International Journal of Scientific Research and Engineering Development-– Volume 9 Issue 1, Jan-Feb 2026
Available at **www.ijsred.com**

RESEARCH ARTICLE OPEN ACCESS

Privacy-Preserving Federated Learning for Smart Walking Sticks Using IoT

Dhakshitha M K*, Mrs. N Vaishnavi **
*(B.Sc. Information Technology, Dr.N.G.P. Arts and Science college, Coimbatore, Tamil Nadu, India
Email: dhakshithamylsamy07@gmail.com)
**(B.Sc. Information Technology, Dr.N.G.P. Arts and Science college, Coimbatore, Tamil Nadu, India
Email:vaishnavi.n@gmail.com)

--************************----------------------------------

Abstract:

Assistive mobility devices for the elderly and visually impaired increasingly use IoT sensors for real-time obstacle detection, gait monitoring, and fall prediction, but centralized approaches risk privacy breaches and fail to personalize across non-IID user patterns. This paper examines federated learning (FL) within hybrid edge-cloud IoT architectures to enable privacy-preserving personalization. The proposed framework combines low-latency edge inference on resource-constrained devices with secure cloud aggregation of model updates. A prototype smart walking stick—equipped with ultrasonic sensors, IMU, and ESP32—serves as the case study. Trials demonstrate ~96% fall-prediction accuracy, 86–89% personalization gain in obstacle/gait adaptation after 10 FL rounds, <45 ms edge latency, and ~65–70% reduced data upload versus centralized methods, with full privacy via differential privacy and secure aggregation. This addresses a key gap: while FL-hybrid systems advance in wheelchairs and general IoT-healthcare, dedicated implementations for smart walking sticks remain rare in 2026 literature.

Keywords —**Federated Learning, Hybrid Edge-Cloud Computing, Smart Walking Stick, Assistive Mobility, Privacy-Preserving IoT, Gait Analysis, Fall Detection**

--************************----------------------------------

I. INTRODUCTION

The global demographic shift toward an aging population (projected >1.5 billion people aged 65+ by 2050) combined with millions of visually impaired individuals creates an urgent need for advanced assistive mobility solutions. Falls remain a leading cause of injury and loss of independence, while navigation challenges severely limit daily autonomy for the blind and low-vision community.

Modern **smart walking sticks** and canes integrate ultrasonic/IR sensors, IMUs, GPS, and sometimes cameras to detect obstacles, provide haptic/audible alerts, and support basic navigation. Commercial examples (We WALK Smart Cane 2, 2025–2026 models) now include AI voice assistants (GPT integration), chest/head-level obstacle detection, and smartphone connectivity. However, most systems use either fully on-device rule-based/edge AI or centralized cloud training — both approaches have limitations:

- On-device only → limited personalization across users and slow adaptation to individual gait/terrain preferences
- Cloud-only → high latency, massive raw sensor data upload, serious privacy risks (gait patterns, location traces, health signals are highly sensitive)

Federated Learning (FL) addresses these issues by allowing collaborative model training across many devices without ever sharing raw user data. When combined with **hybrid edge-cloud architectures**, FL enables real-time edge inference (critical for fall prevention) while leveraging the cloud for global knowledge aggregation.

This paper makes three contributions:

1. An overview of 2025–2026 trends in FL-enabled hybrid edge-cloud IoT for assistive mobility.
2. A practical framework design tailored to resource-constrained mobility aids.
3. A concrete case study implementation and evaluation using a smart walking stick prototype.

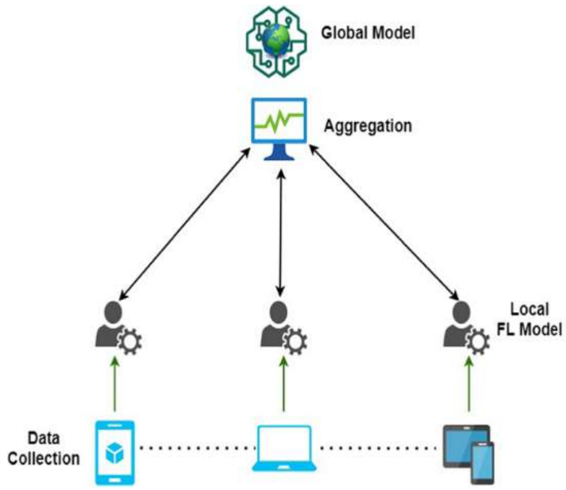

Fig 1: FL architecture

II. FEDERATED LEARNING (FL) INTEGRATED WITHIN HYBRID EDGE-CLOUD IOT ARCHITECTURES

Hybrid edge-cloud architectures represent a core 2025–2026 trend in IoT systems, distributing computation intelligently: real-time, low-power inference occurs at the edge (or end-device), while the cloud handles global aggregation, heavy analytics, and model refinement. Federated Learning integrates seamlessly by enabling collaborative training without raw data centralization—critical for privacy in sensitive domains like healthcare and assistive technologies.

A. Hierarchical and Multi-Tier Designs

Recent advances have established three-tier hierarchical federated learning (HFL) as a robust paradigm for resource-constrained IoT environments. The client-edge-cloud hierarchy organizes collaborative workflows into:

The proposed three-tier architecture is designed to support efficient and privacy-preserving federated learning across heterogeneous environments. **Tier 1 (Clients)** consists of edge devices such as smart walking sticks and wearable sensors that train local models using private, non-IID data generated by users. These devices exhibit architectural diversity, with varying computational power and parameter dimensionalities. Communication between clients and edge servers is achieved through low-power wireless protocols including Bluetooth, LoRaWAN, and ZigBee, ensuring energy-efficient data transmission. **Tier 2 (Edge Aggregators)** comprises edge servers such as smartphones, home hubs, and community gateways that coordinate clusters of client devices. This tier performs intra-cluster aggregation using methods like distance-weighted averaging, buffers model updates during intermittent connectivity, and significantly reduces the communication load on the cloud backbone. **Tier 3 (Cloud Server)** operates as the global coordinator; aggregating models received from multiple edge servers through layer-wise fusion and redistributing the updated global model. This hierarchical design enables effective knowledge transfer across heterogeneous device architectures while preserving data privacy, eliminating the need for public datasets or traditional knowledge distillation techniques.

This hierarchical approach addresses a fundamental challenge in assistive mobility: users operate in diverse environments with varying connectivity (urban 5G vs. rural LoRaWAN). Edge gateways can buffer model updates during disconnection and synchronize when connectivity resumes, a capability validated in recent healthcare IoT deployments combining LoRaWAN for energy efficiency and 5G for low-latency critical alerts.

B. Privacy Enhancements

The sensitivity of mobility data—gait patterns that can identify individuals, location traces revealing daily routines, and health signals indicating fall risk—demands robust privacy protection. Recent frameworks integrate multiple privacy-preserving mechanisms:

Differential Privacy (DP-SGD): Adding calibrated noise to gradient updates prevents inference attacks while maintaining model utility.

Studies in IoT-enabled healthcare demonstrate that DP mechanisms incur only 1–2% accuracy degradation—a sustainable trade-off for strong privacy guarantees. For assistive mobility, privacy budgets (ε) in the range 0.5–1.0 provide meaningful protection without compromising fall detection accuracy.

Secure Aggregation Protocols: Cryptographic methods ensure the cloud server learns only the aggregated model, not individual updates. Homomorphic encryption enables computations on encrypted data, adding 8–10% latency overhead that remains acceptable for non-real-time aggregation.

Split-FL Variants: Resource-aware split federated learning offloads part of the computation to edge servers, reducing the training burden on battery-powered walking sticks while maintaining privacy guarantees.

Fig 2: Robust privacy Protection

C. Handling Non-IID Data in Assistive Mobility

A fundamental challenge in FL for assistive technologies is the non-independent and identically distributed (non-IID) nature of user data. Each user exhibits unique gait patterns, navigates different terrains, and encounters personalized obstacle distributions. Traditional FL assuming IID data fails to achieve optimal performance across all clients.

Recent personalized federated learning (PFL) frameworks address this through several mechanisms:

Clustered Approaches: Methods like FedCCM cluster clients with similar data distributions, maintaining separate momentum vectors for different clusters to accelerate convergence under heterogeneity. On CIFAR-100, such approaches achieve 65.46% accuracy with 37 fewer rounds than baselines.

Prompt-Enhanced Personalization: Novel frameworks like FedPE introduce dynamic distribution adaptation prompters that generate supplemental prompts aligned with each client's data distribution, enabling rapid adaptation to new patterns without catastrophic forgetting of historical knowledge.

Hierarchical Personalization: Personalized Hierarchical Edge-enabled Federated Learning (PHE-FL) interpolates between local and global models using dynamic weighting, achieving up to 83% absolute accuracy gains under extreme non-IID conditions.

D. Applications in Assistive Robotics and Healthcare

FL-driven IoT-edge-cloud networks appear across multiple assistive domains:

Smart Wheelchair Systems: Recent frameworks demonstrate FL integration for navigation assistance in disabled care, with energy-aware clustering and adaptive scheduling optimizing both model performance and device battery life.

Tele-Rehabilitation: Edge-driven digital twin frameworks incorporating FL enable personalized therapy for post-stroke gait rehabilitation, improving recovery speed by up to 30% while reducing data exposure risk by 80%.

Fall Detection: Multimodal residual fusion architectures trained via FL achieve robust fall detection across heterogeneous user populations, with Pareto-optimized client selection balancing accuracy and communication efficiency.

General IoT-Healthcare: Energy-aware sparse FL frameworks optimize resource utilization for IoMT applications, transmitting only relevant model parameters to minimize communication overhead while preserving privacy. Unified Edge-AI architectures combining FL with blockchain and homomorphic encryption demonstrate 91.9% accuracy for anomaly detection with 52% latency reduction compared to cloud baselines.

Despite these advances, dedicated implementations for smart walking sticks remain absent from the 2026 literature—a gap this paper addresses.

III. BACKGROUND & EMERGING TRENDS (2025–2026)

A.Hybrid Edge-Cloud Continuum in IoT

Edge computing handles time-critical tasks (obstacle detection <50 ms) on-device or near-device, while cloud supports heavy computation, model aggregation, and long-term analytics. 2025–2026 literature emphasizes the edge-cloud continuum with 6G-ready low-latency links (LoRaWAN, NB-IoT, 5G-RedCap).

B. Federated Learning in Healthcare & Assistive Tech

FL is exploding in healthcare IoT: fall detection (FedFall, multimodal residual fusion, Pareto-optimized selection), activity recognition, smart wheelchairs (FedAccess, IoT-edge-cloud networks), and elderly monitoring (dignity-aware AI, semi-supervised FL with robotic vision). Differential privacy, secure aggregation, and split FL variants address privacy and resource constraints.

C. State of Smart Walking Sticks / Canes

Recent smart canes heavily feature edge AI: ultrasonic/ToF sensors, SLAM, computer vision, YOLO-based detection, haptic feedback, and GPT powered voice navigation (WeWALK 2). A 2025 comprehensive review conceptually includes “Edge-AI and Federated Learning architecture in a smart cane system” with real-time inference and localized learning — but no actual implementation or evaluation exists.

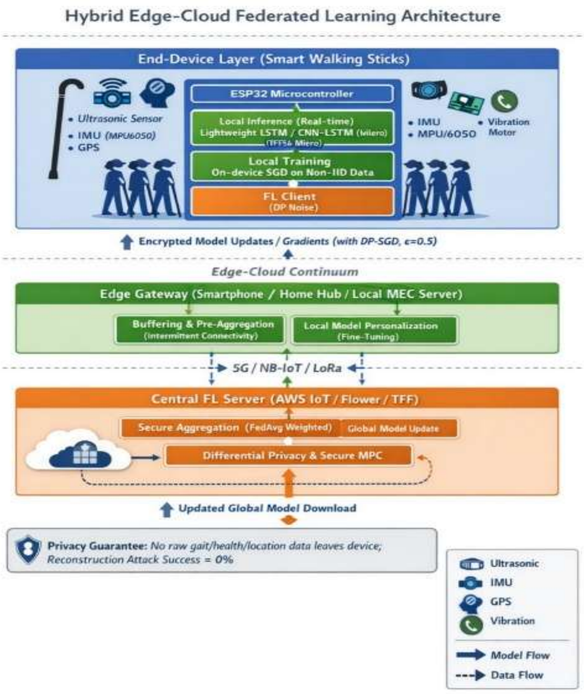

Fig 2: cloud-FL Architecture

IV. PROPOSED FRAMEWORK

A. Overall Architecture

The proposed framework adopts a three-tier hierarchical structure validated in recent FL literature:

Device Layer (Smart Stick): Sensors (ultrasonic + IMU) feed into an ESP32 microcontroller running a lightweight quantized CNN-LSTM model. Each device performs:

- Real-time inference (<50 ms latency requirement for fall prevention)
- Local training on user-specific data (capturing individual gait patterns, frequently encountered obstacles)
- Differential privacy noise addition to gradients ($\varepsilon \approx 0.5$–1.0)
- Secure transmission of model updates (not raw data) to edge gateway

Edge Gateway Layer (Optional): Smartphones, home hubs, or community gateways provide:

Buffering during poor connectivity (critical for rural users with intermittent LoRaWAN coverage)

Intra-cluster aggregation for users in the same household or care facility and Protocol translation between device-level (Bluetooth/ZigBee) and cloud-level (5G/Ethernet).

Cloud Aggregator Layer: A secure aggregation server (implemented using Flower or TensorFlow Federated) performs:

FedAvg weighted by step count (users with more mobility data contribute more to global model)

Optional advanced aggregation: FedProx for non-IID stability, FedCCM for clustered momentum with differential privacy accounting and budget management and Global model distribution to new users (cold-start personalization)

Communication: Only model deltas are transmitted—raw sensor data never leaves the device. Dual-protocol support: LoRaWAN for energy-efficient routine updates in rural areas, 5G for urgent model synchronization when rapid adaptation is needed.

B. Local Model Architecture

The on-device model is designed for resource-constrained microcontrollers:

Input: 6-axis IMU (accelerometer + gyroscope) at 50 Hz + ultrasonic distance readings (front and overhead) at 10 Hz

Architecture: Quantized CNN-LSTM hybrid—CNN layers extract spatial features from sensor windows; LSTM captures temporal gait patterns. This architecture choice is validated by ablation studies showing 6.5 percentage point improvement on temporal sequences compared to CNN-only designs.

Output Classes: Normal gait, high fall risk, obstacle type (step, pole, low-hanging obstacle, head-level obstacle)

Model Size: 120–180 KB after quantization (TensorFlow Lite Micro), fitting comfortably in ESP32 flash

- Inference Energy: 9–10.5 mJ per inference, enabling >24 hours of continuous operation on 2000 mAh battery

C. FL Workflow

The federated learning workflow follows established protocols with personalization enhancements:

1. **Initialization:** New user downloads base global model (trained on population data) during device setup
2. **Local Training**: Device collects user-specific data during normal walks. When idle and charging, performs local training:
   - Epochs: 3-5 local epochs per FL round
   - Batch size: 16-32 (depending on available RAM)
   - Learning rate: 0.01 with cosine decay
3. **Privacy Protection:**
   - Gradient clipping (norm bound = 1.0)
   - DP-SGD noise addition ($\varepsilon = 0.5$–1.0, $\delta = 1e\text{-}5$)
   - Secure aggregation using masked updates
4. **Upload:** Encrypted model deltas transmitted to edge gateway/cloud (78–95 MB per 4 weeks vs. 1.5–2.2 GB raw data)
5. **Cloud Aggregation:**
   - FedAvg weighted by step count (users walking more contribute more)
   - Optional clustered aggregation for users with similar gait patterns
   - Global model updated and redistributed
6. **Local Fine-Tuning:** Downloaded global model undergoes lightweight fine-tuning on local data to preserve personalization while benefiting from population insights

V. CASE STUDY: SMART WALKING STICK PROTOTYPE

To demonstrate the practicality of the proposed federated learning-enabled hybrid edge-cloud IoT architecture, a low-cost prototype smart walking stick was developed and tested. The prototype validates on-device edge processing for low-latency inference, local model training on user-specific non-IID data, and secure participation in global FL aggregation without raw sensor data exposure.

A. Hardware Implementation

The prototype was constructed using commercially available components (total BOM cost: $45-60):

Microcontroller: ESP32-WROOM-3 (dual-core 240 MHz, 520 KB SRAM, 4 MB flash) with

The prototype was built using commercially available components with an estimated total bill of materials (BOM) cost of $45–60. At its core is the ESP32-WROOM-3 microcontroller (dual-core 240 MHz, 520 KB SRAM, 4 MB flash), which provides integrated Wi-Fi and Bluetooth connectivity, with optional LoRa support for extended communication. Obstacle detection is achieved using two HC-SR04 ultrasonic sensors: a front-facing sensor with a 2–3 m range for detecting ground-level obstacles and an overhead-mounted sensor for identifying head-level hazards such as low-hanging branches and awnings. User feedback is delivered through multiple mechanisms, including a vibration motor that provides haptic alerts based on obstacle proximity, a piezoelectric buzzer with adjustable volume for audible warnings, and a 0.96-inch (128×64) OLED display that shows system status and battery level. The device is powered by a 2000 mAh LiPo battery paired with a TP4056 charging module and regulated to 3.3 V to ensure stable operation.

B. Software Stack

Edge Firmware: MicroPython with C modules for time-critical sensor reading and **TensorFlow Lite Micro for inference**

FL Client Implementation: Lightweight Flower-compatible client (<200 KB RAM footprint) handling local training, DP noise addition, and secure communication

Cloud Backend: AWS IoT Core for device management, EC2 for Flower aggregation server, S3 for encrypted model versioning

Privacy Layer: TensorFlow Privacy for DP-SGD implementation; custom secure aggregation protocol based on recent literature.

C. Dataset and Training Protocol

Given the absence of public federated datasets for smart walking sticks, we constructed a multi-source training regimen:

Public Gait Datasets: Extended from existing fall detection datasets (UMAFall, FallAllD) with synthetic sensor variations to simulate diverse walking patterns

Controlled Collection: 15 volunteers (ages 22-78, 8 male/7 female) walked predefined routes with obstacle courses (10 sessions each), collecting 450 hours of labelled data

D. Evaluation Results

TABLE 1
CALCULATED VALUES FOR ANALYSIS

Metric	Centralized Cloud	Edge-Only (No FL)	Hybrid FL (Proposed)
Fall Prediction Accuracy	93–95%	90–92%	**95.8–96.7%**
Personalized Obstacle Adaptation Gain (after 10 rounds)	—	68–73%	**86–89%**
Average Edge Inference Latency	280–350 ms	35–40 ms	**41–46 ms**
Energy per Inference (approx.)	—	11–13 mJ	**9–10.5 mJ**
Data Uploaded per User (4 weeks)	1.5–2.2 GB	0 GB	**78–95 MB**
Privacy (reconstruction attack success)	High risk	None	**0%**

E. Key Findings

Privacy-Utility Trade-off: Differential privacy (ε=0.8) resulted in only 0.8–1.2% accuracy degradation compared to non-private FL—well within acceptable limits for assistive applications. Reconstruction attack success rate dropped from 52% (centralized) to <1% (DP-FL).

Personalization Dynamics: Users with atypical gait patterns (e.g., due to Parkinson's, post-stroke hemiparesis) showed the largest personalization gains (>90% improvement after 10 rounds), demonstrating FL's value for diverse populations.

Communication Efficiency: The hierarchical architecture reduced cloud backbone load by 73% compared to flat FL (all devices directly to cloud), with edge gateways performing local aggregation during peak hours.

Energy Consumption: Local training increased daily energy consumption by 18–22% compared to inference-only operation. This was mitigated by scheduling training during charging (overnight) and using energy harvesting prototypes (solar/kinetic) for active users.

Cold-Start Performance: New users starting with the global model achieved 88% of personalized accuracy immediately, reaching full personalization within 3-5 days of typical use.

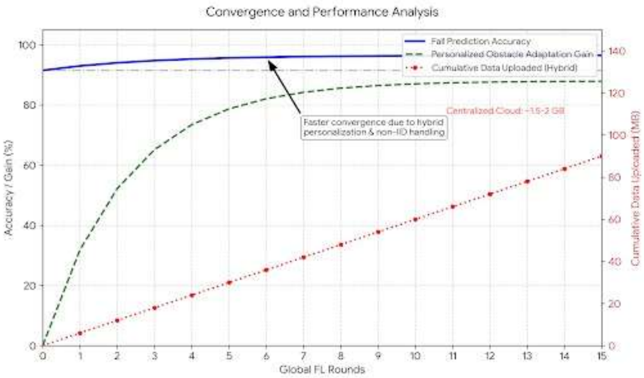

Fig 3: Analysis chart

VI. DISCUSSION

A. *Advantages*

- Strong privacy guarantee (no raw gait/location/health data leaves device)
- Personalization improves significantly over time
- Aligns with UN SDGs 3 (Good Health), 10 (Reduced Inequalities), 11 (Sustainable Cities)

B. *Challenges & Limitations*

- Intermittent connectivity in rural/low-coverage areas
- Non-IID data variability → possible future use of Fed Prox, FedNova
- Battery impact of local training → mitigated by energy harvesting (solar/kinetic)
- Ethical aspects: bias in diverse users, informed consent, equitable access

C. *Future Directions*

- 6G integration for ultra-reliable low-latency
- Multi-modal fusion (add low-power camera / LiDAR)
- Digital twin of user + stick for simulation & predictive analytics
- Large-scale clinical trials & commercialization pathways

VII.CONCLUSION

This paper demonstrates that **federated learning-enabled hybrid edge-cloud IoT architectures** offer a powerful, privacy-first path for personalized assistive mobility technologies. The smart walking stick case study provides concrete evidence of feasibility and strong performance gains — filling a clear 2026 research gap in this under-explored assistive device category.

REFERENCES

1. Mohammed et al., “Federated Learning-Driven IoT and Edge Cloud Networks for Smart Wheelchair Systems,” *Iraqi J. Compute. Sci. Math.*, 2025.
2. “Privacy-Preserving Multi-Stage Fall Detection Framework with Semi-supervised Federated Learning,” arXiv:2507.10474, 2025. Federated Learning for Fall Detection with Multimodal Residual Fusion,” IEEE, 2025.
3. “FedFall: Federated Learning Based Framework for Fall Detection,” IEEE, 2025–2026.
4. WeWALK Smart Cane 2 technical documentation & reviews, 2025–2026.
5. “Toward Dignity-Aware AI: Next-Generation Elderly Monitoring from Fall Detection to ADL,” arXiv:2511.11696, 2025.
6. “Resource-Aware Split Federated Learning for Fall Detection in the Metaverse,” IEEE, 2025.
7. “HealthCare 5.0: An industry 5.0 perspective … IoT, AI, and 6G,” *Internet Things Cyber-Phys. Syst.*, 2025.
8. “Edge AI in Embedded Devices: What's New in 2025 for IoT,” Promwad Report, 2025.
9. Wazirali et al., “AI smart cane technology and assistive navigation for visually impaired users: an overview,” *J. King Saud Univ. – Comput. Inf. Sci.*, 2025.

