Session Calibration for User Identification Using CSI-Derived Wi-Fi Beamforming Feedback

Yuji Makimura, Daichi Hayashi, Naoto Sasaoka, and Masashi Nishiyama
Graduate School of Engineering, Tottori University, Tottori, Japan

***Abstract*—Wi-Fi sensing offers camera-free user identification, but a seat change alters radio paths and can make the same person appear different. Changes between empty-room recordings further obscure identity. We study this problem using beamforming feedback information (BFI), compressed information derived from channel state information (CSI). An empty-room recording captures BFI reports with nobody present; averaging them within each recording block gives an empty-room baseline. Our session calibration subtracts this baseline from both registered observations and new observations to be identified. Calibrated BFI reports retain positive and negative differences for user identification; occupancy detection uses absolute differences to reject empty-room observations. Experiments with different people for training and evaluation show that session calibration improves both tested recognition methods after a seat change and has the largest observed effect in their comparison. These results support session calibration for user identification, but substantial errors still preclude standalone authentication.**

***Index Terms*—User identification, beamforming feedback information, seat change, session calibration, occupancy detection.**

I. Introduction

Shared consumer devices need user identification before loading a personal profile or granting access. Wi-Fi sensing uses body-induced changes in radio propagation to identify people without cameras, wearables, or deliberate recognition actions [1]. Wi-Fi measurements can reveal human location, activity, and identity. *Channel state information* (CSI) supports user identification [2], [3]; its compressed derivative, *beamforming feedback information (BFI)* [4], supports localization [5], activity recognition [6], and user identification [7].

At a shared workstation or household computer, a returning user—a registered person using it again—may sit elsewhere. Enrollment collects BFI reports with a known identity; a query is a later observation to identify. In Fig. 1, this *seat change* leaves the person, room, and devices unchanged but places enrollment and query at different seats. Identification should follow the person, not require returning to the enrollment seat.

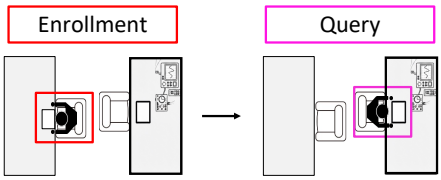

Fig. 1. User identification after a seat change. A known user is recorded during enrollment and later identified from a query at another seat. Only one person occupies the monitored area at a time.

Existing studies leave a basic requirement unresolved: can a registered person still be identified after a seat change? Because the body blocks and reflects different radio paths, high fixed-seat accuracy may reflect a person–seat combination rather than identity alone. Turetta *et al.* [8] use a one-dimensional convolutional neural network (1-D CNN) with fixed desk geometry. WiPT [9] uses a convolutional autoencoder (CAE) for authentication from posture changes; BFId [7] uses a long short-term memory (LSTM) network for walking-based BFI identification. None tests a seated person moving within an unchanged room. We quantify this loss and reduce it without query-user labels or retraining.

To improve user identification after a seat change, we propose *session calibration*: subtracting the corresponding mean of empty-room BFI reports from enrollment and query BFI reports. An *empty-room recording* captures BFI reports with nobody present; their mean within a block is the *empty-room baseline*. Each session has its own recording. The resulting *calibrated BFI reports* preserve positive and negative differences, while seat-dependent blocking and reflection remain. *Occupancy detection* rejects empty queries using absolute differences. We use *identity-disjoint evaluation*, with different people for training and testing, and vary *gallery size*, the number of registered people, with the model fixed. The contributions are:

- **Seat change:** We isolate the seat change while keeping the person, room, and devices fixed. Common-input comparisons expose severe losses across CNN, CAE, and LSTM models, not just one architecture, and demonstrate improved user identification with our method.
- **Session calibration:** A measured empty-room baseline is shared across seats without query-user labels or retraining. Comparing two recognition methods with and without session calibration isolates its benefit and identifies it as the largest observed effect in this comparison.
- **Occupancy detection:** The same empty-room recording supports user identification and rejection of empty queries. Comparing two CNNs with and without absolute differences tests the detector design. Matched-query evaluation then separates added detection errors from the remaining confusion between users after a seat change.

II. Problem Setting and BFI Acquisition

A. User identification after a seat change

BFI reports reflect both the person and the radio environment. Different people at a fixed seat can produce distinct patterns,

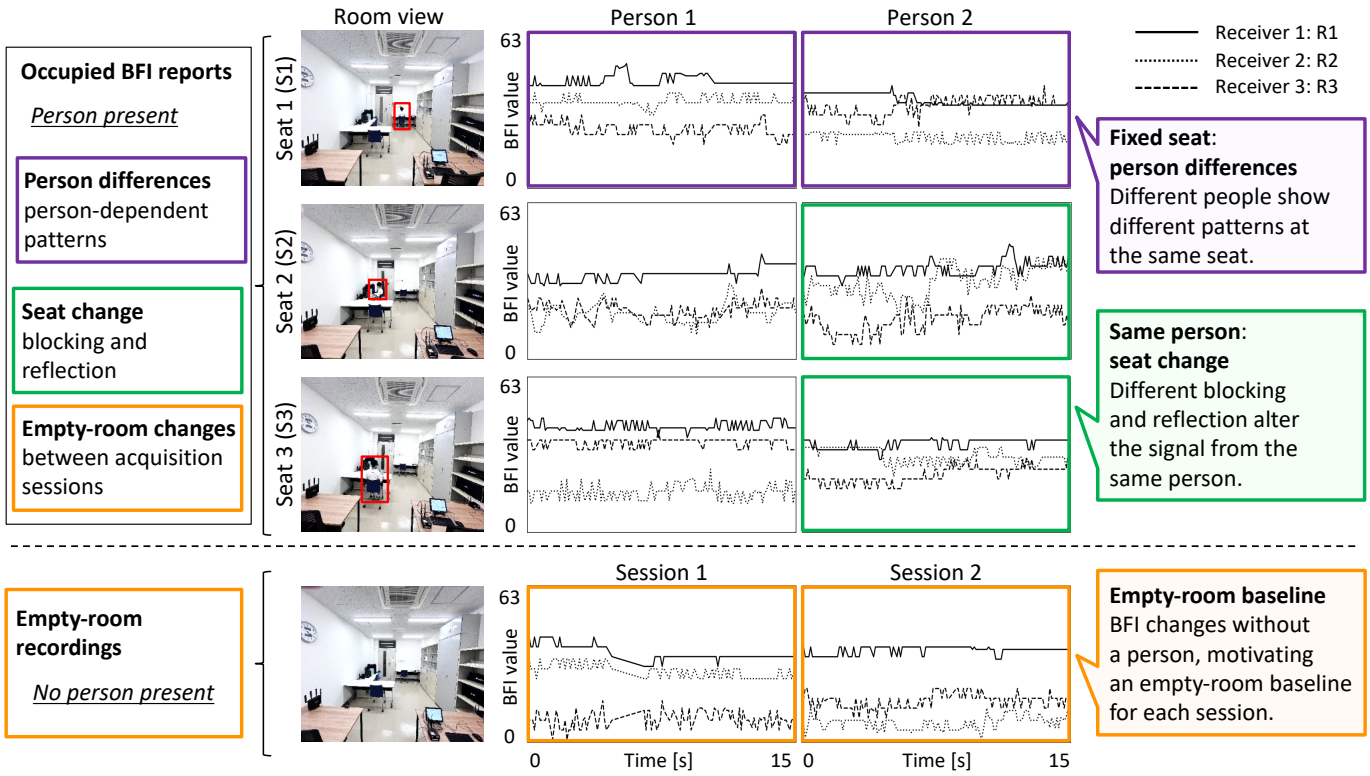

Fig. 2. Why a seat change makes user identification difficult. Occupied rows compare people horizontally and the same person across seats vertically; the bottom row compares empty-room recordings. All eight plots share the ranges 0–63 (BFI value) and 0–15 s (time). Curves show a selected BFI component from each receiver R1–R3 before normalization.

but moving the same person changes which radio paths the body blocks or reflects. An uncalibrated model may therefore recognize a person–seat combination. Empty-room BFI reports also change between acquisition sessions, even with fixed devices.

Figure 2 makes these difficulties visible. Each occupied row compares two people at one seat; each person column shows the same person across S1–S3. The bottom row shows changes with nobody present. Both effects can obscure person-dependent patterns. These examples motivate our method; they do not separately measure signal components or establish recognition accuracy.

An empty-room recording provides a baseline without the participant's contribution. In contrast, the mean of an occupied recording contains person-dependent information that subtraction could remove. Session calibration instead uses the empty-room baseline to reduce changes unrelated to the person while retaining differences caused by the occupied state.

B. Target scenario and acquisition system

This section describes the measurement layout and how BFI reports are obtained using Wireshark, software for capturing and inspecting network packets. The access point (AP), receivers, and capture node stay fixed, with one occupant at a time. Enrollment is confirmed occupied; queries undergo occupancy detection before user identification.

Figure 3 shows the approximately $3.9\,\mathrm{m}$ by $5.0\,\mathrm{m}$ room. A Buffalo WXR-5700AX7S AP is the beamformer; Jetson Nano systems are beamformees R1–R3. An Ubuntu 18.04 notebook with Wireshark 3.4.8 is the monitor-mode capture node C. S1–S3 mark the three candidate seats in Fig. 3. The task and devices stay fixed, so the seat is the intentionally varied acquisition condition. Communication uses 802.11ax [10] at $5\,\mathrm{GHz}$ over $20\,\mathrm{MHz}$. We passively capture its standard feedback, not 802.11bf sensing operations [11]. No AP firmware modification, custom waveform, camera, wearable, or specialized sensing radio is required.

Wireshark records approximately nine compressed beamforming reports per second per receiver, each providing 640 BFI values. The combined BFI reports (Section II-C) have rate $f_{\mathrm{BFI}} \approx 9$ reports/s. Each model input contains T consecutive BFI reports.

C. CSI-derived BFI representation

In this paper, a *BFI report* is one observation vector formed by collecting the BFI values from all three receivers in a fixed order. During channel sounding, each beamformee derives feedback matrices from CSI and reports compressed ϕ and ψ angles. We extract quantized integer codes from the *Feedback Matrices* field of high-efficiency (HE) compressed beamforming reports. Full complex CSI is not reconstructed; the system does not actively scan transmit beams.

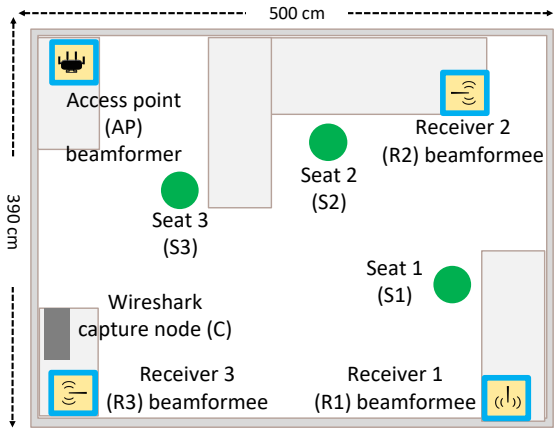

Fig. 3. Measurement layout for the seat-change study.

The 640 BFI values per receiver (Section II-B) comprise five ϕ and five ψ codes at each of 64 frequency indices for a 4×2 feedback matrix. Combining R1–R3 gives one BFI report with $D = 3 \times 640 = 1920$ dimensions. Each dimension is Min–Max scaled using only the people available at the training stage. Session calibration operates in this normalized space, putting the 6-bit ϕ and 4-bit ψ codes on comparable numerical scales. No circular unwrapping or physical-angle reconstruction is applied.

D. Acquisition protocol

This section describes the recording order and pairing of occupied and empty-room recordings. Twenty-six male participants provided written informed consent and completed three acquisition sessions, each a separate recording period. In each session, a 30 s empty-room recording was followed by a 30 s recording of BFI reports at each of S1–S3 while the participant used a keyboard and mouse. The seat order was randomized, and the devices were not moved. Each recording was divided into two non-overlapping 15 s blocks. Each occupied block is paired with the empty-room block at the corresponding position within the same session's recordings. All seats share this pairing of separately acquired blocks.

III. Session Calibration for User Identification

A. Overview

Figure 4 shows the overview of our method. Step 1 trains the occupancy detector offline, from previously collected BFI reports before enrollment and query. Step 2 acquires enrollment BFI reports from a known, present user. Step 3 rejects query windows predicted empty. Step 4 calibrates enrollment and accepted query BFI reports. Step 5 extracts identification features and performs centroid matching. The occupancy detector and identification encoder use development identities for training and validation, then remain fixed for enrollment and query BFI reports from separate evaluation identities.

Let r uniquely index acquisition sessions, s the occupied seat, and b blocks paired as described in Section II-D. The baselines $\boldsymbol{\mu}_{r,b}^{\mathrm{det}}$ and $\boldsymbol{\mu}_{r,b}^{\mathrm{cal}}$ serve occupancy detection and session calibration, respectively. Both are means of BFI reports recorded separately with nobody present. Query occupancy is inferred from the query BFI reports.

$\mathbf{x}_{r,s,b,i} \in \mathbb{R}^D$ is a scaled occupied BFI report at seat s. $\mathbf{c}_{r,b,n} \in \mathbb{R}^D$ is a scaled BFI report from the paired block of a separate empty-room recording made with nobody present in session r. $\mathbf{q}_{r,b,i} \in \mathbb{R}^D$ is a scaled query BFI report with unknown occupancy. Indices i, n count BFI reports in capture-time order. Each input contains T consecutive scaled BFI reports within one block, indexed by $\mathcal{W} = \{i_0, \ldots, i_0 + T - 1\}$, with starting report i_0. Its nominal duration $\tau = T/f_{\mathrm{BFI}}$ seconds uses the average report rate, not exact capture timestamps.

B. Occupancy detection

Step 1 trains the detector; Step 3 applies it. Each empty-room block supplies disjoint reports for baseline estimation and empty detector examples. Let $\mathcal{I}_{r,b}^{\mathrm{det}}$ index the baseline-estimation reports; $|\mathcal{I}_{r,b}^{\mathrm{det}}|$ counts them. Superscript det denotes occupancy detection. The empty-room baseline is $\boldsymbol{\mu}_{r,b}^{\mathrm{det}} = |\mathcal{I}_{r,b}^{\mathrm{det}}|^{-1} \sum_{n \in \mathcal{I}_{r,b}^{\mathrm{det}}} \mathbf{c}_{r,b,n}$. For a scaled query BFI report $\mathbf{q}_{r,b,i}$, the detector input is the element-wise absolute difference

$$\mathbf{d}_{r,b,i} = \left| \mathbf{q}_{r,b,i} - \boldsymbol{\mu}_{r,b}^{\mathrm{det}} \right|. \tag{1}$$

A lightweight multiscale Conv1D network processes the vectors $\mathbf{d}_{r,b,i} \in \mathbb{R}^D$ for $i \in \mathcal{W}$. Training uses these absolute differences with known occupancy labels. Empty examples use the disjoint segment of the same block, a later block in that session, or another session to limit reliance on temporal adjacency. Separate validation identities select the threshold. In Step 3, $\boldsymbol{\mu}_{r,b}^{\mathrm{det}}$ comes from the query session, not the training session. The fixed detector rejects predicted-empty windows; accepted windows pass their original scaled BFI reports $\mathbf{q}_{r,b,i}$, not the absolute differences $\mathbf{d}_{r,b,i}$, to session calibration in Step 4.

C. Session calibration

In Step 4 of Fig. 4, session calibration uses the empty-room baseline obtained by averaging all scaled BFI reports $\mathbf{c}_{r,b,n}$ from session r, block b. Superscript cal denotes session calibration:

$$\boldsymbol{\mu}_{r,b}^{\mathrm{cal}} = \frac{1}{L_{r,b}} \sum_{n=1}^{L_{r,b}} \mathbf{c}_{r,b,n}, \tag{2}$$

where $L_{r,b}$ counts all BFI reports in that empty-room recording block. During enrollment, session calibration subtracts this baseline from $\mathbf{x}_{r,s,b,i}$, the i-th scaled BFI report acquired with the known user at seat s:

$$\widetilde{\mathbf{x}}_{r,s,b,i} = \mathbf{x}_{r,s,b,i} - \boldsymbol{\mu}_{r,b}^{\mathrm{cal}}, \qquad i \in \mathcal{W}. \tag{3}$$

The result $\widetilde{\mathbf{x}}_{r,s,b,i}$ is a calibrated enrollment BFI report. For an accepted query, $\widetilde{\mathbf{q}}_{r,b,i} = \mathbf{q}_{r,b,i} - \boldsymbol{\mu}_{r,b}^{\mathrm{cal}}$ gives its calibrated BFI report. Both paths retain positive and negative differences and use their own session's baseline, shared across seats within a block.

Seat-dependent blocking, reflection, and baseline-estimation errors can remain; subtraction does not isolate person and room components. It nevertheless enables feature comparison

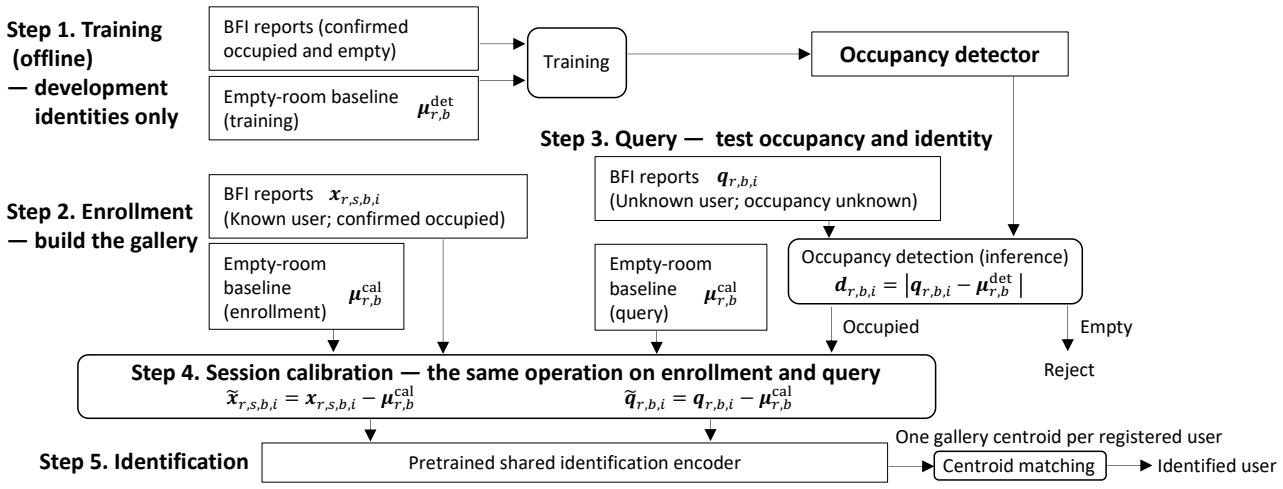

Fig. 4. Overview of our method. Step 1 trains the occupancy detector offline using development identities with known occupancy labels. Step 2 acquires enrollment BFI reports from a known, present user. Step 3 rejects query windows predicted empty using the trained detector. Step 4 subtracts each session's empty-room baseline $\mu^{\mathrm{cal}}_{r,b}$ from enrollment and accepted query BFI reports, preserving positive and negative differences. In Step 5, the pretrained shared identification encoder extracts features; enrollment features are averaged and normalized into one gallery centroid per user, and query features are matched to these centroids by cosine distance.

relative to each session's baseline without query-user labels or retraining.

D. User identification

In Step 5 of Fig. 4, the *pretrained shared identification encoder* maps calibrated BFI report windows to embeddings (feature vectors for enrollment and query) using the same fixed, offline-trained weights. Its architecture comprises Conv1D branches with kernels 3 and 7, a projection layer, max pooling, and two Transformer blocks [12]. Global average pooling and a dense layer produce a unit-length, 128-dimensional embedding. P-K batches contain $P = 8$ people and $K = 2$ samples per seat per person, totaling 48 samples. CosFace [13] uses normalized class weights, scale $\gamma = 8$, and margin $m = 0.10$; we call this method P-K CosFace.

Enrollment embeddings are averaged and renormalized into one gallery centroid per registered person. Each accepted query embedding is matched to the nearest centroid by cosine distance, one minus the dot product of unit-length vectors.

IV. Experimental Evaluation and Discussion

A. Experimental settings and identity-disjoint evaluation

Correct match rate (CMR) is closed-set rank-1 accuracy: the fraction of queries assigned to the correct person. Equal error rate (EER) is where false-acceptance and false-rejection rates coincide. Distances to the true/other centroids are genuine/impostor scores. We sweep one threshold over pooled scores and average the two error rates at their closest point. A rejected occupied query counts as a CMR error and a false rejection, with no impostor scores.

Each of five trials splits 26 people into 13 development identities and 13 evaluation identities. Nine development identities train detectors and encoders; four select epochs and detector thresholds with architectures fixed. P-K CosFace uses learning rate 5×10^{-5} and up to 60 epochs, selected by highest validation seat-change CMR, then lowest EER. All detectors and encoders are retrained on the 13 development identities. In Fig. 4, Step 1 trains the detector; Steps 2–5 use BFI reports from the 13 evaluation identities. The pretrained shared identification encoder in Step 5 is trained exclusively on the 13 development identities.

Each seat serves as the enrollment seat. Ten enrollment windows per person come from one of six blocks; ten queries per person and query seat come from each remaining block. Fixed-seat queries use the enrollment seat; seat-change queries use the other two. Query windows may overlap, but exclude the enrollment block. We average over six enrollment-block choices, then report mean $\pm$ sample standard deviation (SD) over five trials and three enrollment seats (15 trial–seat units). SDs are descriptive because units share people and models; means and SDs are rounded only for display.

All comparisons use nominal duration $\tau = 2\,\mathrm{s}$, among the windows tested by Turetta *et al.* [8]. This equals the monitoring interval in WiPT [9], not its $2.5\,\mathrm{s}$ authentication input. At the BFI report rate $f_{\mathrm{BFI}} \approx 9$ reports/s given in Section II-B, we use $T = 18$ consecutive reports, from $2 \times 9 = 18$. Each report combines 640 BFI values from each of three receivers. The occupancy baseline uses the leading half of each empty-room block; the other half supplies disjoint empty examples. This report partition is separate from the participant split.

B. Comparison with existing methods

Table I compares our method with the seated-user 1-D CNN of Turetta *et al.* [8], WiPT's 2-D CAE [9], and the BFId-style LSTM [7]. All use the same BFI reports, identity-disjoint evaluation, and centroid matching, with settings fixed on development identities. Source tasks and signals differ, so these are our comparison results, not published source scores. Our method combines session calibration with P-K CosFace; this comparison uses occupied queries.

TABLE I
EXISTING METHODS AND OUR METHOD UNDER COMMON BFI REPORTS AND EVALUATION SETTINGS (%; MEAN ± SD).

Method	Seat change		Fixed seat	
	CMR ↑	EER ↓	CMR ↑	EER ↓
1-D CNN [8]	17.1±3.0	39.9±2.4	58.8±6.9	20.0±3.8
2-D CAE [9]	12.5±2.7	43.6±2.9	72.7±6.5	14.9±3.6
BFId-style LSTM [7]	19.6±3.6	36.7±2.4	67.5±5.0	16.0±3.5
Our method	**39.0±5.8**	**27.3±2.6**	**79.9±3.4**	**11.4±1.9**

TABLE II
EFFECT OF SESSION CALIBRATION AND RECOGNITION METHOD (%; MEAN ± SD OVER 15 TRIAL–SEAT UNITS). ALL USE CENTROID MATCHING. OUR METHOD IN TABLE I IS O4; THE CAE RESULTS ARE FROM SEPARATE RUNS.

ID	Session calibration	Recognition	Seat change		Fixed seat	
			CMR ↑	EER ↓	CMR ↑	EER ↓
O1	No	2-D CAE	14.3±3.2	42.5±2.1	75.9±4.8	13.9±2.9
O2	Yes	2-D CAE	34.7±6.3	28.9±2.1	71.8±4.4	14.1±2.2
O3	No	P-K CosFace	14.7±3.4	41.8±2.0	68.2±8.2	16.0±3.9
O4	Yes	P-K CosFace	**39.0±5.8**	**27.3±2.6**	**79.9±3.4**	**11.4±1.9**

Our method has the highest mean CMR and lowest mean EER in both seat conditions. After a seat change, its 39.0% CMR and 27.3% EER compare with 12.5–19.6% and 36.7–43.6% for existing models, whose CMR drops by 41.7–60.2 points and EER rises by 19.9–28.7 points. Thus, both correct-person selection and genuine–impostor separation improve for unseen people without query-user labels or retraining, although seat effects remain.

C. Effect of session calibration

Table II isolates session calibration by comparing each recognition method with and without it. The 2-D CAE has the highest fixed-seat CMR and lowest fixed-seat EER among existing models in Table I; comparing it with P-K CosFace tests whether session calibration also helps an existing model. All four conditions use centroid matching:

- **O1: 2-D CAE without session calibration.** Measures the failure after a seat change.
- **O2: 2-D CAE with session calibration.** Comparison with O1 isolates session calibration.
- **O3: P-K CosFace without session calibration.** Comparison with O1 tests changing the recognition method alone.
- **O4: P-K CosFace with session calibration.** Comparison with O3 tests session calibration; comparison with O2 tests recognition after session calibration.

From the displayed means in Table II, session calibration raises seat-change CMR by 20.4 and 24.3 points and lowers EER by 13.6 and 14.5 points for 2-D CAE and P-K CosFace, respectively. Changing the recognition method raises CMR by only 0.4 points without session calibration and 4.3 points with it. Session calibration therefore has the largest observed effect in this comparison. At a fixed seat, O2 has lower mean CMR than O1.

TABLE III
OCCUPANCY-DETECTION ABLATION (%; MEAN ± SD OVER FIVE TRIALS). ABSOLUTE INDICATES WHETHER EQ. (1) TAKES THE ABSOLUTE VALUE. D4 IS THE CURRENT DETECTOR CONFIGURATION.

ID	Absolute	CNN	Bal. acc. ↑	Recall ↑	Specificity ↑
D1	No	Single-scale	96.4±0.7	94.1±1.8	**98.7±1.1**
D2	Yes	Single-scale	98.5±2.1	99.7±0.6	97.3±4.3
D3	No	Multiscale	96.1±0.9	94.4±1.5	97.9±0.9
D4	Yes	Multiscale	**99.2±0.8**	**99.9±0.1**	98.5±1.7

D. Occupancy detection and user identification

Table III tests whether absolute differences and multiple convolutional scales improve occupancy detection. No omits only the absolute value in Eq. (1). The single-scale CNN uses kernel 5 with 64 channels; the multiscale CNN uses kernels 3 and 7 with 32 channels each. Remaining layers and parameter counts are identical. All four conditions are retrained with common splits, windows, baselines, training limits, and validation-selection rules. Each trial pools the three empty-query conditions in Section III-B; balanced accuracy averages occupied recall and empty specificity.

Absolute differences improve mean balanced accuracy by 2.1/3.1 points and occupied recall by 5.6/5.5 points for single-scale/multiscale CNNs. D4 has the highest mean balanced accuracy and recall; D1 has slightly higher specificity. Multiple scales add 0.7 points in balanced accuracy with absolute differences, but lose 0.3 points without them. Absolute differences provide the clearer benefit.

Table IV evaluates the D4 detector configuration in a separate O4 run against known occupancy (accepting every occupied query), using identical query windows. Nearly unchanged performance in both seat conditions shows that empty-query rejection need not materially reduce user identification.

In the full 14-class task (13 users plus empty), seat-change empty-state recall is 98.5 ± 1.6%, but macro recall, which

TABLE IV
USER IDENTIFICATION WITH KNOWN OCCUPANCY OR OCCUPANCY DETECTION IN A SEPARATE RUN (%).

Evaluation	Change CMR ↑	Change EER ↓	Fixed CMR ↑	Fixed EER ↓
Known occupancy	38.6 ± 7.6	27.0 ± 3.4	80.7 ± 5.7	10.8 ± 2.5
Occupancy detection	38.5 ± 7.6	27.0 ± 3.4	80.6 ± 5.8	10.8 ± 2.5

TABLE V
SEAT-CHANGE CMR VERSUS GALLERY SIZE IN A SEPARATE O4 RERUN (%; MEAN ± SD OVER 15 TRIAL–SEAT UNITS).

G	Chance CMR	CMR ↑
4	25.0	63.1 ± 4.5
7	14.3	50.7 ± 5.3
10	10.0	43.7 ± 5.6
13	7.7	39.0 ± 5.8

weights classes equally, is $42.8 \pm 7.1\%$. The remaining limitation is confusion between users, not empty-query rejection.

E. Effect of gallery size

To quantify how the number of registered people affects identification after a seat change, we vary gallery size G with the learned model fixed. We test $G \in \{4, 7, 10, 13\}$ in a separate O4 run, since the original weights were not retained. Original identity splits, selected epochs, seeds, preprocessing, architecture, and hyperparameters are retained. Each trial–seat–block setting reuses its encoder, enrollment/query windows, centroids, and distances across G. Each of the $\binom{13}{G}$ subsets uses only its members' queries and centroids, without retraining or model selection. We average over subsets, then six enrollment blocks, and report mean ± SD over 15 trial–seat units, not independent subset replicates.

Table V shows CMR rising from 39.0% at $G = 13$ to 63.1% at $G = 4$. The $G = 13$ row rounds to the same mean and SD as O4 in Table II but uses different predictions. Chance CMR, $(100/G)\%$, also rises as gallery size decreases. This quantifies reduced identification difficulty, not an additional effect of session calibration or improved tolerance to a seat change.

F. Practical use and limitations

One $30\,\mathrm{s}$ empty-room recording serves all seats in a session. Startup or detected sustained absence could support acquisition, but automatic updating, exclusion of occupied samples, and baseline validity are untested. Device or furniture changes may require a new empty-room recording.

The $2\,\mathrm{s}$ window is not measured processing latency, and the prototype needs a separate capture node. Evaluation covers one room, three receivers, one person at a time, short-term sessions, and an all-male cohort. Unknown-user rejection, spoofing, multiple occupants, and long-term or multiroom transfer remain untested. Even the smaller gallery size leaves substantial errors: the method is a supporting cue for user identification, not standalone authentication.

V. CONCLUSION

Session calibration improves user identification after a seat change by subtracting an empty-room baseline from enrollment and query BFI reports. It has the largest observed effect in our recognition-method comparison. The occupancy ablation supports absolute differences, while occupancy detection adds negligible identification error under the tested conditions. Unlike processing complex CSI [14] or learning an adaptation model [15], our method corrects BFI reports without CSI reconstruction, query-user labels, or query-time retraining. Gallery-size evaluation reveals remaining errors. Future work will examine baseline validity over time and when a new empty-room recording is needed. The authors thank Dr. Michiko Inoue, Dr. Shunsuke Yoneda, Mr. Shouki Sakemoto, and Mr. Teppei Kobiki for their discussions and assistance.

REFERENCES

[1] Z. Wei, W. Chen, S. Ning, W. Lin, N. Li, B. Lian, X. Sun, and J. Zhao, "A survey on WiFi-based human identification: Scenarios, challenges, and current solutions," *ACM Trans. Sensor Netw.*, vol. 21, no. 1, pp. 10:1–10:32, Jan. 2025.

[2] F. Wang, J. Han, F. Lin, and K. Ren, "WiPIN: Operation-free passive person identification using Wi-Fi signals," in *Proc. IEEE GLOBECOM*, 2019, pp. 1–6.

[3] M. Han, L. Guo, J. Zhang, H. Ji, Z. Diao, and J. Sun, "WiID: Precise WiFi-based person identification via bio-electromagnetic information," in *Proc. ICPR*, 2022, pp. 1105–1112.

[4] K. F. Haque, F. Meneghello, and F. Restuccia, "Wi-BFI: Extracting the IEEE 802.11 beamforming feedback information from commercial Wi-Fi devices," in *Proc. 17th ACM Workshop Wireless Netw. Testbeds, Exp. Eval. Characterization (WiNTECH)*, 2023, pp. 104–111.

[5] S. Kondo, S. Itahara, K. Yamashita, K. Yamamoto, Y. Koda, T. Nishio, and A. Taya, "Bi-directional beamforming feedback-based firmware-agnostic WiFi sensing: An empirical study," *IEEE Access*, vol. 10, pp. 36 924–36 934, 2022.

[6] K. F. Haque, M. Zhang, F. Meneghello, and F. Restuccia, "BeamSense: Rethinking wireless sensing with MU-MIMO Wi-Fi beamforming feedback," *Comput. Netw.*, vol. 258, p. 111020, 2025.

[7] J. Todt, F. Morsbach, and T. Strufe, "BFId: Identity inference attacks utilizing beamforming feedback information," in *Proc. ACM SIGSAC Conf. Comput. Commun. Security (CCS)*, 2025, pp. 2399–2413.

[8] C. Turetta, F. Demrozi, P. H. Kindt, A. Masrur, and G. Pravadelli, "Practical identity recognition using WiFi's channel state information," in *Proc. DATE*, 2022, pp. 76–79.

[9] S. Liu, Y. Chen, H. Wang, H. Liang, and L. Chen, "A low-calculation contactless continuous authentication based on postural transition," *IEEE Trans. Inf. Forensics Security*, vol. 17, pp. 3077–3090, 2022.

[10] *IEEE Std 802.11ax-2021: Wireless LAN Medium Access Control and Physical Layer Specifications—Amendment 1: Enhancements for High-Efficiency WLAN*, IEEE, 2021.

[11] *IEEE Std 802.11bf-2025: Wireless LAN Medium Access Control and Physical Layer Specifications—Amendment 4: Enhancements for Wireless LAN Sensing*, IEEE, 2025.

[12] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Adv. Neural Inf. Process. Syst.*, vol. 30, 2017, pp. 5998–6008.

[13] H. Wang, Y. Wang, Z. Zhou, X. Ji, D. Gong, J. Zhou, Z. Li, and W. Liu, "CosFace: Large margin cosine loss for deep face recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2018, pp. 5265–5274.

[14] X. Wang, L. Gao, and S. Mao, "CSI phase fingerprinting for indoor localization with a deep learning approach," *IEEE Internet Things J.*, vol. 3, no. 6, pp. 1113–1123, Dec. 2016.

[15] J. Zhang, Z. Chen, C. Luo, B. Wei, S. S. Kanhere, and J. Li, "MetaGanFi: Cross-domain unseen individual identification using WiFi signals," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 6, no. 3, pp. 152:1–152:21, 2022.