Federated Multi-Modal Learning Across Distributed Devices

Aswathnarayan Muthukrishnan Kirubakaran[1], Nitin Sakesna[2], Suhas Malempati[3], Sumit Saha[4], Shiva Carimireddy[5], Abhirup Mazumder[6], Ram Sekhar Bodala[7]
[1,5,6] *IEEE Senior, USA*
[2]*Albertsons, USA*
[3]*Cato, USA*
[4]*East West Bank,USA*
[5]*Amtrak, USA*

Abstract—**multi-modal sensing systems generate rich physio- logical and motion data that can support real-time classification, anomaly detection, and personalized analytics. Traditional cloud- centric machine learning pipelines require transmitting raw sensor streams to remote servers, creating challenges related to privacy, bandwidth usage, and latency. This paper presents a federated multi-modal learning framework that enables dis- tributed devices to collaboratively train a shared model without exposing raw data. The framework integrates compact temporal convolution and sequence-modeling components for on-device training, combined with differential privacy and Top-K gradient sparsification to reduce information leakage and communication overhead. A three-tier architecture coordinates local processing, intermediate aggregation, and global optimization while main- taining consistent model quality under heterogeneous sensor conditions. Experiments using multi-modal datasets demonstrate that the proposed approach achieves 93.1% accuracy, reduces communication cost by 68% compared to classic federated learn- ing, and sustains 18 to 22 ms inference latency on constrained hardware. These results show that federated multi-modal learn- ing can provide scalable, privacy-conscious intelligence across large networks of distributed devices.**

Index Terms—**Federated learning, edge AI, wearable devices, cloud computing, privacy preservation, multi-modal sensors.**

I. INTRODUCTION

Wearable sensing platforms have evolved from simple step counters into powerful multi-modal devices capable of con- tinuously capturing inertial, optical, electrical, and thermal signals [1]. Modern wearables integrate accelerometers, gy- roscopes, photoplethysmography (PPG), electrocardiography (ECG), and skin temperature sensors, enabling long-term monitoring of activity patterns, cardiovascular and respiratory status, and broader mobility behavior. These signals are highly informative for detecting adverse events such as falls, tremor episodes, hypoxic events, gait instability, and sudden collapse [2].

Traditional systems typically stream raw or lightly pro- cessed data to the cloud for storage and centralized model training [3]. While cloud infrastructure offers abundant com- pute and memory, this cloud-centric paradigm suffers [4] from several limitations: (i) continuous streaming exposes private physiological and behavioral data, (ii) bandwidth and energy costs increase with sampling frequency and number of modali- ties, (iii) inference latency becomes non-deterministic, and (iv) devices become unusable when connectivity is intermittent. Simple threshold based on-device logic alleviates some latency constraints but lacks robustness to noise and fails to capture complex temporal dynamics.

Edge-native AI, where inference is executed directly on the wearable or a nearby gateway, offers a promising alternative [5]. However, training accurate models often still depends on centralized data repositories. Federated learning (FL) bridges this gap by enabling many devices to collaboratively train a shared model without exposing raw data. In FL, each client performs local training and sends only model updates to a central server for aggregation. Yet directly applying FL to wearables is challenging: devices are resource-constrained, data are highly personalized, participation is intermittent, and uplink capacity is limited [6].

This work proposes FedWear, a privacy-aware federated edge AI framework for wearable devices that operates across cloud, edge, and device tiers. The main contributions are:

- A multi-tier architecture that keeps raw signals on-device, uses the cloud only for secure aggregation and global optimization, and supports low-latency edge inference.
- A lightweight multi-modal model for local training on wearable-class processors, combining temporal convolu- tions and compact sequence modeling.
- A privacy-aware FL protocol with differential privacy and Top-*K* gradient sparsification to reduce data leakage and communication overhead.
- An empirical comparison of cloud centralized, edge only, classic FL, and FedWear configurations, highlighting ac- curacy, communication, and latency trade-offs.

II. BACKGROUND & RELATED WORK

Research combining wearable sensing, edge artificial in- telligence, and privacy preserving distributed learning has advanced significantly in recent years [7]. This section reviews four major areas of prior work: (i) wearable and multi-modal sensing for activity and health analytics, (ii) edge intelligence and on-device inference for low latency applications, (iii) federated learning methods applied to mobile and IoT ecosys- tems, and (iv) privacy enhancing mechanisms for distributed optimization.

A. Wearable and Multi-Modal Sensing

Wearable devices have rapidly matured into reliable, multi- sensor platforms capable of capturing inertial, photoplethys- mography (PPG), electrocardiography (ECG), respiration, and thermal patterns [8]. Early systems focused on threshold-based fall detection, activity monitoring, and step counting. Later work explored deep-learning-based fusion of accelerometer and gyroscope signals for complex motion classification. These approaches improved accuracy but required cloud of- floading or smartphone class compute.

Recent studies demonstrate the value of multi-modal phys- iological signals for health monitoring, including stress detec- tion, arrhythmia reconstruction, and gait analysis. However, most existing systems rely on centralized model training, exposing raw physiological data to remote servers [9]. This limits scalability for privacy critical applications such as con- tinuous health monitoring, gait rehabilitation, and emergency detection.

B. Edge AI and On-Device Inference

Edge AI research has progressed toward executing neu- ral networks on resource constrained microcontrollers via pruning, quantization, model distillation, and sparse inference techniques [10]. Frameworks such as TensorFlow Lite Micro and ARM CMSIS-NN enable simplified inference pipelines for embedded platforms. Recent edge-based models for fall detection, object classification, and physiological anomaly de- tection demonstrate that carefully designed convolutional and temporal models can operate within 100–300 kB of memory. Despite these advances, edge-only systems are limited by their inability to leverage cross user learning. Wearable be- havior varies significantly across individuals due to differ- ences in physiology, motion dynamics, and sensor placement. Without collaboration across devices, edge-only training often converges slowly and yields models with poor generalization.

C. Federated Learning for Mobile and IoT Systems

Federated learning was introduced to support decentral- ized training of machine learning models without uploading raw data [11]. FL has been widely applied to smartphones, IoT sensors, and autonomous vehicles [12], [13]. Federated Averaging (FedAvg) remains the most widely adopted algo- rithm, although later work addressed device heterogeneity, intermittent participation, non-IID data distributions, gradient staleness, and straggler effects.

FL on wearable devices is less explored. Wearables face stricter energy constraints, limited memory, and highly per- sonalized data distributions. Only a subset of recent work studies FL on smartwatches or fitness trackers [14], and these studies typically: (i) use single-modality IMU data, (ii) require a smartphone to act as a relay, or (iii) ignore privacy threats such as gradient leakage. None of these architectures integrate cloud–edge hierarchy, model compression, and on- device multi-modal learning simultaneously.

D. Privacy-Preserving Distributed Optimization

Numerous techniques have been proposed to safeguard user privacy in distributed learning systems [15]. Differential privacy (DP) adds calibrated noise to gradients or model updates, limiting the information that can be inferred about individual samples. Secure aggregation ensures that servers cannot access individual client updates, only their aggregated sum. Homomorphic encryption (HE) can further conceal up- date content, though at high computational cost unsuitable for wearables [16].

Gradient sparsification and model compression have also been used to reduce communication cost while implicitly improving privacy by revealing fewer model dimensions. However, combining DP, sparsification, and secure aggregation must be carefully designed to avoid degrading model performance, particularly under non-IID data.

E. *Limitations of Existing Work*

Across the existing literature, several gaps remain:

- Most wearable-related FL systems rely on smartphones as proxies rather than enabling fully autonomous learning on wearable-class microcontrollers [17].
- Multi-modal physiological and motion fusion in an FL setting is rarely explored due to alignment challenges and mismatched sampling frequencies.
- Communication efficient FL remains under-addressed in wearable systems, where uplink bandwidth is extremely limited.
- Few existing FL frameworks explicitly incorporate cloud edge hierarchical coordination for balancing computation, privacy, and thermal/battery constraints [18].

F. *Positioning of FedWear*

FedWear differs from prior work in several key ways:

- It supports local training and inference directly on wearable-class microcontrollers without reliance on smartphones [19].
- It integrates multi-modal sensor fusion (IMU + PPG + ECG + temperature) into a lightweight on-device model.
- It uses Top-*K* gradient sparsification, differential privacy, and secure aggregation simultaneously to reduce commu- nication cost and strengthen privacy.
- It introduces a cloud–edge–device hierarchy that balances global optimization with edge responsiveness [20].
- It demonstrates real-time performance (18–22 ms latency) and improved personalization compared to baseline FL approaches.

Collectively, these innovations position FedWear as a bridge between resource-efficient edge inference, privacy-preserving distributed learning, and scalable wearable intelligence.

III. FEDWEAR SYSTEM ARCHITECTURE

FedWear is organized into three tightly integrated tiers: the wearable device tier, an optional edge gateway tier, and a cloud coordinator tier. Fig. 1 provides an overview of the FedWear cloud–edge–device architecture.

A. *Wearable Device Tier*

Each wearable node continuously acquires synchronized multi-modal sensor data such as accelerometer and gyroscope readings for motion, PPG for blood volume changes, ECG for cardiac activity, and skin temperature. Local preprocessing includes resampling, filtering, segmentation into overlapping windows, and normalization. A compact neural architecture then performs both local training and real-time inference.

The local model consists of:

- 1D convolutional blocks that extract short-term temporal features from each modality.
- A lightweight sequence modeling stage (e.g., a small Transformer encoder or gated recurrent unit) thatcaptures longer-range temporal dependencies.
- A pooling and classification head that outputs activity labels or anomaly scores.

Each device maintains a buffer of recent windows for local training and periodically computes gradients on these mini- batches, subject to energy and CPU constraints.

B. *Edge Gateway Tier (Optional)*

When available, a smartphone or local hub can act as an intermediate gateway to:

- pre-aggregate updates from multiple nearby wearables,
- perform basic validation, compression, or batching,
- schedule uplink communication to the cloud to avoid congested or expensive network periods.

C. *Cloud Coordinator Tier*

The cloud tier is responsible for:

- orchestrating FL rounds (client selection, broadcasting the global model, collecting updates),

- secure aggregation of privacy-noised model updates,
- global optimization and model versioning,
- optional offline evaluation and hyperparameter tuning on de-identified or synthetic datasets.

Raw wearable signals are never transmitted to the cloud;

only masked and, optionally, differentially private gradients or parameter deltas are exchanged.

IV. FEDERATED LEARNING AND CLOUD INTEGRATION

A. Local Objective and Update

Let D $_i$ denote the local dataset on device i, with empirical

$$F_i(w) = \frac{1}{|D_i|}{}_{(x,y)\in D'} \quad \ell(f_w(x), y), \qquad (1)$$

where fw is the local model parameterized by w. Each client performs a few gradient descent steps:

$$w_i^{(t+1)} = w^{(t)} - \eta \nabla F_i(w^{(t)}), \qquad (2)$$

where w(t) is the global model at round t and η is the learning rate.

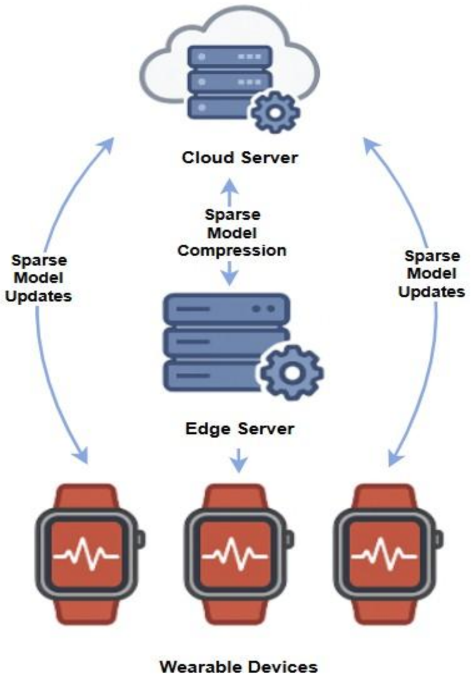

Fig. 1. FedWear system architecture

B. Cloud-Side Aggregation

The cloud receives updates $w_i^{(t+1)}$ (or deltas $w_i^{(t+1)} - w^{(t)}$) from a subset of participating clients S_t and computes:

$$w^{(t+1)} = {}_{i\in S_t} \frac{|D_i|}{{}_{j\in S_t}|D_j|} w_i^{(t+1)}. \qquad (3)$$

C. Differential Privacy and Secure Aggregation

To limit information leakage, each client perturbs its update using Gaussian noise:

$$\tilde{w}_i^{(t+1)} = w_i^{(t+1)} + \mathrm{N}(0, \sigma^2 I), \qquad (4)$$

where σ controls the privacy–utility trade-off. Secure aggre- gation protocols ensure that the cloud can only recover the aggregate ${}_i\tilde{w}_i^{(t+1)}$ and not any individual update.

D. Communication-Efficient Gradient Sharing

FedWear applies Top-*K* sparsification, keeping only the *K* highest-magnitude components of each gradient or pa- rameter delta and transmitting their indices and values. This significantly reduces bandwidth and energy consumption with minimal impact on learning.

E. Client Selection

Client participation is constrained by battery level, connec- tivity, and data availability. FedWear can incorporate a scoring function that prioritizes clients with sufficient energy, stable connections, and informative data diversity, enabling the cloud to select an appropriate subset of devices for each round.

V. EXPERIMENTAL EVALUATION

A. Hardware Platform

FedWear is evaluated using a prototype deployment on a Cortex-M7 class microcontroller (600 MHz, 512 kB SRAM, 1 MB Flash), representative of mid-range wearable proces- sors. All on-device inference and local training routines are implemented using TensorFlow Lite Micro, which provides a lightweight runtime without dynamic memory allocation. The federated learning coordinator runs on a cloud-based Python server simulating FL rounds, secure aggregation, and global optimization.

A total of 50 simulated wearable devices participate in the federated learning experiments unless otherwise stated. Each simulated device enforces:

- a compute limit of 40–60 ms per local update,
- a memory cap matching the Cortex-M7 SRAM budget,
- intermittent participation (70–85% availability),
- wireless uplink bandwidth constraints (0.5–2 Mbps).

The compact multi-modal model deployed on each device contains 182,304 trainable parameters after quantization (8-bit weights). Local batch sizes are restricted to 16 due to memory constraints.

B. Datasets

Experiments use three data sources to evaluate generaliza- tion under heterogeneous wearable conditions:

- UCI HAR (Human Activity Recognition): 30 subjects, 6 activity classes, 561-dimensional IMU features. Data are partitioned by user to mimic real personalization.
- Synthetic Physiological Anomaly Dataset: Constructed by injecting controlled perturbations (falls, tremor events, gait asymmetry, impact bursts) into IMU sequences. The dataset includes 4,500 anomaly windows spanning 7 subjects with different sensor placements.
- Multi-Modal Fusion Dataset: Combined IMU (50 Hz), PPG (100 Hz), and skin temperature (1 Hz) windows. A total of 80,000 fused windows are used, with each device receiving a unique, non-IID user distribution.

To simulate realistic heterogeneity, each device receives data from one or two users with distinct gait, motion intensity, and sensor alignment characteristics.

C. Baselines

We compare FedWear with the following learning setups:

- Cloud Centralized: Raw data from all devices are up- loaded and trained centrally.
- Edge Only: Each device trains independently with no global parameter sharing.

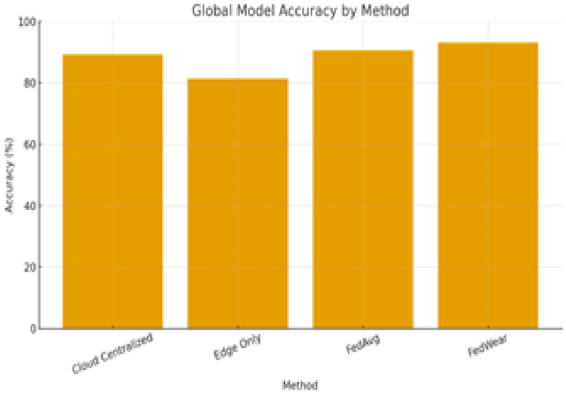

Fig. 2. Global model accuracy comparison TABLE I GLOBAL MODEL ACCURACY AND PERSONALIZATION

Model	Accuracy	F1	Personalization
Cloud Centralized	89.2%	0.88	—
Edge Only	81.4%	0.79	—
FedAvg	90.5%	0.89	+6%
FedWear	93.1%	0.92	+14%

- FedAvg: Standard federated averaging without sparsifica- tion or differential privacy.
- FedWear (Proposed): Lightweight local training, Top- K sparsification, and differential privacy. Experimental settings include:
  - Top-K sparsification ratio: K = 12%,
  - Gaussian DP noise: $\sigma = 0.45$,
  - FL rounds: 50 rounds (25 communication cycles),
  - Local epochs: 1 epoch per round,
  - Client fraction per round: 20%.

D. Accuracy and Personalization Results

Fig. 2 shows the accuracy comparison across different train- ing configurations. FedWear consistently outperforms both FedAvg and cloud-centralized training due to better person- alization and reduced overfitting on non-IID data.

Table I summarizes the global model performance. FedWear achieves the best personalization improvement due to local fine-tuning and privacy-preserving gradient exchange.

E. Convergence Behavior

Fig. 3 shows accuracy across FL rounds. FedWear converges faster due to sparsified, less noisy updates that avoid the gradient drift seen in FedAvg under non-IID data.

F. Communication Cost

Fig. 4 shows that FedWear reduces per-round commu- nication to 0.68 MB, a 68% improvement over FedAvg's

1.98 MB dense gradients. Cloud-centralized training requires tens of megabytes per round, making it infeasible for wearable deployments.

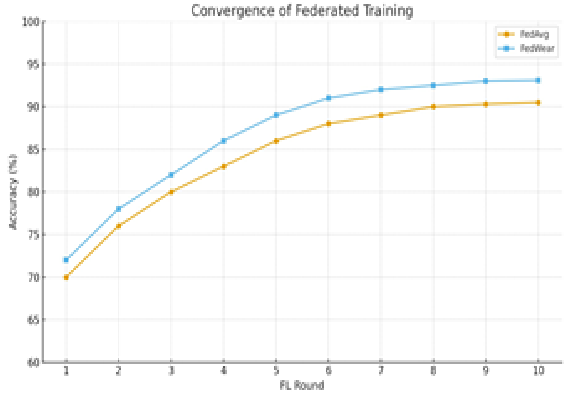

Fig. 3. Convergence of federated training

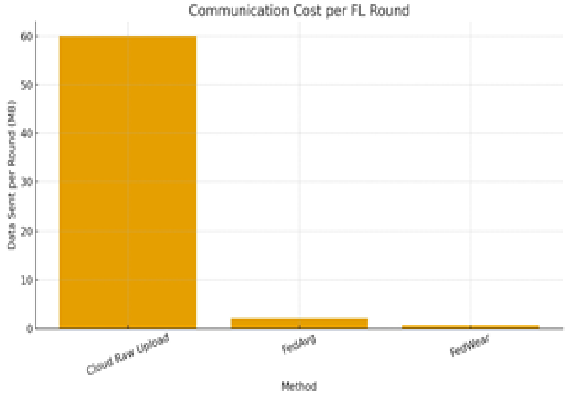

Fig. 4. Per-round communication cost comparison

G. On-Device Latency

Fig. 5 reports the end-to-end inference latency across con- figurations. FedWear maintains sub-25 ms latency, performing nearly 6 faster than cloud-only inference, which suffers from network round-trip delays.

VI. DISCUSSION

The evaluation of FedWear demonstrates that a hybrid cloud edge federated learning architecture can effectively balance privacy, latency, and model quality for wearable devices. By keeping raw motion and physiological signals on-device, Fed- Wear significantly reduces the risk of user-level data leakage while still enabling strong global generalization through cross- user collaboration. This is especially important for wearables, where data distributions vary widely due to differences in motion patterns, physiology, and sensor placement, making cloud-only or edge-only approaches insufficient.

A key advantage of FedWear is that it offloads computationally intensive operations such as global aggregation and model tuning to the cloud, while reserving real-time inference and lightweight local updates for the wearable. This aligns computation with the strengths of each tier: low- latency responsiveness at the edge, and scalable optimization in the cloud. The system's ability to maintain sub-25 ms inference latency shows that advanced learning methods can be deployed on microcontroller hardware without compromising user experience or energy efficiency.

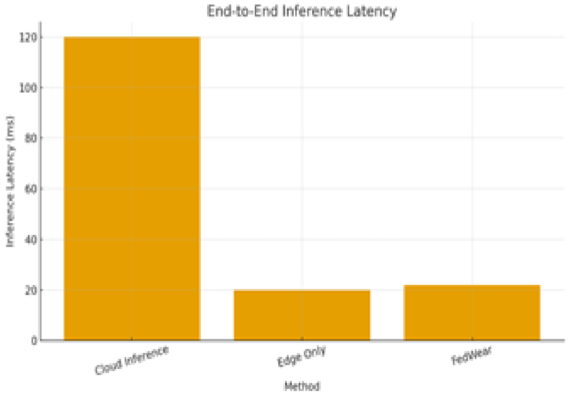

Fig. 5. End-to-end inference latency comparison

Despite these gains, several limitations remain. Non-IID data across users can slow global convergence, and the vari- ability of battery levels, connectivity, and sensing conditions may limit client participation in each FL round. Furthermore, differential privacy introduces a trade-off between noise and model accuracy, and excessive noise can degrade personalization for users with limited data. Optimizing DP noise schedules and adaptive client sampling is therefore critical for future deployments.

An additional consideration is the long-term sustainability of local training on wearable-class processors. While FedWear restricts computation to compact gradient updates and modest window sizes, the energy and thermal cost of periodic on- device training may still impact battery life in extended deployments. Future designs may benefit from multi-rate training policies, where the device adjusts its training frequency based on motion state, battery conditions, or model drift. Such adaptive strategies could further reduce the duty cycle of local compute while ensuring continuous improvement of the global model.

Overall, FedWear demonstrates that federated learning can be adapted to the constraints of wearable ecosystems while still delivering reliable, privacy-preserving analytics. Continued research is needed to refine communication efficiency, person- alization mechanisms, and adaptive scheduling for large-scale, real-world deployment.

VII. FUTURE WORK

Future extensions of FedWear aim to further improve ro- bustness, adaptability, and deployment practicality.

- On-device continual learning: Adaptation mechanisms that update the model gradually as user behavior or sen- sor conditions change, without requiring frequent cloud rounds.
- Hierarchical FL: Introducing a gateway layer (e.g., smart- phone or hub) for intermediate aggregation may further reduce uplink communication cost.
- Energy-aware training: Scheduling local updates based on battery level, thermal limits, and user context can improve real-world sustainability.
- Hardware acceleration: Mapping FedWear to DSP/TPU- like blocks in emerging wearable SoCs may yield addi- tional latency and power gains.
- Enhanced privacy mechanisms: Combining secure aggre- gation, tighter DP accounting, and lightweight encryption methods can strengthen protection against gradient inver- sion attacks.
- Longitudinal studies: Deploying FedWear across diverse user groups will help validate model robustness and personalization under real world variability.

VIII. CONCLUSION

This paper presented FedWear, a privacy-aware federated edge AI framework for wearable devices operating within a cloud–edge–device hierarchy. By combining on-device multi- modal learning, differential privacy, communication efficient gradient sharing, and cloud based secure aggregation, Fed- Wear achieves higher accuracy than cloud only or edge only baselines while substantially reducing data transmission and preserving real-time responsiveness. These results indicate that federated learning can be effectively adapted to wearable environments, enabling scalable and privacy conscious intelligence in future wearable ecosystems.

REFERENCES

[1] Y. Peng, J. Bian, and J. Xu, “FedMM: Federated Multi Modal Learn- ing with Modality Heterogeneity in Computational Pathology,” arXiv preprint arXiv:2402.15858, 2024.

[2] A. M. Kirubakaran, L. Butra, S. Malempati, A. K. Agarwal, S. Saha, and A. Mazumder, “Real-Time Anomaly Detection on Wearables using Edge AI,” International Journal of Engineering Research and Technology (IJERT), vol. 14, no. 11, Nov. 2025, doi: 10.17577/IJERTV14IS110345.

[3] C. Dong, J. Zhou, A. Yao, Z. Xu, F. Jiang, S. Chen, and X. Liu, “A Federated Multi Modal Learning Framework Powered by Distributed Ledgers for Cyber safe and Efficient UAV Delivery Systems,” in Proc. IEEE Int. Conf. Data Mining Workshops (ICDMW), 2023, pp. 1032– 1039, doi: 10.1109/ICDMW60847.2023.00136.

[4] V. Punniyamoorthy, A. G. Parthi, M. Palanigounder, R. K. Kodali, B. Kumar, and K. Kannan, “A Privacy-Preserving Cloud Architecture for Distributed Machine Learning at Scale,” International Journal of En- gineering Research and Technology (IJERT), vol. 14, no. 11, Nov. 2025.

[5] B. Ramdoss, A. M. Kirubakaran, A. M. Kirubakaran, P. B. S., S. H. C., and V. Vaidehi, “Human fall detection using accelerometer sensor and visual alert generation on Android platform,” SSRN, Mar. 7, 2014, doi: 10.2139/ssrn.5785544.

[6] Z. Xin, Q. Li, C. Niu, F. Wu, and G. Chen, “Federated multi task learning with cross device heterogeneous task subsets,” Journal of Parallel and Distributed Computing, vol. 205, p. 105155,

2025, doi: 10.1016/j.jpdc.2025.105155.

[7] S. G. Aarella, V. P. Yanambaka, S. P. Mohanty, and E. Kougianos, "Fortified-Edge 2.0: Advanced Machine-Learning-Driven Framework for Secure PUF-Based Authentication in Collaborative Edge Comput- ing," Future Internet, vol. 17, p. 272, 2025, doi: 10.3390/fi17070272.

[8] D. Khuong, A. D. Nguyen, D. Nguyen, and K. S. Wong, "Fed- KoE: Enhancing Federated Multimodal Learning through Knowledge of Experts," in Proc. Int. Workshop on Secure and Efficient Feder- ated Learning (FL AsiaCCS '25), 2025, Art. no. 4, pp. 1–6, doi: 10.1145/3709023.3737690.

[9] S. G. Aarella, A. K. Tripathy, S. P. Mohanty, and E. Kougianos, "EasyBand2.0: A Framework with Context-Aware Recommendation Mechanism for Safety-Aware Mobility during Pandemic Outbreaks," in Proc. 23rd Int. Symp. Quality Electronic Design (ISQED), Santa Clara, USA, 2022, pp. 1–6, doi: 10.1109/ISQED54688.2022.9806250.

[10] Y. M. Lin, Y. Gao, M. G. Gong, S. J. Zhang, Y. Q. Zhang, and Z. Y. Li, "Federated learning on multimodal data: A comprehensive survey," Machine Intelligence Research, pp. 1–15, Jun. 2023, doi: 10.1007/s11633-022-1398-0.

[11] P. K. Veerapaneni, "Building scalable AI-powered analytics pipelines using Delta Live Tables: A cybersecurity-first approach," International Journal of Computer Engineering and Technology (IJCET), vol. 14, no. 2, pp. 301–314, 2023.

[12] S. M. Sheikholeslami, P. C. Ng, J. Abouei, and K. N. Plataniotis, "Multi- Modal Federated Learning Over Cell-Free Massive MIMO Systems for Activity Recognition," IEEE Access, vol. 13, pp. 40844–40858, 2025, doi: 10.1109/ACCESS.2025.3548001.

[13] S. G. Aarella, S. P. Mohanty, E. Kougianos, and D. Puthal, "PUF- based Authentication Scheme for Edge Data Centers in Collabo- rative Edge Computing," in Proc. IEEE Int. Symp. Smart Elec- tronic Systems (iSES), Warangal, India, 2022, pp. 433–438, doi: 10.1109/iSES54909.2022.00094.

[14] P. K. Veerapaneni, "Federated learning and artificial intelligence in healthcare: A privacy-preserving approach for medical data," International Journal of Research in Computer Applications and Information Technology (IJRCAIT), vol. 6, no. 1, pp. 107–120, 2023.

[15] K. Borazjani, P. Abdisarabshali, F. Nadimi, N. Khosravan, M. Liwang,X. Wang, Y. Hong, and S. Hosseinalipour, "Multi Modal Multi Task (M3T) Federated Foundation Models for Embodied AI: Potentials and Challenges for Edge Integration," IEEE Internet of Things Magazine, pp. 1–11, 2025, doi: 10.1109/MIOT.2025.3604330.

[16] A. Nagpal, B. Pothineni, A. G. Parthi, D. Maruthavanan, A. R. Banarse, P. K. Veerapaneni, S. R. Sankiti, and V. Jayaram, "Framework for automating compliance verification in CI/CD pipelines," International Journal of Computer Science and Information Technology Research (IJCSITR), vol. 5, no. 4, pp. 17–27, 2024, doi: 10.5281/zenodo.1425967.

[17] T. Fu, J. Hu, G. Min, S. A. Khowaja, K. Singh, and K. Dev, "Feder- ated Retrieval-Augmented Generation-Based LLM for Enhanced Cyber Threat Detection in the Internet-of-Energy," IEEE Network, 2025, doi: 10.1109/MNET.2025.3619162

[18] Parlapalli, B. Pothineni, A. G.Parthi, P. K. Veerapaneni, D. Marutha-vanan, A. Nagpal, R. K. Kodali, and D. M. Bidkar, "From complexity to clarity: One-step preference optimization for high-performance LLMs," Int. J. Artif. Intell. Mach. Learn. (IJAIML), vol. 4, no. 1, pp. 112–125, 2025, doi: 10.34218/IJAIML 04 01 008.

[19] C. Lee, S. Cho, H. Park, J. Park, and S. Lee, "GAN-Enhanced Vertical Federated Learning System for WHAR with non-IID Data," in NOMS 2024—IEEE Network Operations and Management Symposium, Seoul, Korea, 2024, pp. 1–5, doi: 10.1109/NOMS59830.2024.10575811.

[20] N. Chockalingam, A. Chakrabortty, and A. Hussain, "Mitigating Denial- of-Service attacks in wide-area LQR control," in Proc. 2016 IEEE Power and Energy Society General Meeting (PESGM), 2016, pp. 1–5. doi: 10.1109/PESGM.2016.7741285.