





International Journal of Innovative Research in Computer and Communication Engineering

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





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

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

A Review on Privacy-Preserving Multimodal Patient Monitoring using Edge AI

Shreya Roy¹, Mayank Balpande², Prachi Mishra³, Roma Dive⁴

B. Tech Scholar, Department of Computer Engineering, St. Vincent Pallotti College of Engineering & Technology, Nagpur, Maharashtra, India¹

B. Tech Scholar, Department of Computer Engineering, St. Vincent Pallotti College of Engineering & Technology, Nagpur, Maharashtra, India²

B. Tech Scholar, Department of Computer Engineering, St. Vincent Pallotti College of Engineering & Technology, Nagpur, Maharashtra, India³

B. Tech Scholar, Department of Computer Engineering, St. Vincent Pallotti College of Engineering & Technology, Nagpur, Maharashtra, India⁴

ABSTRACT: Remote healthcare monitoring is increasingly expected to do two things at once: catch behavioural anomalies early and keep patient data private. Most existing systems rely on cloud-based processing, where raw physiological and behavioural signals travel to remote servers for analysis, adding latency, communication overhead, and privacy exposure along the way. This review looks at recent work across five connected areas: Edge Artificial Intelligence (Edge AI), multimodal sensing, behavioural anomaly detection, Federated Learning (FL), Differential Privacy (DP), and voice-assisted healthcare interaction. Across the studies surveyed, edge-based processing consistently lowers latency and reduces how much raw data leaves the device, and wearable-sensor deep learning models have reported strong performance in detecting falls, abnormal gait, and related behavioural anomalies. Federated learning lets multiple devices train a shared model collaboratively without pooling raw patient data in one place, and differential privacy adds a further layer of protection by perturbing the model updates that are shared. What stands out, though, is how separately these strands have been developed — very few studies bring anomaly detection, privacy-preserving distributed learning, and voice-mediated intervention together into one closed-loop system. This gap motivates the framework proposed here: an integrated, privacy-preserving multimodal patient-monitoring system built around lightweight CNN-LSTM models, federated learning, differential privacy, and voice-assisted escalation.

KEYWORDS: Edge AI, Privacy-Preserving Healthcare, Multimodal Sensor Fusion, Behavioural Anomaly Detection, Federated Learning, Differential Privacy, Voice-Assisted Intervention, Remote Patient Monitoring.

I. INTRODUCTION

Continuous observation is especially valuable for elderly and chronically ill patients, which is part of why remote patient monitoring has grown into such a significant area of healthcare research. Wearable devices and IoT-based sensors now collect a wide range of patient information — heart rate, ECG, blood-oxygen saturation, body temperature, and motion data from accelerometers and gyroscopes among them. Together, these signals allow a system (or the clinicians reviewing it) to track a patient's health and flag abnormal behavioural events such as falls, prolonged inactivity, abnormal gait, and signs of distress.

Most traditional monitoring systems handle this analysis in the cloud, which means patient information is continuously sent off to remote servers. That continuous transmission comes at a cost: heavier network dependency, greater communication overhead, and legitimate privacy concerns around sensitive health data. Edge computing offers a way around this by processing data closer to the patient rather than shipping it out, cutting down on how much raw information ever needs to leave the device. The payoff is faster response times, lower bandwidth use, and sensitive information that can largely stay local.



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

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

AI has found a natural fit in this space as well. Deep learning models are now routinely used to analyse physiological and motion sensor data and pick out abnormal events, and CNN-LSTM architectures in particular are well suited to this because they extract useful features from multimodal sensor streams while also capturing the temporal patterns behind a patient's behaviour. For these models to actually run on wearable or edge hardware, though, they generally need to be made lighter — techniques like quantization and pruning help fit them within the memory and compute limits of resource-constrained devices.

Privacy is the other half of this problem. Federated Learning addresses it by letting multiple edge devices contribute to a shared model without ever transferring raw patient data off the device, and Differential Privacy takes this further by injecting controlled noise into the model updates before they are shared, making it harder to reconstruct individual patient information from what is exchanged. Together, these two techniques offer a way to run collaborative healthcare analytics without centralizing patient data in the first place.

Detecting an anomaly accurately is only useful if it is followed by timely intervention, and this is where voice-assisted systems come in — particularly for elderly patients who may struggle with conventional mobile apps. Voice interfaces offer a more accessible way to deliver alerts, walk patients through emergency guidance, and keep communication open between patients and caregivers. They can also be used to check in with a patient once an anomaly is flagged, confirming whether the situation is genuinely urgent before an alert gets escalated to caregivers or clinicians.

A fair amount of work already exists in each of these areas individually — Edge AI, behavioural anomaly detection, privacy-preserving machine learning, Federated Learning, and voice-assisted healthcare — but they have mostly developed along separate tracks. Only a handful of frameworks attempt to bring real-time multimodal sensing, on-device anomaly detection, Federated Learning combined with Differential Privacy, and voice-assisted intervention together into a single system.

This review works toward closing that gap by proposing a Privacy-Preserving Multimodal Patient Monitoring framework that brings Edge AI, Federated Learning, Differential Privacy, and voice-assisted healthcare together under one architecture. The goal is a system that detects anomalies accurately and in real time, keeps sensitive patient data protected, relies less on cloud infrastructure, and still offers timely assistance through intelligent voice-based interaction.

II. LITERATURE REVIEW

The related work below is grouped into four areas: Edge AI for healthcare monitoring, behavioural anomaly detection, privacy-preserving healthcare, and voice-assisted healthcare interaction.

A. Edge AI and Remote Patient Monitoring

By processing healthcare data directly on local devices, Edge AI cuts down on cloud dependence and helps keep sensitive patient information protected [1], [2], [10], [16]. Khalid et al. explored lightweight, interpretable AI models for applications such as cardiac arrhythmia detection [1], while Salimzhanov et al. looked at IoMT-based remote patient monitoring combined with Federated Learning, which keeps patient data local while still allowing collaborative model training [2]. Broader reviews of remote patient monitoring point to AI's growing role in distributed healthcare monitoring, though they also flag ongoing challenges around clinical validation and generalisability [10].

Deploying telemonitoring systems at scale also depends on interoperability across healthcare standards [3], and digital health technologies more broadly have been studied for how they improve healthcare quality and safety through data-driven monitoring [4]. This is not a new concern — earlier work on wearable biosensor systems had already established how valuable continuous, unobtrusive monitoring could be [18], and Clifton et al. showed the benefit of combining clinical observations with wearable sensor data for predictive monitoring [17]. Satyanarayanan's account of the emergence of edge computing supplies much of the technological groundwork for processing this kind of data closer to the patient [16].



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

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

B. Healthcare Anomaly Detection and Multimodal Monitoring

AI-based anomaly detection has been studied in both IoT and Internet of Health Things settings for spotting abnormal healthcare events [6], [9], and — as noted above — lightweight AI approaches have similarly been applied to problems like cardiac arrhythmia detection [1].

Hernandez et al. combined different types of healthcare information into a multimodal tensor-based method to support integrative, predictive analysis [12]. AI's reach in this space extends beyond anomaly detection, too, with applications reported in drug-drug interaction detection [5] and pharmacogenomics [8] — evidence that AI is being used fairly broadly across healthcare decision support, not just for monitoring.

Wearable devices remain central to all of this, offering continuous health data collection outside traditional clinical settings [17], [18]. What is missing across most of this research, though, is integration — studies tend to treat monitoring or anomaly detection as standalone tasks rather than combining multimodal sensing with privacy-preserving learning and direct patient intervention.

C. Privacy-Preserving Healthcare and Federated Learning

Federated Learning allows multiple healthcare devices or institutions to jointly train a model without ever sharing raw patient data directly [2], [14], which is exactly why it fits healthcare so well — sensitive information never has to leave its local device or institution.

Chen et al. add a further safeguard through Differential Privacy, which injects controlled noise into shared information to help protect individual patient records [11]. Combined with this, Federated Learning has become one of the more established approaches for privacy-preserving digital healthcare, as Rieke et al. discuss [14].

Even so, most of this literature treats privacy-preserving learning and healthcare monitoring as separate problems, and combining them with real-time anomaly detection and patient intervention remains largely an open task.

D. Voice-Assisted and Explainable Healthcare

Alshudukhi's work on voice-based healthcare systems shows how voice can serve as an alternative interaction method, supporting communication between patients and their care systems [13]. This is especially useful for patients who struggle with conventional touchscreen or app-based interfaces.

Explainability matters just as much in this context. Tjoa and Guan survey explainable AI for medical applications, with the goal of making AI decisions easier to understand [15], and Zhang et al. apply similar ideas to wearable-based mental-health monitoring [7]. Both point to the same underlying need — AI-supported healthcare decisions have to be interpretable to patients and clinicians, not just accurate.

III. PROBLEM STATEMENT

Cloud-based processing remains the default for most remote patient monitoring systems, and it brings latency, network dependency, and privacy risk along with it as recurring drawbacks. As the literature above shows, Edge AI, anomaly detection, Federated Learning, Differential Privacy, and voice-assisted healthcare have each been studied fairly extensively, but almost always in isolation from one another.

What is needed instead is a unified system — one that can perform real-time multimodal anomaly detection while keeping sensitive patient data local and still intervening in time when something goes wrong.

The work proposed here addresses this by integrating Edge AI, Federated Learning, Differential Privacy, and voice-assisted intervention into a single privacy-preserving patient-monitoring framework.

IV. PROPOSED WORK

Building on this gap, the proposed system integrates Edge AI, Federated Learning, Differential Privacy, and voice-assisted healthcare into a single privacy-preserving multimodal patient-monitoring pipeline. Raw patient data is preprocessed and kept entirely on the edge device rather than being transmitted elsewhere, and the underlying model is designed to identify events such as falls, prolonged inactivity, and other forms of abnormal behaviour. Since this has to



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

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

run on wearable hardware, model optimization through quantization and pruning is factored in from the start to keep memory and computational demands within reach of resource-constrained devices. Once an anomaly is flagged, a voice-assistant module checks in with the patient directly — if the response indicates everything is fine, monitoring simply continues; if there is no response, or if the event looks critical, an encrypted alert carrying only the essential information is sent on to authorized caregivers or clinicians.

V. CONCLUSION

Taken together, the literature reviewed here points to real potential in edge-based inference for cutting down latency and limiting how much raw data leaves the patient's device, and it confirms that voice-assisted systems offer a practical, accessible channel for verifying a patient's condition and intervening when needed. The recurring issue, though, is fragmentation — these capabilities keep showing up as separate strands of research rather than as parts of one system. The gap this review identifies, then, is the lack of integration between multimodal sensing, on-device anomaly detection, privacy-preserving distributed learning, and voice-mediated intervention within a single closed-loop framework. Addressing that gap is the basis for the Differential Privacy and voice-assisted escalation approach proposed in this paper.

REFERENCES

- [1] S. Khalid et al., "Lightweight and interpretable AI for cardiac arrhythmia detection," 2026.
- [2] R. Salimzhanov et al., "AI-based remote patient monitoring using IoMT and federated learning," arXiv, 2025.
- [3] R. Martins et al., "Interoperability in telemonitoring systems using healthcare standards," Frontiers in Digital Health, 2025.
- [4] H. Hu et al., "Digital health technologies for improving healthcare quality and safety," Frontiers in Public Health, 2025.
- [5] L. Huang et al., "AI-based drug-drug interaction detection for safer medication," Scientific Reports, 2025.
- [6] F. Alohalı et al., "AI-driven anomaly detection in Internet of Health Things systems," IEEE Access, 2025.
- [7] Y. Zhang et al., "Explainable AI for mental health monitoring using wearable devices," 2025.
- [8] H. Taherdoost and A. Ghofrani, "Artificial intelligence in pharmacogenomics: Applications and challenges," 2024.
- [9] M. Khan and A. Alkhathami, "Machine learning-based anomaly detection in IoT healthcare systems," Scientific Reports, 2024.
- [10] M. Shaik et al., "Remote patient monitoring using artificial intelligence: Current state and challenges," WIREs Data Mining and Knowledge Discovery, 2023.
- [11] X. Chen et al., "Privacy-preserving healthcare using differential privacy," IEEE Access, 2022.
- [12] J. Hernandez et al., "Multimodal tensor-based method for integrative and predictive healthcare analytics," 2021.
- [13] M. Alshudukhi, "Smart and interactive healthcare system based on voice recognition," 2021.
- [14] P. Rieke et al., "The future of digital health with federated learning," NPJ Digital Medicine, 2020.
- [15] E. Tjoa and C. Guan, "A survey on explainable artificial intelligence (XAI): Toward medical XAI," IEEE Transactions on Neural Networks and Learning Systems, 2020.
- [16] M. Satyanarayanan, "The emergence of edge computing," Computer, 2017.
- [17] D. A. Clifton et al., "Predictive monitoring of mobile patients by combining clinical observations with wearable sensor data," IEEE Journal of Biomedical and Health Informatics, 2013.
- [18] A. Pantelopoulos and N. G. Bourbakis, "A survey on wearable biosensor systems for health monitoring," IEEE Transactions on Systems, Man, and Cybernetics, 2012.



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

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

Table 1: Survey Table

Sr. No.	Author(s) & Year	Concept Used	Performance Evaluation Parameter	Claims by Concerned Authors	Our Findings
1.	Prabha et al. (2021)	Edge-AI, blockchain and homomorphic encryption for healthcare monitoring	Accuracy, latency overhead	Reported an accuracy of 91.9%, with encryption adding roughly 8.7% latency overhead.	Shows that edge-based anomaly detection and privacy mechanisms can work together without derailing real-time processing.
2.	Fall detection study (2021)	Class-imbalanced deep learning for wearable-sensor fall detection	Detection accuracy	Applied deep learning specifically to handle the imbalance between routine activity data and rare fall events.	Reinforces that treating falls as an anomaly-detection problem, rather than relying on fixed thresholds, fits well with continuous monitoring.
3.	Shaikh et al. (2025)	Edge-AI Empowered Health IoT	Real-time healthcare monitoring, security	Focuses on combining Edge AI with Health IoT to enable secure, real-time eHealth monitoring.	Backs the case for using Edge AI to cut reliance on cloud-based healthcare processing.
4.	Vankayalapati et al. (2025)	Federated Edge Computing	Privacy preservation, healthcare analytics	Applies federated edge computing to healthcare analytics in a way that avoids sharing sensitive data directly.	Reinforces the choice of Federated Learning for the proposed privacy-preserving monitoring framework.
5.	Chen et al. (2022)	Differential Privacy	Privacy protection, privacy-utility trade-off	Shows that Differential Privacy can prevent individual patient information from being reconstructed from shared model parameters.	Supports pairing Differential Privacy with Federated Learning to strengthen patient-data protection.
6.	Alshudukhi (2021)	Voice Recognition, Smart Healthcare	Voice interaction and accessibility	Finds that voice-based healthcare interaction can make monitoring more accessible for elderly and chronically ill patients.	Supports including a voice-assisted intervention step for patient interaction after anomaly detection.
7.	Salimzhanov et al. (2025)	AI-based Remote Patient Monitoring, IoMT, Federated Learning	Privacy-preserving collaborative model training	Shows IoMT architectures using Federated Learning can keep multi-sensor patient data local while still supporting collaborative model training.	Reinforces Federated Learning as a way to keep patient data local while still enabling collaborative learning.
8.	Alohali et al. (2025)	AI-driven Anomaly Detection, Internet of Health Things	Anomaly detection performance	AI-driven anomaly detection methods for IoT and Internet-of-Health-Things settings are shown to outperform static threshold rules in these environments.	Backs the choice of AI-based anomaly detection over static, threshold-based detection.



INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN COMPUTER & COMMUNICATION ENGINEERING

 9940 572 462  6381 907 438  ijircce@gmail.com



www.ijircce.com

Scan to save the contact details