Poster Abstract: Privacy–Generalisation Trade-offs in Federated Learning for Cross-Room Wi-Fi Sensing

Rod Koo[1], Thomas Gosman[1], Deepak Mishra[1], Gustavo Batista[2], Hassan Habibi Gharakheili[1], Aruna Seneviratne[1]

[1]School of Electrical Engineering & Telecommun., University of New South Wales, Sydney, NSW 2052, Australia

[2]School of Computer Science and Engineering, University of New South Wales, Sydney, NSW 2052, Australia

Emails: {rod.koo, t.gosman, d.mishra, g.batista, h.habibi, a.seneviratne}@unsw.edu.au

***Abstract*—Wi-Fi Channel State Information (CSI) enables device-free activity recognition in privacy-sensitive environments such as hospitals and smart buildings, where camera-based sensing is undesirable and raw data sharing is restricted. However, models trained in one room often degrade sharply in unseen environments due to layout changes, multipath variation, and interference. Federated learning (FL) enables on-device retention of raw CSI while supporting collaborative training. However, standard aggregation methods struggle with heterogeneous, non-identical signal distribution conditions, which in CSI are driven by propagation physics rather than semantic variation. This poster investigates whether server-side adaptive aggregation, without modifying client models or sensing pipelines, can improve cross-room robustness. Using the SHARP multi-room benchmark in a federated setup, replacing FedAvg with FedOpt-Adam improves unseen-room accuracy to 71.11%, a +7.49% gain over FedAvg and within 2.34% of centralised training, while preserving data locality. We further compare against non-FL generalisation baselines to contextualise these gains. Finally, we evaluate memory usage and communication cost on Raspberry Pi 4-class hardware and incorporate these parameters into NVFLARE simulations. The results show that adaptive server aggregation narrows the cross-room generalisation gap, highlighting the promise and limits of FL for deployable, privacy-preserving Wi-Fi sensing under real-world distribution shifts.**

***Index Terms*—Federated Wi-Fi sensing, edge AI, privacy-aware**

I. Introduction

Wi-Fi Channel State Information (CSI) has emerged as a practical modality for device-free human activity recognition in privacy-sensitive spaces such as hospitals, offices, and smart buildings. Momentum is growing around the IEEE 802.11bf sensing standard, positioning Wi-Fi as a scalable sensing infrastructure [1]. Yet despite this progress, CSI-based machine learning models remain highly environment-dependent. Even modest changes in layout, multipath structure, interference patterns, or transceiver placement can reshape temporal–spectral motion signatures, leading to sharp performance degradation when models are deployed in unseen rooms, limiting scalability. Pooling raw CSI across rooms can partially mitigate these distribution shifts, but centralised aggregation introduces governance and privacy concerns. Although CSI does not contain visual imagery, it implicitly encodes occupancy patterns, motion timing, and environmental characteristics that organisations may be unwilling or unable to share across sites. Federated Learning (FL) offers a promising alternative by enabling collaborative training while retaining raw CSI locally [2]. But standard aggregation strategies struggle under the strong non-independent and identically distributed (i.i.d.) conditions inherent to cross-room CSI, where each client's updates reflect its unique propagation environment.

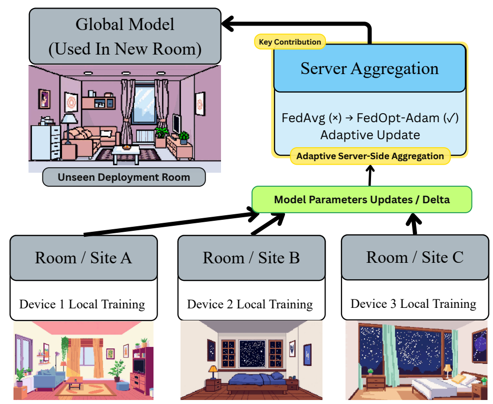

Fig. 1: Proposed Federated Wi-Fi sensing workflow where clients train locally on CSI features and server applies adaptive aggregation.

This poster reframes cross-room Wi-Fi sensing as a privacy–generalisation trade-off problem. Rather than modifying client models or sensing pipelines, we isolate a systems-level design lever (server-side adaptive aggregation) and examine whether replacing FedAvg with FedOpt-Adam can meaningfully reduce unseen-room performance degradation while preserving data locality. Using the SHARP multi-room benchmark and Raspberry Pi 4-class edge profiles, we quantify: (i) the magnitude of cross-room degradation, (ii) the extent to which adaptive aggregation narrows this gap relative to centralised and non-FL generalisation baselines, and (iii) the communication and resource overhead required to achieve it. By explicitly measuring these trade-offs, we provide novel insights into when federated Wi-Fi sensing improves deployability, where its limitations remain, and how aggregation design shapes generalisation across different environments.

II. System Description and Federated Setup

A. Cross-Room Federated Configuration

Fig. 1 illustrates the cross-room federated learning setup. Each room functions as an independent client with locally collected CSI shaped by its unique multipath geometry. Raw CSI is segmented into 1-second windows and transformed into

1000×30 Doppler spectrograms using the SHARP pipeline [2] to ensure consistent feature extraction across environments. Clients then train identical lightweight classifiers and send only model-parameter deltas to the server, keeping all CSI strictly on-device. Federated training proceeds in synchronous rounds. The server broadcasts the current global model; each room performs four local epochs of optimisation (batch size: 32; learning rate: 10^{-4}) and returns updated parameters for aggregation. Training is performed for 20 federated rounds.

To isolate coordination effects under heterogeneous CSI, we compare standard FedAvg against FedOpt-Adam, a server-side adaptive aggregator that applies momentum and second-moment scaling to reduce the impact of unstable room-specific updates. Crucially, client models, sensing pipelines, and hyperparameters remain unchanged, allowing aggregation to be evaluated as a systems-level design lever.

B. Experimental Design, Baselines and System Profiling

We adopt a leave-one-room-out evaluation to quantify cross-room generalisation. Baselines include: (i) Local-only (no collaboration), (ii) Centralised pooled training (upper bound), (iii) FedAvg, FedProx, SCAFFOLD, (iv) FedOpt-Adam (adaptive aggregation), (v) Non-FL generalisation techniques (fine-tuning and statistical alignment). This design enables direct measurement of the privacy–generalisation trade-off, comparing FL against both centralised training and domain-adaptation baselines rather than only against other FL variants.

The evaluation is designed to reflect a realistic indoor deployment condition, where models trained in a limited number of rooms must operate reliably across varying room layouts and environmental changes. The five-class activity recognition task captures both static (Empty room, Sitting) and dynamic (Walking, Running, Jumping) motion patterns, providing a balanced test of motion sensitivity and robustness. By evaluating performance on both seen rooms and a held-out room, we assess how well the trained global model transfers to new physical environments, which is a key requirement in practical multi-room Wi-Fi sensing deployments.

To quantify deployability, we profile compute latency, memory footprint, and communication cost on Raspberry Pi 4-class devices and inject these measurements into NVFLARE simulations. This provides a realistic characterisation of per-round latency and communication overhead using the same model sizes and data flows as physical clients.

III. Preliminary Results & Trade-Off Analysis

A. Cross-Room Generalisation Gap

Table I summarises cross-room performance. Centralised training achieves 73.45% unseen-room accuracy, representing the upper bound. FedAvg exhibits substantial degradation, achieving only 63.62% accuracy. Substituting FedAvg with FedOpt-Adam yields a +7.49-point improvement on unseen-room performance and narrows the gap to centralised training to 2.34 points. This quantifies a clear privacy–generalisation trade-off: adaptive FL substantially narrows cross-room degradation while maintaining strict on-device data locality, although measurable gap remains relative to centralised pooling.

TABLE I: Cross-environment performance.

Method	**Train**	**Alter**	**Unseen**	**F1**
Centralized (early stop)	98.90	84.19	73.45	0.74
FedOpt-Adam [3]	89.21	79.31	71.11	0.71
FedAvg [4]	86.20	79.06	63.62	0.62
FedProx [5]	89.06	74.54	60.40	0.60
SCAFFOLD [6]	84.42	74.69	58.87	0.59

B. Comparison, Convergence Stability and System Overhead

Including fine-tuning and statistical alignment baselines shows that while domain-adaptation methods help, adaptive aggregation achieves competitive robustness without centralising CSI. This provides a stronger privacy–utility balance than post-hoc adaptation alone. Across altered-condition tests (e.g., TX shift, new occupant), FedOpt-Adam also achieves the strongest performance in Table I (79.31%), indicating improved robustness under moderate environmental variation.

FedOpt-Adam adds server-side optimizer state (first and second moments) but does not change the client model, communication payload, or on-device inference cost. All methods remain within memory and latency limits on Raspberry Pi,4-class devices. By smoothing and rescaling aggregated updates, the adaptive server rule reduces update variability and improves stability under heterogeneous, room-dependent CSI distributions, leading to stronger unseen-room robustness.

IV. Conclusion & Future Work

This poster reframes cross-room Wi-Fi sensing as a privacy–generalisation trade-off rather than an architectural modelling problem. By modifying only the server-side aggregation rule and preserving on-device CSI processing, FedOpt-Adam improves unseen-room accuracy over FedAvg by 7.49 points and approaches centralised performance without requiring CSI pooling. These findings underscore that aggregation design is a critical systems lever for mitigating non-i.i.d. CSI effects across rooms. Future work will evaluate partial participation, communication-efficient aggregation, and multi-room physical deployments to measure latency, energy consumption, and scheduling behaviour under real-world constraints.

Acknowledgements

Supported by the Australian NISDRG grant: NI240100147.

References

[1] R. Koo *et al.*, "A Survey on Green Wireless Sensing: Energy-Efficient Sensing via WiFi CSI and Lightweight Learning," *Energies*, vol. 19, no. 2, 2026. [Online]. Available: https://www.mdpi.com/1996-1073/19/2/573

[2] F. Meneghello *et al.*, "Sharp: Environment and person independent activity recognition with commodity ieee 802.11 access points," *IEEE Transactions on Mobile Computing*, vol. 22, no. 10, pp. 6160–6175, 2023.

[3] S. Reddi *et al.*, "Adaptive federated optimization," 2020. [Online]. Available: https://arxiv.org/abs/2003.00295

[4] B. McMahan *et al.*, "Communication-efficient learning of deep networks from decentralized data," in *Proc AISTATS*, 2017, pp. 1273–1282.

[5] T. Li *et al.*, "Federated optimization in heterogeneous networks," in *Proc. Machine Learning and Systems (MLSys)*, Mar. 2020, pp. 429–450.

[6] S. P. Karimireddy *et al.*, "Scaffold: stochastic controlled averaging for federated learning," in *Proc. ICML, PMLR 119:5132-5143*, 2020.