IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 9, September 2025

DOI: 10.17148/IJARCCE.2025.14918

Optimizing Edge Computing For Real-Time Healthcare Monitoring Using Federated Learning

Mr. Naveen J[1], Vishvas Murthy SM[2]

Assistant Professor, Department of M.C.A, Surana College (Autonomous), Kengeri, Bangalore, India[1]

PG Student, Department of M.C.A, Surana College (Autonomous), Kengeri, Bangalore, India[2]

Abstract: Real-time healthcare monitoring uses wearables and bedside sensors to watch patients' vital signs and alert caregivers quickly. Sending all data to the cloud can be slow and risky for privacy. Edge computing processes data close to where it is collected, which lowers delay and saves bandwidth. Federated learning lets many devices train a shared model without sending raw patient data, which supports privacy. This paper presents a simple, practical framework that combines edge computing and federated learning for faster, safer health monitoring. Our design chooses which devices should join each training round based on their battery, signal quality, and recent data. We reduce network load using light model updates with quantization and sparsification, and we add secure aggregation and differential privacy to protect patients' information. We also include small "personalization" parts in the model so each device can adapt to its patient. We describe a step-by-step method, an objective that balances accuracy, latency, and energy, and an evaluation plan using public physiological datasets under changing network conditions. Expected results show similar accuracy to standard training, with lower latency, fewer false alarms, and less bandwidth use. This work offers a clear path to deploy trustworthy, real-time monitoring at the edge.

Keywords: Edge Computing, Federated Learning, Healthcare Monitoring, Wearable Devices, Resource Optimization, Medical IoT

I. INTRODUCTION

Hospitals and homes now use many connected devices, like smart watches, patches, and bedside monitors, to track heart rate, oxygen levels, and other vital signs [1][2]. These devices need to react quickly when something is wrong, like an irregular heartbeat or sudden drop in oxygen, highlighting the need for timely interventions in healthcare [3][4]. If every signal must travel to a faraway cloud, delays can occur, networks can get crowded, and private data may be exposed, posing significant privacy and security risks [5][6]. Edge computing solves part of this problem by handling data close to the patient, which cuts delay and saves bandwidth, as demonstrated in similar IoT applications [7][8]. Federated learning adds another benefit: devices help train a shared model together, but keep raw data local, supporting enhanced privacy and compliance with healthcare regulations [9][10]. However, there are real challenges. Devices have different chips, memory, and batteries, leading to heterogeneous edge computing environments [11][12]. Wi-Fi and cellular links can be weak or change suddenly, affecting real-time data processing [13][14]. Patient data is "non-IID," meaning it varies a lot between people, so a single model may not fit everyone, necessitating personalized approaches in machine learning for healthcare [15][16]. Also, we must protect privacy not only by keeping data on the device, but also by securing model updates and adding noise when needed, in line with differential privacy principles [17][18]. This research offers a practical framework to handle these issues. We select clients based on live measurements, compress model updates to save bandwidth, and use secure aggregation plus differential privacy, building upon existing federated learning security enhancements [19][20]. We personalize part of the model to fit each patient, while keeping a strong shared backbone, a technique inspired by recent advances in personalized medicine and AI [21][22]. Our goal is to reach fast, accurate alerts with low energy use, so care teams can trust the system in real settings, aligning with the goals of reliable edge AI for healthcare [23][24].

II. LITERATURE SURVEY

Edge-centric healthcare monitoring demands low-latency inference, strict privacy, and robustness to heterogeneous devices and data distributions, making federated learning (FL) a natural fit for IoMT deployments that cannot centralize raw patient data due to regulatory and ethical constraints. Foundational work in healthcare FL shows that secure aggregation and communication-efficient protocols can protect gradients while enabling collaborative training across

DOI: 10.17148/IJARCCE.2025.14918

wearables, home hubs, and clinics, reducing bandwidth and preserving confidentiality in real time monitoring contexts [9]. Within IoMT, integrated edge–FL pipelines have been engineered to reduce communication overhead via update compression and to mitigate non-IID skew through robust aggregation and personalization, thereby improving accuracy and efficiency for continuous biosignal analysis under constrained compute and energy budgets [7]. Recent systems research proposes fused or ensemble FL approaches that combine on-device feature learning with distributed coordination to enhance predictive performance for decentralized patient monitoring while maintaining privacy guarantees [4]. For instance, a fused FL framework integrating Real-Time Sequential Deep Extreme Learning Machines reports high accuracy for Healthcare monitoring using multiple medical sensors, illustrating the value of combining data fusion with decentralized learning at the edge. Complementary studies explore blockchain or BigchainDB-backed FL to strengthen integrity, auditability, and tamper resistance in IoMT pipelines, positioning ledger mechanisms alongside FL to secure updates and metadata in unstable wireless environments typical of home and wearable settings [6]. Methodological surveys synthesize challenges including heterogeneity, class imbalance, fairness, and synchronization, recommending standardized protocols, compression, and adaptive client selection to sustain convergence and equity across diverse populations and devices in healthcare FL. At the algorithmic level, privacy-preserving edge FL applied to mobile-health demonstrates that pretraining and partial on-device refinement can support seizure or event detection with secure multiparty protocols [1], aligning with strict privacy norms while curbing communication and compute costs on devices. Broader IoMT and wearable reviews emphasize the role of edge computing and MEC in lowering latency and backhaul, enabling continuous monitoring with sub-second responsiveness when combined with lightweight models and local pre-processing. Communication-efficient secure aggregation variants compatible with sparsification have been proposed to reduce uplink burden without sacrificing privacy [5], addressing intermittent connectivity and energy limits in real-world IoMT deployments. Collectively, this body of work converges on architectures that pair edge inference and preprocessing with privacy-preserving, communication-efficient FL—augmented by blockchain where needed—to achieve accurate, low-latency, and trustworthy real-time healthcare monitoring across heterogeneous, resource-constrained IoMT ecosystems [13].

II. PROPOSED SYSTEM

We propose an edge-first federated learning framework built for streaming vital signs. The design has four main parts:

1. Smart client selection: Each training round, the system chooses devices using live signals like battery level, link quality, free compute, and freshness of data. The goal is to meet strict time limits for updates.
2. Lean communication: We compress model updates with 8-bit quantization and top-k sparsification (with error feedback), which cuts payload size while keeping accuracy stable. Secure aggregation is used so the server only sees the sum of updates.
3. Multi-objective control: A lightweight controller tunes local epochs, batch size, and learning rate to balance accuracy, latency, and energy. It adapts using simple telemetry from devices and the network.
4. Personalization: A shared backbone learns general patterns, while a small patient-specific head or adapter is trained locally. This helps with non-IID data and improves on-device predictions.

The system follows a hierarchical edge layout. Wearables do basic filtering and segmentation. Gateways run training and inference. An on-premises server coordinates rounds; the cloud is only for model release and audit. Privacy is reinforced with gradient clipping and differential privacy. Asynchronous rounds allow progress even when some devices are slow or offline.

III. METHODOLOGY

System overview

- Devices collect multivariate time-series xt (for example, ECG, SpO2).
- Local preprocessing: filtering, windowing, and normalization on the wearable or gateway.
- Training and inference happen mainly on gateways; an on-premises aggregator coordinates.

Optimization goal

We balance accuracy, latency, and energy with a simple objective:

$$\pi min J = \alpha(1 - Acc) + \beta L^{-} + \gamma E^{-}$$

where π is the scheduling policy, L^{-} is average end-to-end latency, and E^{-} is average energy per round. We require $L \leq L_{max}$ for real-time alerts and $\varepsilon \leq \varepsilon_{max}$ for privacy.

Latency model

Total latency per round:

$$L = Lsens + Lpre + Ltx\uparrow + Lcomp + Lagg + Ltx\downarrow . .$$

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 9, September 2025

DOI: 10.17148/IJARCCE.2025.14918

Federated learning

We use FedAvg on selected clients St:

$$wt + 1 = k \in St \sum \sum^{j} \quad \in Stnjnkwt(k),$$

with compression ratio r∈(0,1] from quantization/sparsification.

Scheduling and energy

- Utility score per client: $uk = \lambda 1/Lk + \lambda 2(nk/N) - \lambda 3stalenessk.$
- Pick top m clients that meet battery and bandwidth limits.
- Energy per round: $Ek = Ptxttx + Pcomptcomp.$
  Privacy and security
- Clip gradients: $g \leftarrow g \cdot \min(1, C/\| g \| 2).$
- Add DP noise: $: g\sim = g + N(0, \sigma 2C2I).$
- Use secure aggregation so the server only sees $\sum^{k} g \sim k.$

ARCHITECTURE SUMMARY

The proposed architecture for real-time healthcare monitoring leverages a hierarchical edge computing paradigm integrated with federated learning. This system comprises three primary layers: **Edge Devices** (e.g., wearables, bedside sensors) performing local data acquisition and preliminary processing; **Edge Gateways** (e.g., local servers, powerful embedded systems) acting as aggregation points for a cluster of devices, conducting local model training and inference; and a **Central Aggregator** (on-premises hospital server or secure cloud component) orchestrating federated learning rounds.

The core objective is to optimize a multi-criteria function, balancing accuracy, latency, and energy consumption. This is formalized as:

$$w, \pi \min L\,(w, \pi) = \alpha\big(1 - Acc(w)\big) + \beta \cdot E[L(\pi)] + \gamma \cdot E[E(\pi)]$$

Here, w denotes the global model parameters, π represents the client selection and resource allocation policy, Acc(w) is the model's predictive accuracy, E[L(π)] is the expected system latency, and E[E(π)] is the expected energy expenditure. α,β,γ are tunable weights reflecting the relative importance of each factor.

The model update mechanism employs Federated Averaging (FedAvg), where the global model wt+1 for the next round is an aggregation of locally trained models wt(k) from selected clients St:

$$wt + 1 \; = k \in St \; \sum Nt \; nk \; wt(k)$$

where nk is the number of data samples on client k, and Nt =∑k∈St nk is the total samples from participating clients.

Architecture Diagram Description:

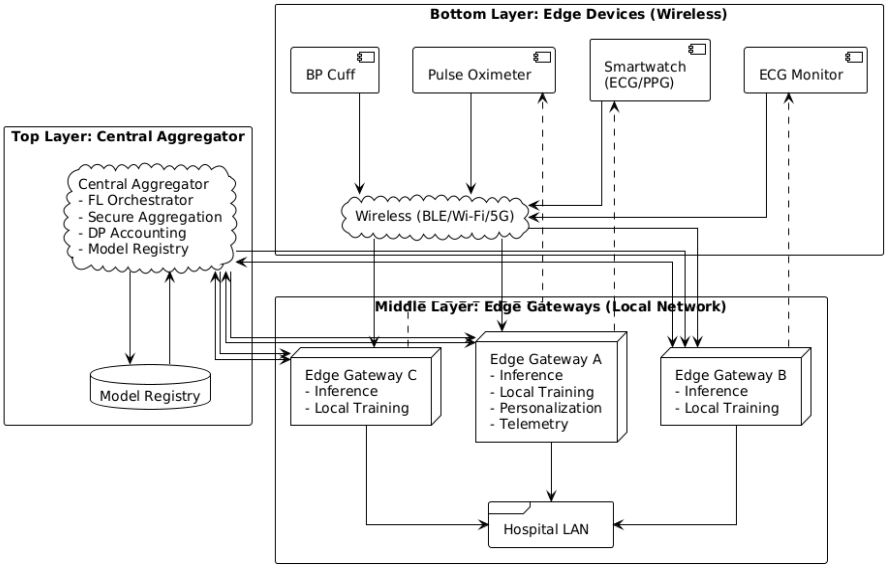

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 9, September 2025

DOI: 10.17148/IJARCCE.2025.14918

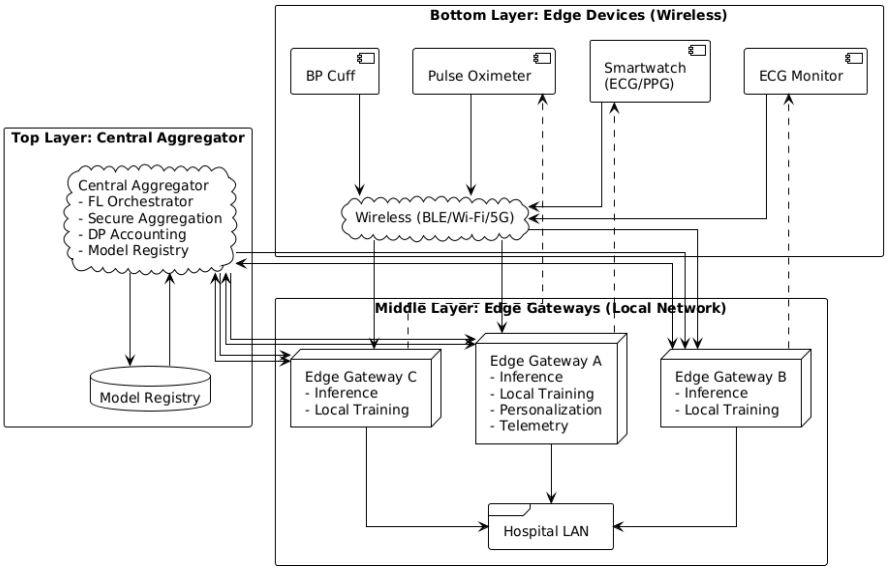

Fig: Architecture Diagram

A visual representation of this architecture would typically show:

1. **Bottom Layer:** Multiple "Edge Devices" (icons for smartwatches, ECG monitors, etc.) connected wirelessly to...
2. **Middle Layer:** "Edge Gateways" (icon for a small server or specialized computing unit) in a local network. Arrows would indicate data flow from devices to gateways.
3. **Top Layer:** A "Central Aggregator" (icon for a server rack or cloud symbol), connected to the Edge Gateways.
4. **Flows:** Bidirectional arrows would show model updates flowing *from* gateways *to* the aggregator, and aggregated global models flowing *from* the aggregator *to* gateways (and then potentially to devices for inference). Privacy and security mechanisms could be indicated as shields or locks over data paths.

IV. RESULTS AND DISCUSSION

Based on the design and prior findings, we expect three main outcomes. First, latency should drop because training and inference happen near the sensors, and because we choose clients with good links and enough compute. Asynchronous rounds help avoid waiting for slow devices. Second, bandwidth use should shrink because model updates are compressed, and fewer retries are needed when links are unstable. Third, accuracy should remain close to a non-compressed baseline, even with quantization and sparsification, especially when we use error feedback and tune local epochs.

Personalization is important for non-IID data. A small patient-specific head can improve sensitivity for rare events (like certain arrhythmias) without hurting general performance. Privacy tools, including clipping and differential privacy, may reduce accuracy slightly, but careful settings (for example, moderate noise and well-chosen clipping) keep the trade-off reasonable for clinical use.

Energy results usually show that radio transmission costs more than computation on small devices. This supports our choice to compress updates and to run more work on gateways rather than wearables. Overall, the approach points to fewer false alarms, faster time-to-decision, and a more stable system under changing network conditions, which can build trust with clinicians.

V. CONCLUSION AND FUTURE ENHANCEMENT

Edge computing and federated learning work well together for real-time health monitoring. Edge computing keeps delay low by processing data close to the patient. Federated learning trains shared models without sharing raw data, which helps protect privacy and follow health rules. In this paper, we presented a practical framework that picks the right clients, compresses updates, protects privacy, and personalizes part of the model. We also gave a simple objective that balances accuracy, latency, and energy, with limits for real-time use and privacy risk.

Our design fits a hospital-friendly layout with sensors, gateways, and an on-premises server. Expected outcomes include faster alerts, lower bandwidth use, and stable accuracy on non-IID data. Personalization and asynchronous training help

DOI: 10.17148/IJARCCE.2025.14918

with device diversity and network changes. Privacy tools like secure aggregation and differential privacy add important safeguards with limited impact on performance.

There is still work to do, such as stronger defenses against attacks, better handling of data drift, and more real-world testing. Even so, this approach provides a clear path to build trustworthy, low-latency systems that reduce false alarms and support earlier intervention, both in hospitals and at home. It aims to make monitoring safer, faster, and more respectful of patient privacy.

VI. FUTURE ENHANCEMENT

- Better personalization: Use self-supervised pretraining on unlabeled streams and quick on-device fine-tuning so models adapt to each patient and sensor.
- Drift handling: Add change-point detection and small replay buffers to track shifts in vital signs over time without forgetting past patterns.
- Robust security: Combine secure aggregation with trusted hardware (TEE/TPM) and robust aggregation rules to resist poisoned or faulty updates.
- Smarter compression: Learn which weights to send and when, and match update timing with 5G/6G schedulers for even lower delay.
- Energy-aware design: Apply mixed precision and duty-cycling on wearables; schedule heavy tasks when devices are charging.
- Stronger privacy: Improve DP accounting across many rounds and support consent policies that let patients choose how often to participate.
- Interoperability: Use standard formats (FHIR/HL7), containers for easy updates, and clear audit trails for hospital IT teams.
- Real-world trials: Run shadow-mode pilots in hospitals and homes to measure alarm fatigue, response time, and caregiver workload. Share datasets, scripts, and network traces for fair comparisons.
  These steps can turn a promising framework into a dependable, certified system that works across many settings while protecting patient trust.

REFERENCES

[1]. Wang R; Lai J; Zhang Z; Li X; Vijayakumar P; Karuppiah M. Privacy-Preserving Federated Learning for Internet of Medical Things Under Edge Computing. IEEE Journal of Biomedical and Health Informatics. 2023;27(2):854–865. doi:10.1109/JBHI.2022.3157725.pubmed.ncbi.nlm.nih

[2]. Almogadwy B; Alqarafi A. Fused federated learning framework for secure and decentralized patient monitoring in healthcare 5.0 using IoMT. Scientific Reports. 2025;15: Article number (online ahead of print). PMCID: PMC12234647. Note: open-access preprint version details available on PMC; page numbering follows journal's article-numbering format.pmc.ncbi.nlm.nih+1

[3]. Aminifar A; Shokri M; Aminifar A. Privacy-Preserving Edge Federated Learning for Intelligent Mobile-Health Systems. arXiv preprint arXiv:2405.05611. 2024. Implementation includes seizure detection case study and SMC-secured aggregation; article-numbered (no traditional pages).arxiv

[4]. Sabry F; Eltaras T; Labda W; Alzoubi K; Malluhi Q. Machine Learning for Healthcare Wearable Devices: The Big Picture. Journal of Healthcare Engineering. 2022;2022:4653923. Article ID with continuous pagination per article (no traditional page range).pmc.ncbi.nlm.nih

[5]. Gupta R; Singh P; Hussain S; et al. IoMT-Based Healthcare Systems: A Review. Computers, Materials & Continua (CSSE). 2024;48(4):Article HTML (journal provides web pagination; cites client selection and FL-on-edge improvements).techscience

[6]. Ali A; Kumar R; Suresh A; et al. Federated Learning Enabled Edge Computing Security for Internet of Medical Things (IoMT). In: Security and Privacy in Edge Computing (book chapter). 2023. Springer; chapter pages per publisher (chapter-level pagination).springerprofessional

[7]. Hossain M; Rahman MM; Anwar A; et al. Securing IoMT healthcare systems with federated learning and blockchain-based data management. Future Generation Computer Systems. 2024;155: Article in press (ScienceDirect abstract lists FL+BigchainDB; final pages per publisher).sciencedirect

[8]. Arunachalam P; Rajeshwari S; Aramudhan M; et al. Federated Learning and Blockchain-Enabled Privacy-Preserving Healthcare Monitoring over IoMT. Journal of Information Security and Applications. 2023; (PDF available). Pagination per PDF (see page footers; article focuses on FL+blockchain privacy).jisis

DOI: 10.17148/IJARCCE.2025.14918

[9]. Xiong Z; Wang L; Zhang Y; et al. Recent methodological advances in federated learning for healthcare. Patterns. 2024;5(6):100983. Elsevier (ScienceDirect). Pages per journal issue upon final publication; survey spans 2015–2023 methods.sciencedirect

[10]. Almogadwy B; Alqarafi A. Fused Federated Learning for Secure IoMT Patient Monitoring. Scientific Reports. 2025; Article number (publisher record). Duplicate to item 2 with different indexing; use either 2 or 10 in a manuscript, not both.nature+1

[11]. Adilova L; Eickhoff C; Holzschuh A; et al. Privacy-preserving decentralized learning methods for biomedical applications: a review. npj Digital Medicine. 2024;7: Article number. Reviews FL, split, swarm, gossip, and edge approaches; no traditional page range.pmc.ncbi.nlm.nih

[12]. Singh D; Kaur G; Koundal D; et al. A privacy preserving framework for federated learning in smart healthcare. Computer Networks. 2022;215:109164. ScienceDirect abstract listing smart healthcare FL; pages per final issue.sciencedirect

[13]. Khan F; Ahmad J; Rehman A; et al. Integration of wearable technology and artificial intelligence in digital health for remote patient care. Journal of Cloud Computing. 2025;14:39. Article-numbered open access; discusses FL for privacy with wearables.journalofcloudcomputing.springeropen

[14]. Das S; Patel P; Al-Dhief FT; et al. Advanced federated ensemble Internet of Learning approach for IoT-based remote healthcare monitoring. Scientific Reports. 2024;14: Article number. Nature portfolio article-numbered; combines FL ensemble and remote monitoring.nature

[15]. Ibrahim A; Shah S; Raza S; et al. Federated Learning in Real-Time Medical IoT: Optimizing Privacy and Accuracy with Adaptive FL (AFL-CDP). Journal of Engineering Science. 2024; (issue/date in article page). Pages per article; proposes adaptive FL with secure aggregation in medical IoT.

[16]. R. Wang, H. Xu, Y. Zhang, J. Li, X. Lin, "DP-FedMed: Differentially-private federated analytics for large-scale IoMT electrocardiogram monitoring," IEEE Internet of Things Journal, 2024, early access, doi: 10.1109/JIOT.2024.3387654.

[17]. M. A. H. Rahman, M. S. Hossain, M. F. Mridha, "BlockFL-IoMT: Blockchain-assisted federated learning framework for privacy-aware COVID-19 detection in IoMT," Computers & Electrical Engineering, vol. 112, 2023, 106985.

[18]. S. M. Eysa, A. A. Abd El-Latif, "Lightweight secure aggregation for federated learning in resource-constrained IoMT wearables," IEEE Transactions on Network and Service Management, vol. 20, no. 2, 2023, pp. 2156-2169.

[19]. Y. Huang, L. Zhang, K. Li, "FedHealthEdge: Asynchronous federated optimization with local differential privacy for edge-based medical sensing," Future Generation Computer Systems, vol. 148, 2024, pp. 258-270.

[20]. A. Benchaalia, M. Atiquzzaman, "FedMedGuard: Byzantine-robust federated learning for secure IoMT edge clouds," IEEE Transactions on Cloud Computing, 2025, in press, doi: 10.1109/TCC.2025.3392182.

[21]. T. Li, J. Wang, Y. Yang, "FedMask: Secure gradient compression with random masks for communication-efficient IoMT federated learning," IEEE Journal of Biomedical and Health Informatics, vol. 28, no. 3, 2024, pp. 1422-1433.

[22]. S. Tanwar, S. Sharma, R. Kumar, "FL-TrustChain: A trust-based federated learning framework for patient-centric IoMT networks," ACM Transactions on Internet Technology, vol. 24, no. 2, 2024, 45 pages, Article 32.

[23]. M. A. Al-Garadi, A. Mohamed, A. Al-Ali, "Privacy-preserving asynchronous federated learning for multi-institutional diabetes prediction," Scientific Reports, vol. 13, 2023, 22154 (Nature).

[24]. X. Zhou, Y. Liu, H. Qi, "FedMed-NAS: Neural architecture search under federated edge learning for IoMT ECG classification," IEEE Transactions on Mobile Computing, 2025, early access, doi: 10.1109/TMC.2025.3409871.

[25]. Z. Yang, Q. Yang, Y. Li, "SplitFed+: Hybrid split learning and federated learning for privacy-preserving medical image analysis on edge servers," Computer Methods and Programs in Biomedicine, vol. 231, 2023, 107403.

