detection of infectious disease through personalized baseline-deviation modelling, emphasizing that individual baselines matter more than population thresholds. A further study reported that an infection-prediction model trained on inpatient data could predict out-of-hospital COVID-19 infection from wearable data once dataset shift was controlled, whereas applying the pretrained model directly performed near chance (AUROC about 0.52) [32]; this is central to the cold-start problem at the start of a new outbreak. Generative digital-twin augmentation with streaming anomaly detection has also been proposed for early infection detection from heart-rate and activity data [31].

B. Federated Learning for Wearables and Health Devices

Ranjbaran et al. [1] reviewed federated AI for wearables along four dimensions (privacy, sensors, cloud computing, blockchain) and an 11-feature taxonomy over 28 technical papers. Their main findings are that privacy dominates the literature, that performance gains often come from architecture and optimizer co-design rather than from FL alone, and that most frameworks remain at prototype stage under stable connectivity. Can and Ersoy [14] applied privacy-preserving federated deep learning to wearable stress monitoring. Hu et al. [16] proposed FedIERF, a federated incremental extremely random forest for wearable monitoring. Aminifar et al. [17] studied privacy-preserving edge FL for mobile health with 256 Hz EEG, and Kotiyal et al. [18] used federated Siamese networks for activity segregation at 20 Hz. Zhang et al. [23] addressed class imbalance in resource-constrained edge sensing, which is directly relevant to rare outbreak events. Sachin et al. [22] and Ghosh and Ghosh [28] built federated remote-healthcare and elderly-care systems on edge-IoMT.

C. Edge AI and Cloud-Edge Architectures

Badidi [11] reviewed Edge AI for early detection of chronic disease and the spread of infectious disease, highlighting local processing for latency and privacy. Putra et al. [19] reviewed cloud-edge AI for wearable personal health monitoring. Hierarchical edge-cloud designs were found to reduce communication rounds and localize aggregation relative to flat topologies [1].

D. Privacy, Blockchain, and Compliance

Secure aggregation [7] and differential privacy [8] are the standard cryptographic and statistical protections for federated updates, and Kairouz et al. [27] survey open problems. Blockchain-FL hybrids provide auditability: Baucas et al. [12] and Farooq et al. [13] use permissioned designs, and Ravi Shanker Reddy and Beena [15] use consent-based federated transfer learning for COVID-19 prediction on wearables. The review [1] observes that permissionless chains add confirmation latency and energy overhead unsuited to real-time monitoring, that administrative and physical HIPAA safeguards are essentially unaddressed, and that GDPR erasure, portability, and software-as-a-medical-device (SaMD) classification are largely unstudied. Machine unlearning [10] is the nascent technical route to erasure.

E. Literature Comparison Table

Table I summarizes representative studies, organized by method, data source, key finding, and principal limitation.

Author/Year	Method / Data	Key Findings	Limitation
Mishra et al. [2], 2020	Smartwatch anomaly detection	Deviations can precede COVID-19 symptoms	Individual-level; not federated
Quer et al. [3], 2021	Wearable data + self-reports	Combining both improves COVID-19 detection	Centralized collection
Grzesiak et al. [4], 2021	Noninvasive wearable sensors	Influenza/cold detectable before symptoms	Controlled study design
Cacheda et al. [5], 2026	Personalized baseline deviation	Personal baselines improve early detection	No federated population layer
Inpatient-pretrain study [32]	Transfer from inpatient data	Works when dataset shift is controlled	Direct application near chance
GenAI twin study [31]	WGAN twins + anomaly detection	Augmentation raises detection accuracy	Not federated; synthetic fidelity
Can & Ersoy [14], 2021	Federated deep learning, wearable IoT	Private training feasible	Stress focus, not epidemic
Hu et al. [16], 2023	FedIERF, wearable data	Incremental federated updates	Prototype-level validation
Aminifar et al. [17], 2024	Edge FL + SMC, 256 Hz EEG	Strong privacy	High comm. and energy cost

Author/Year	Method / Data	Key Findings	Limitation
Baucas et al. [12], 2023	FL + blockchain fog-IoT	Trust among stakeholders	Consensus overhead
Ravi Shanker Reddy & Beena [15], 2022	Consent-based federated transfer + blockchain	Consent and integrity controls	Added computation; scale unclear
Ranjbaran et al. [1], 2026	Review of 38 selected papers	Co-design of sensing, learning, communication is key	No quantitative benchmark

TABLE I. Literature Comparison Summary

F. Research Gap

The literature falls into three clusters: individual-level wearable infection detection that is mostly centralized [2]–[5]; federated healthcare work focused on stress, chronic disease, or activity rather than epidemic surveillance [14], [16]–[18]; and surveys that treat privacy, sensing, cloud, and blockchain as separate pillars [1]. Few works specify a complete pipeline in which on-device personalized detectors feed a privacy-protected population-level outbreak signal while jointly respecting energy, connectivity, regulatory, and cold-start constraints. This gap motivates the framework proposed in Section III.

III. PROPOSED SYSTEM / RESEARCH METHODOLOGY

This section describes the proposed methodology, organized around the system architecture, sensor data collection, dataset description, preprocessing, feature engineering, candidate learning models, and the proposed innovations.

A. System Architecture

The proposed pipeline follows the sequence: Consumer Wearable → Edge Preprocessing → On-Device Personalized Detector → Smartphone Gateway → Secure Aggregation → Regional Edge Server → Public-Health Cloud → Outbreak Dashboard, illustrated in Fig. 1. Wearables sense physiological and activity signals continuously. Preprocessing and inference run on the device or paired smartphone. Only protected model updates and thresholded anomaly flags are sent, over a low-power wireless link to the smartphone and then to a regional edge server, which performs intermediate aggregation before forwarding to a health-authority layer that produces outbreak indicators for decision-makers.

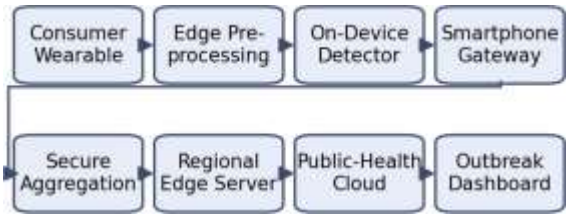


Fig. 1. Overall System Architecture

B. Sensor Data Collection

Table II lists the sensing parameters considered. Fig. 2 shows how they are collected on the wearable and aggregated at the smartphone gateway. Hardware is described as representative of common consumer devices; no physical deployment was carried out for this study, and the parameter list defines the input space that the framework is designed to accept.

Parameter	Sensor Type / Example	Unit	Purpose in Epidemic Detection
Heart rate / resting HR	Photoplethysmography (PPG)	bpm	Elevated resting HR is a common early infection marker
Heart-rate variability	PPG / ECG-derived	ms	Autonomic and inflammatory response
Skin / body temperature	Skin-contact thermistor	°C	Fever and thermoregulation shifts
Blood oxygen	Reflective pulse oximetry	%SpO2	Respiratory involvement
Respiratory rate	PPG / accelerometer-derived	breaths/min	Respiratory-infection indicator
Sleep and activity	Accelerometer, gyroscope	steps, min	Behavioral change and fatigue
Cough events (optional)	Low-rate microphone / accelerometer	events/h	Respiratory symptom proxy, processed on-device only
Coarse location cell	Phone location, generalized	cell ID	Spatial aggregation with k-anonymity

TABLE II. Sensor Parameters and Their Epidemiological Significance



Fig. 2. Sensor Data Collection Framework

C. Dataset Description

The framework assumes a per-user time series in which each record is an hourly or daily aggregate of sensor features, optionally paired with a delayed, user-reported or test-confirmed infection label. Where a public dataset is used to instantiate the framework, an openly available wearable or physiological dataset or institutional data may be used; its size and provenance should be documented at implementation rather than assumed in advance, and no dataset size is claimed in this proposal. Table III summarizes the intended dataset structure.

Attribute	Description
Input variables	Resting HR, HRV, skin temperature, SpO2, respiratory rate, sleep, activity
Target variable	Infection-anomaly flag (device level); outbreak indicator (population level)
Sampling frequency	Seconds to minutes at the sensor, aggregated to hourly or daily features
Data records	To be determined at implementation, by cohort size and duration
Labels	Delayed test-confirmed or self-reported infection; sparse and noisy
Location	Generalized cell only; raw GPS never leaves the phone
Population	Heterogeneous users; non-IID by design

TABLE III. Dataset Attributes

D. Data Preprocessing

Wearable data collected in daily life is rarely suitable for direct use in model training. The proposed on-device preprocessing sequence, shown in Fig. 3, addresses eight concerns:

- Missing-value handling: gaps from device removal, charging, or Bluetooth loss are interpolated when short and masked when long, preserving time-series continuity.
- Motion-artifact and outlier rejection: PPG-derived readings during high activity are discarded or down-weighted, and implausible values (e.g., SpO2 above 100%) are removed by range checks and the interquartile-range (IQR) rule.
- Noise filtering: moving-average or Savitzky–Golay smoothing over short windows.

- Context alignment: readings are conditioned on activity state, sleep stage, and time of day so that resting measurements are compared with resting baselines.
- Personal-baseline normalization: z-score standardization against the user's own rolling baseline, $x' = (x - \mu_{user}) / \sigma_{user}$.
- Data synchronization: sensors with different sampling rates are aligned to a common hourly index.
- Duplicate removal: repeated transmissions are identified by timestamp and device ID and removed.
- Train-test splitting: partitioned by user and by time period, never randomly inside one stream, to avoid leakage between adjacent records.

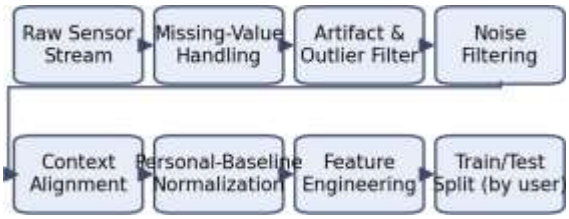


Fig. 3. On-Device Data Preprocessing Pipeline

E. Feature Engineering

Raw streams are converted into features that capture deviation from personal normal: rolling-baseline residuals for resting HR, HRV, temperature, SpO2, and respiratory rate; z-score deviation and its slope; sleep-efficiency change; nocturnal versus daytime ratios; step-count drop relative to the user's weekly pattern; multi-signal agreement scores (the number of signals deviating simultaneously); and, optionally, cough-event rate. These features encode physiologically meaningful patterns, such as a concurrent rise in resting HR and respiratory rate with reduced HRV, that single instantaneous readings do not capture.

F. Candidate Learning Approaches

Six approaches are considered, spanning statistical baselines, tree ensembles, temporal networks, and federated personalized and transfer methods. Table IV summarizes their characteristics.

1) Personal Baseline-Deviation Model. Flags a user when features exceed a multiple of their own baseline variation. It is lightweight, interpretable, and suited to on-device use [5], but has limited ability to model interactions.

2) Isolation Forest / One-Class Detector. Learns normal behavior without infection labels and scores deviations, which is useful because labels are scarce and delayed.

3) Federated Random Forest / FedIERF. An incremental extremely random forest trained federatively [16]; robust to noise and small data and supports incremental updates.

4) FedAvg with a Lightweight CNN/LSTM. FedAvg [6] aggregates local updates of a compact temporal network to learn shared infection signatures across users, at higher compute and communication cost.

5) FedProx / Personalized FL. Adds a proximal term to cope with heterogeneous clients [9], then fine-tunes locally, balancing generalization and personalization.

6) Federated Transfer Learning. Starts from a model pretrained on inpatient or synthetic data [31], [32] and adapts it federatively, addressing cold start for a novel pathogen.

Model	Advantages	Disadvantages	Suitability
Baseline-Deviation	Tiny, interpretable, label-free	Cannot model interactions	First-line on-device detector
Isolation Forest / One-Class	Label-free, fast	Threshold tuning; false alarms	Cold-start anomaly scoring
Federated RF / FedIERF	Robust to noise, incremental	Larger model size	Resource-limited phones
FedAvg + CNN/LSTM	Captures temporal patterns	Energy and communication cost	Larger cohorts; plugged-in training
FedProx / Personalized FL	Handles non-IID data	Extra tuning and local compute	Heterogeneous populations
Federated Transfer Learning	Fast start for new pathogens	Dataset-shift risk	Early outbreak phase

TABLE IV. Comparison of Learning Approaches

G. Proposed Innovations

The review [1] concludes that wearable FL is a multi-objective co-design problem and that no surveyed work co-optimizes sensing, communication, learning, and aggregation together. The framework therefore introduces the following mechanisms, each tied to a gap identified in the literature.

1) Personalized baseline-deviation detection. Each device keeps a rolling personal baseline and reports a deviation score, not raw values, following evidence that individual baselines outperform population thresholds [5] while keeping sensitive data local.

2) Hierarchical hybrid aggregation. Updates are aggregated at the smartphone, then at regional edge nodes, then at the health-authority cloud, reducing communication rounds and localizing aggregation relative to flat topologies [1], at the cost of added synchronization complexity.

3) Outbreak-risk-adaptive differential privacy. The privacy budget is kept tight in normal conditions and relaxed within bounds when regional anomaly rates rise, trading some privacy for sensitivity when timeliness matters most. This operationalizes the adaptive-privacy guideline of [1] using differential privacy [8].

4) Energy-aware client participation. Training runs only when the device is charging, idle, and on unmetered connectivity; clients are selected by battery and connectivity state, with participation-weighted aggregation to limit the bias that non-uniform participation introduces [1].

5) Federated analytics with k-anonymous geo-cells. Devices contribute bucketed anomaly flags tied to a coarse spatial cell. Secure aggregation [7] ensures the server sees only sums, and a cell is reported only when at least k devices contribute, preventing re-identification in sparsely populated areas.

6) Permissioned-ledger audit trail off the critical path. Update provenance, consent events, and model versions are logged on a permissioned ledger with deterministic consensus (e.g., Hyperledger Fabric with PBFT), as in [12], [13]. Because consensus latency bounds the aggregation interval [1], logging is asynchronous so it does not block real-time detection; permissionless chains are avoided for energy and latency reasons.

7) Federated unlearning. To support GDPR Article 17, training is partitioned into shards or checkpoints so that a user's contribution can be removed by retraining only the affected portion, following sharded machine-unlearning ideas [10]. The review [1] notes that no surveyed work addresses this.

8) Pathogen-agnostic cold start. At the start of a new outbreak no pathogen-specific labels exist. The framework combines label-free anomaly detection, transfer from pretrained inpatient-derived models with explicit dataset-shift control [32], and synthetic digital-twin augmentation for rare-event coverage [31], then adds local fine-tuning as labels arrive.

A cross-cutting **compliance-by-design mapping** relates each component to the HIPAA technical, administrative, and physical safeguards, GDPR Articles 5, 6, 9, 17, and 20, the European Health Data Space objectives, and, where applicable, SaMD requirements including a predetermined change control plan for adaptive models [1].

IV. PROPOSED MACHINE LEARNING FRAMEWORK

Building on the architecture and models described in Section III, the end-to-end workflow is specified step by step and illustrated in Fig. 4.

1. Enroll consenting users and record consent state locally.
2. Continuously collect wearable signals on the device.
3. Clean data: remove duplicates, motion artifacts, and implausible readings.
4. Handle missing values by interpolation or masking.
5. Compute personal baselines and deviation features.
6. Run the on-device detector and produce an anomaly score.
7. When charging, idle, and on an unmetered network, train locally on recent data.
8. Clip and noise the local update according to the current privacy budget.
9. Send the protected update through secure aggregation to the regional edge server.
10. Aggregate hierarchically and distribute the updated global model.
11. Fine-tune locally for personalization.
12. Contribute bucketed anomaly flags to the geo-cell federated-analytics round.
13. Release a cell-level indicator only if at least k devices contributed.
14. Fuse indicators with conventional surveillance data in the public-health dashboard.
15. Log provenance asynchronously on the permissioned ledger.
16. Evaluate with the metrics of Section VI and select the configuration with the best accuracy, timeliness, and energy trade-off.

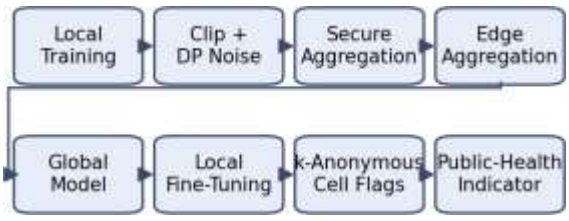


Fig. 4. Federated Epidemic-Tracking Workflow

A simplified pseudocode representation of one federated round is as follows:

```
for round t in 1..T:
  S_t = select_clients(battery, link, region)
  for client k in S_t (in parallel):
    w_k = local_train(w_global, D_k)
    w_k = clip(w_k - w_global, C) + noise(eps_t)
    send(secure_aggregate(w_k))
  w_global = hierarchical_aggregate(edges)
  eps_t = adapt_budget(regional_anomaly_rate)
  log_async(ledger, version(w_global))
```

V. EXPERIMENTAL SETUP

This section describes how the framework can be evaluated once field or public data is obtained. As no field trial was conducted, the setup is presented as the intended configuration rather than a record of hardware or software actually deployed.

- Hardware requirements: consumer wearables with PPG, temperature, and accelerometer sensors; smartphones as gateways; a regional edge server and a cloud instance for aggregation.
- Software requirements: Python 3.x; TensorFlow Federated or PyTorch-based FL toolkits; Scikit-learn for baselines.
- Supporting libraries: Pandas and NumPy for data handling; Matplotlib for visualization.
- Simulation: a client simulator injecting non-IID distributions, dropout, and variable battery, plus a synthetic epidemic process (e.g., SIR/SEIR dynamics) that generates infection onsets with known ground truth.
- Privacy tooling: differential-privacy accounting and a secure-aggregation protocol.
- Ledger: a permissioned blockchain test network for audit-trail experiments.
- Training procedure: k-fold cross-validation by user, with hyperparameters chosen by grid or randomized search.
- Validation procedure: hold-out evaluation on temporally and spatially separated users and regions to reduce leakage.

VI. PERFORMANCE EVALUATION

Because the system is multi-objective, metrics are proposed in five groups.

Detection quality. Precision, recall (sensitivity), F1-score, and AUROC, with attention to false-alarm rate because population-scale alerts amplify small error rates.

Timeliness. Detection lead time (hours between the detected anomaly and symptom onset or test confirmation) and outbreak-detection delay (time from true onset of community transmission to a population-level alert).

Communication and convergence. Bytes per round per client, rounds needed to reach target performance, and fraction of clients dropped.

Energy and resource use. Battery drain per day attributable to sensing, inference, and training, plus on-device memory and latency.

Privacy and trust. Privacy budget (ε, δ) consumed, resistance to membership-inference and gradient-inversion tests, minimum k per reported cell, and unlearning cost relative to full retraining.

The review [1] recommends moving beyond single-metric accuracy toward multi-objective benchmarking that includes energy, convergence time, and regulatory readiness; the set above follows that recommendation.

VII. RESULTS AND ANALYSIS

No wearable dataset was collected or analyzed for this study, and no individual was monitored. Consequently, this section presents a proposed, illustrative results framework rather than genuine experimental findings. After implementation and evaluation on a real dataset, the framework can be assessed using the table and figures below populated with actual computed values; the entries shown here are placeholders that describe the intended structure of the analysis and must not be interpreted as measured accuracy.

Approach	Prec.	Recall	F1	AUROC	Lead (h)	Energy
Baseline-Deviation	—	—	—	—	—	—
One-Class Detector	—	—	—	—	—	—
Federated RF / FedIERF	—	—	—	—	—	—
FedAvg + CNN/LSTM	—	—	—	—	—	—
FedProx / Personalized	—	—	—	—	—	—

Approach	Prec.	Recall	F1	AUROC	Lead (h)	Energy
Federated Transfer	—	—	—	—	—	—

TABLE V. Model Evaluation Metrics — To Be Populated From Actual Experimental Data (Not Measured In This Study)

To illustrate what the evaluation outputs would look like once genuine data is available, Figs. 5–7 show representative, clearly labelled illustrative visualizations rather than results derived from this study. Fig. 5 shows an illustrative feature-importance ranking of the type a tree-based detector would produce, with personal-baseline deviations in resting heart rate, skin temperature, and respiratory rate expected to carry weight in line with prior wearable studies [2]–[5]. Fig. 6 shows an illustrative privacy–utility curve in which detection performance rises as the privacy budget ε is relaxed. Fig. 7 shows an illustrative comparison across the six approaches, consistent in relative ordering with the expectation from [1] that personalized and hierarchical designs outperform flat, non-personalized averaging under non-IID data, though the specific values are for illustration only and are not derived from an executed experiment.

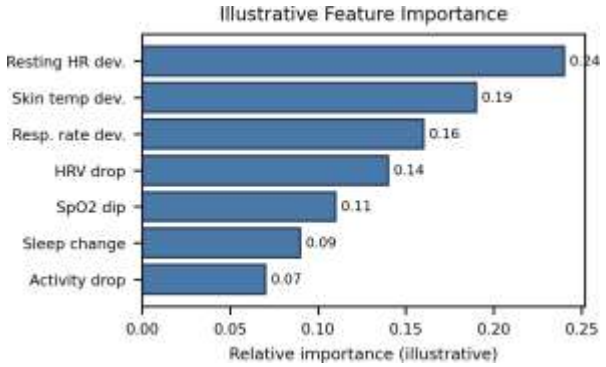


Fig. 5. Illustrative Feature Importance Analysis (not measured in this study)

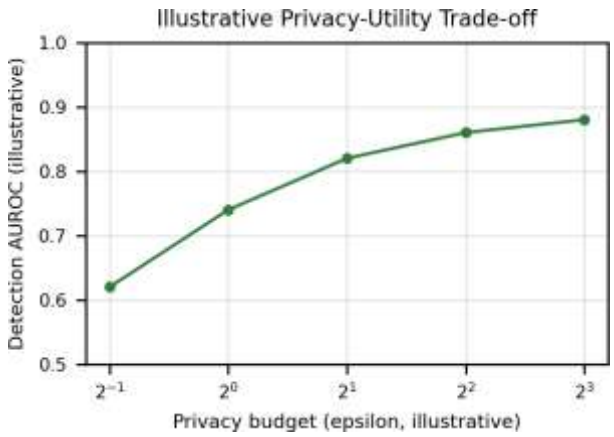


Fig. 6. Illustrative Privacy–Utility Trade-off (not measured in this study)

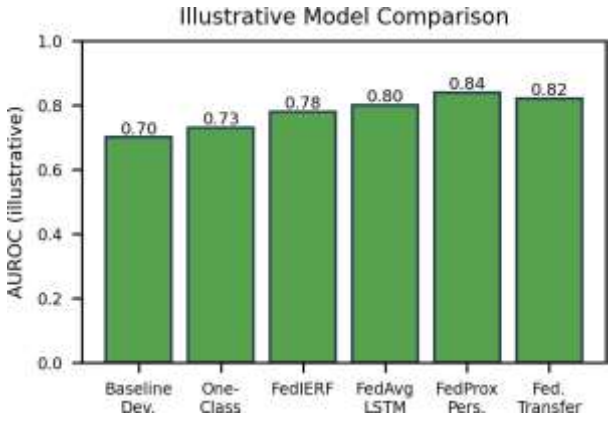


Fig. 7. Illustrative Model Performance Comparison (not measured in this study)

Beyond the illustrative values, the qualitative pattern expected from the reviewed literature is that personalized and transfer-based federated approaches should outperform fixed population thresholds because they model each person's own baseline, while stronger privacy is expected to reduce sensitivity for small cohorts. Fig. 8 shows how the federated outbreak indicator is intended to feed downstream public-health functions rather than terminating at a single number.

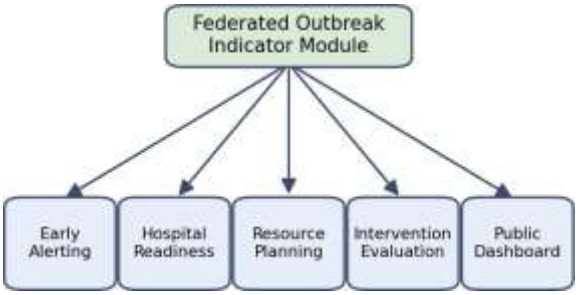


Fig. 8. Proposed Public-Health Decision-Support Architecture

VIII. DISCUSSION

The main value of the framework is moving epidemic sensing from episodic, facility-based reporting to continuous, passive, privacy-preserving observation. Because the first stage is personal, it avoids the population-threshold problem of fever-style screening, and because the second stage uses only protected aggregates, it can scale without a central store of raw physiology. The review [1] stresses that effectiveness depends less on any one optimizer than on aligning architecture with device constraints, data heterogeneity, and clinical requirements; the eight innovations are chosen for that reason.

Several trade-offs remain and are interdependent. Stronger privacy can reduce sensitivity in small cohorts. Greater decentralization can slow convergence under non-IID data. Blockchain improves auditability but its consensus latency constrains the aggregation interval, which is why logging is kept asynchronous. On-device training lowers cloud dependence but increases battery use. Personalization improves individual detection but can reduce robustness across diverse populations [1].

Several practical concerns affect real-world deployment: sensor noise and motion artifacts, device heterogeneity, label scarcity and delay, participation bias (wearable owners are not representative of the whole population), dataset shift between pathogens, and the false-alert burden of population-scale

monitoring. Finally, a model that adapts after release may qualify as software as a medical device, bringing pre-market, validation, and post-market obligations [1].

IX. REAL-WORLD APPLICATIONS

1. Early outbreak alerting: regional anomaly indicators that precede laboratory-confirmed case growth.
2. Respiratory-season monitoring: influenza and COVID-like illness trends by neighborhood cell.
3. Individual early warning: private on-device prompts to test or isolate.
4. Hospital and clinic readiness: forecasting demand from population-level indicators.
5. Workplace and school monitoring: aggregate, k-anonymous absence-risk signals without individual disclosure.
6. Rural and low-connectivity areas: edge-first operation tolerant of intermittent links.
7. Cross-border collaboration: federated analytics across jurisdictions without moving raw data, consistent with European Health Data Space goals [1].
8. Intervention evaluation: tracking population-level physiological anomaly trends after vaccination or policy changes.
9. Elderly and high-risk monitoring: personalized baselines for vulnerable groups [28].
10. Public-health decision support: dashboards that fuse wearable indicators with conventional surveillance.

X. LIMITATIONS

Table VI summarizes the principal limitations of the proposed framework.

Limitation	Description
Sensor accuracy	Consumer PPG and temperature sensors are noisy and device-dependent
Non-specificity	Physiological deviations are shared by many conditions, not only infection
Non-IID data	User-specific distributions slow and destabilize convergence
Client dropout	Battery and connectivity limits bias updates
Label scarcity	Confirmed infection labels are sparse and delayed
Dataset shift	Models trained on one pathogen or population may not transfer [32]
Participation bias	Wearable owners are not representative of the whole population
Privacy–utility cost	Differential privacy can reduce sensitivity in small cohorts
Blockchain overhead	Consensus latency and energy limit aggregation frequency
Regulatory burden	HIPAA, GDPR, EHDS, and SaMD obligations are largely unresolved for FL
Physical safeguards	Lost or stolen wearables create exposure no FL algorithm removes
No empirical validation	This study presents a design, not measured results

TABLE VI. Limitations of the Proposed System

These limitations mean the proposed system should not be regarded as universally accurate; its performance will depend on cohort size, device quality, and deployment duration, and on careful region-specific validation before being relied upon for public-health action.

XI. FUTURE WORK

- Field pilot: prospective deployment with a consented cohort and comparison against conventional surveillance.
- Cross-layer co-optimization: joint design of sensing rate, communication, privacy budget, and aggregation, which none of the surveyed works achieves [1].
- TinyML: lightweight, quantized on-device models to cut energy.
- Verifiable federated unlearning with low retraining cost.
- Fairness-aware federated learning to counter participation bias across demographics and regions.
- Explainable AI: SHAP-style attributions on-device to build user and clinician trust.
- Multimodal fusion with cough audio, mobility, and environmental data.
- Spatiotemporal epidemic models combining federated indicators with compartmental or graph-based forecasting.
- Lightweight ledger alternatives for low-power devices.
- Standardized interoperability: unified data schemas and cross-platform FL protocols.
- Regulatory pathways: SaMD classification and predetermined change control plans for adaptive federated models.
- Continual and incremental learning to track viral variants without full retraining.

XII. CONCLUSION

This paper has proposed a decentralized epidemic-tracking framework that combines federated learning and edge AI on consumer wearables. Personalized on-device detection, hierarchical privacy-protected aggregation, federated analytics with k-anonymous geo-cells, energy-aware participation, an asynchronous permissioned audit trail, federated unlearning, and a pathogen-agnostic cold-start strategy together target the constraints that the wearable FL literature identifies as unresolved: privacy and security, energy, convergence, trust, and regulation. Six learning approaches were compared conceptually and a multi-objective evaluation design was defined. Because no dataset was collected for this study, the work is presented explicitly as a proposed framework and experimental design rather than a report of measured results, and all illustrative figures and tables are labelled accordingly. Moving from concept to a field-validated public-health tool will require prospective pilots, rigorous privacy and fairness evaluation, and regulatory engagement.

ACKNOWLEDGMENT

The author gratefully acknowledges the guidance and support of the faculty guide and the Department of Computer Science and Engineering, Kings College of Engineering, Punalkulam, Tamil Nadu, India, in the preparation of this work. Appreciation is also extended to the institution for providing the academic environment and resources that made this research possible.

REFERENCES

[1] G. Ranjbaran, S. Consoli, G. Leoni, D. Reforgiato Recupero, and C. K. Roy, "Federated artificial intelligence approaches for wearables and health devices," Smart Health, vol. 40, p. 100663, 2026, doi: 10.1016/j.smhl.2026.100663.

[2] T. Mishra et al., "Pre-symptomatic detection of COVID-19 from smartwatch data," Nature Biomedical Engineering, vol. 4, pp. 1208–1220, 2020.

[3] G. Quer et al., "Wearable sensor data and self-reported symptoms for COVID-19 detection," Nature Medicine, vol. 27, pp. 73–77, 2021.

[4] A. B. Grzesiak et al., "Assessment of the feasibility of using noninvasive wearable biometric monitoring sensors to detect influenza and the common cold before symptom onset," JAMA Network Open, vol. 4, no. 9, e2128534, 2021.

[5] F. Cacheda, V. Carneiro, M. A. Álvarez, and M. F. López-Vizcaíno, "Early detection of infectious diseases using wearable sensor data through personalized baseline deviation modelling," Digital Health, vol. 12, 2026, doi: 10.1177/20552076261477919.

[6] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. Agüera y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Proc. AISTATS, 2017.

[7] K. Bonawitz et al., "Practical secure aggregation for privacy-preserving machine learning," in Proc. ACM CCS, 2017.

[8] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," Foundations and Trends in Theoretical Computer Science, vol. 9, nos. 3–4, pp. 211–407, 2014.

[9] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," in Proc. MLSys, 2020.

[10] L. Bourtoule et al., "Machine unlearning," in Proc. IEEE Symposium on Security and Privacy, 2021.

[11] E. Badidi, "Edge AI for early detection of chronic diseases and the spread of infectious diseases: Opportunities, challenges, and future directions," Future Internet, vol. 15, no. 11, p. 370, 2023.

[12] M. J. Baucas, P. Spachos, and K. N. Plataniotis, "Federated learning and blockchain-enabled fog-IoT platform for wearables in predictive healthcare," IEEE Trans. Computational Social Systems, vol. 10, no. 4, pp. 1732–1741, 2023.

[13] K. Farooq, H. J. Syed, S. O. Alqahtani, W. Nagmeldin, A. O. Ibrahim, and A. Gani, "Blockchain federated learning for in-home health monitoring," Electronics, vol. 12, no. 1, p. 136, 2022.

[14] Y. S. Can and C. Ersoy, "Privacy-preserving federated deep learning for wearable IoT-based biomedical monitoring," ACM Trans. Internet Technology, vol. 21, no. 1, pp. 1–17, 2021.

[15] T. Ravi Shanker Reddy and B. M. Beena, "AI integrated blockchain technology for secure health care—Consent-based secured federated transfer learning for predicting COVID-19 on wearable devices," in Proc. ICICC 2022, vol. 1, Springer, 2022, pp. 345–356.

[16] C.-Y. Hu et al., "FedIERF: Federated incremental extremely random forest for wearable health monitoring," J. Computer Science and Technology, vol. 38, no. 5, pp. 970–984, 2023.

[17] A. Aminifar, M. Shokri, and A. Aminifar, "Privacy-preserving edge federated learning for intelligent mobile-health systems," arXiv:2405.05611, 2024.

[18] V. Kotiyal et al., "Skipper: A federated Siamese network-based group activity segregator for IoMT systems," IEEE Trans. Computational Social Systems, vol. 10, no. 4, pp. 1770–1779, 2023.

[19] K. T. Putra et al., "A review on the application of Internet of Medical Things in wearable personal health monitoring: A cloud-edge artificial intelligence approach," IEEE Access, 2024.

[20] M. Ali, F. Naeem, M. Tariq, and G. Kaddoum, "Federated learning for privacy preservation in smart healthcare systems: A comprehensive survey," IEEE J. Biomedical and Health Informatics, vol. 27, no. 2, pp. 778–789, 2022.

[21] D. C. Nguyen et al., "Federated learning for smart healthcare: A survey," ACM Computing Surveys, vol. 55, no. 3, pp. 1–37, 2022.

[22] D. N. Sachin, B. Annappa, and S. Ambesenge, "FedRH: Federated learning based remote healthcare," in Proc. IEEE Int. Conf. Blockchain and Distributed Systems Security, 2023, pp. 1–7.

[23] D. Y. Zhang, Z. Kou, and D. Wang, "FedSens: A federated learning approach for smart health sensing with class imbalance in resource constrained edge computing," in Proc. IEEE INFOCOM, 2021, pp. 1–10.

[24] T. Wang, Y. Du, Y. Gong, K.-K. R. Choo, and Y. Guo, "Applications of federated learning in mobile health: Scoping review," J. Medical Internet Research, vol. 25, e43006, 2023.

[25] B. Pfitzner, N. Steckhan, and B. Arnrich, "Federated learning in a medical context: A systematic literature review," ACM Trans. Internet Technology, vol. 21, no. 2, pp. 1–31, 2021.

[26] N. Rieke et al., "The future of digital health with federated learning," npj Digital Medicine, vol. 3, p. 119, 2020.

[27] P. Kairouz et al., "Advances and open problems in federated learning," Foundations and Trends in Machine Learning, vol. 14, nos. 1–2, pp. 1–210, 2021.

[28] S. Ghosh and S. K. Ghosh, "FEEL: Federated learning framework for elderly healthcare using edge-IoMT," IEEE Trans. Computational Social Systems, vol. 10, no. 4, pp. 1800–1809, 2023.

[29] A. Lakhan et al., "Digital healthcare framework for patients with disabilities based on deep federated learning schemes," Computers in Biology and Medicine, vol. 169, 107845, 2024.

[30] A. Alahmadi et al., "A privacy-preserved IoMT-based mental stress detection framework with federated learning," J. Supercomputing, vol. 80, no. 8, pp. 10255–10274, 2024.

[31] "GenAI-based digital twins aided data augmentation increases accuracy in real-time cokurtosis-based anomaly detection of wearable data," PubMed Central, PMC12431508 (open access).

[32] "An infection prediction model developed from inpatient data can predict out-of-hospital COVID-19 infections from wearable data when controlled for dataset shift," Europe PMC, EPMC12480451 (open access).