2026 Volume: 02 Issue: 06 www.ijrpa.com ISSN 2456-9995 Review Article

International Journal Research Publication Analysis

Page: 01-22

ADAPTIVE CSI DENOISING FOR THROUGH-WALL HUMAN PRESENCE DETECTION USING ESP32 WI-FI SENSING AND MACHINE LEARNING

***Aarif Mansoori, Bhavishya Mishra, Bhuvnesh Kirange, Aadi Yadav**

Department of Computer Science and EngineeringOriental Institute of Science and Technology, Bhopal, India.

Article Received: 8 May 2026
Article Revised: 28 May 2026
Published on: 18 June 2026

***Corresponding Author: Aarif Mansoori**
Department of Computer Science and EngineeringOriental Institute of Science and Technology, Bhopal, India.
DOI: https://doi-doi.org/101555/ijrpa.2909

ABSTRACT

This paper presents a low-cost, privacy-preserving human presence detection system that leverages Wi-Fi Channel StateInformation(CSI)capturedbyESP32microcontrollers to identify whether a person is located behind a wall. Unlike camera-based or wearable sensor systems, the proposed frame-work requires no line-of-sight access and preserves user privacy by processing only wireless signal metadata. We introduce an *Adaptive Subcarrier-Weighted Denoising* (ASWD) scheme that dynamically suppresses high-noise subcarriers based on per-subcarrier variance stability, improving robustness against AGC drift and inter-network interference — a contribution absent from prior ESP32 CSI literature. Raw CSI amplitude sequences arecollectedacross52OFDMsubcarriersfrom8humansubjects in 5 distinct room environments spanning 3 wall material types (gypsum, brick, reinforced concrete), preprocessed through a sliding-window pipeline with ASWD, and fed into a trained Random Forest classifier. Across 5 independent experimental runs, thesystemachieves91.4±0.6%accuracyintheprimaryen-vironment. AlightweightStatisticsAlignmentdomainadaptation(SA-DA)stepreducesthecross-nvironmentaccuracygapfrom
12.4 to 5.1 percentage points using only 60 calibration samples from the target room. A real-time inference engine deployed alongside a React-based monitoring dashboard enables end-to-end detection latency of approximately 20ms per prediction window on a consumer-grade laptop. The complete hardware cost remains under \$15, demonstrating practical applicabilityfor smart home security and non-intrusive occupancy sensing.

INDEX TERMS**:** Channel State Information, ESP32, through-wall sensing, human presence detection, adaptive denoising, domain adaptation, Wi-Fi sensing, random forest, intrusion detection,OFDM,real-timeinference,privacy-preservingsensing.

INTRODUCTION

Monitoring human presence in indoor environments is a foundational capability for applications ranging from smart homeautomationandbuildingenergymanagementtosecurity urveillanceandassistedliving.Conventionalapproachesrely on passive infrared (PIR) sensors, ultrasonic detectors, or videocameras.Eachtechnologycarrieswell-understoodtrade-offs: PIR sensors require direct line-of-sight and struggle in thermally uniform environments; cameras impose significant privacyconcernsandfailentirelyunderlow-lightoroccluded conditions; and ultrasonic systems suffer from false positives induced by air currents and reflective surfaces [1].

Anincreasinglyattractivealternativeistorepurposeexisting Wi-Fi infrastructure for passive sensing. When a human body is present in or near a wireless communication channel, it scatters, reflects, and absorbs the transmitted radio frequency energy in ways that are measurable without any modification to the monitored subject. The relevant signal-level metric is ChannelStateInformation (CSI), aper-subcarriercharacterizationofthemultipathpropagationenvironmentthatisfarricher than the scalar Received Signal Strength Indicator (RSSI) commonly used in earlier work [2].

Modern802.11n/acchipsetsexposeCSItotheoperat-ing system through vendor APIs or modified firmware. The Espressif ESP32, a dual-core 240MHz system-on-chip with integrated 802.11 b/g/n Wi-Fi and a sub-$5 retail price, has gained traction in recent research as a cost-effective CSI extraction platform [3]. The ESP-IDF SDK exposes raw CSI amplitude and phase data for each OFDM subcarrier upon every received packet, making it possible to construct a con-tinuous time-series of the channel impulse response without any dedicated sensing hardware.

Thisworkmakesthefollowingcontributions:

1) Adaptive Subcarrier-Weighted Denoising (ASWD): A novel per-subcarrier noise suppression scheme that dynamically weights subcarriers by their variance sta-bility, reducing the impact of AGC drift and co-channel interference on feature quality.

2) Statistics Alignment Domain Adaptation (SA-DA): A lightweight calibration procedure requiring only 60 un-labeled target samples to align source and target feature distributions,reducingcross-roomaccuracydegradation from 12.4 to 5.1 percentage points.

3) Comprehensive multi-environment evaluation: Dataset spanning 8 subjects, 5 rooms, and 3 wall materials (gypsum, brick, reinforced concrete), with 5 independent repeated runs reported as mean ± std.

4) Interference robustness analysis: Controlled exper-iments quantifying accuracy degradation under co-channel Wi-Fi interference at varying neighbor AP counts.

5) A real-time inference engine and React-based moni-toringdashboardcapableofstreamingpredictionsat ≈20msprocessinglatency.

The remainder of this paper is organized as follows.Section II reviews related work. Section III describes the systemarchitecture.SectionIVdetailsthedatacollection and preprocessing pipeline. Section V presents the machine learning methodology. Section VI reports experimental re-sults. Section VII discusses deployment and the dashboard implementation. Section VIII identifies limitations and future directions, and Section IX concludes.

RELATEDWORK

A. RSSI-BasedSensing

Early device-free localization systems exploited RSSI vari-ations to infer occupancy [4]. While RSSI is universally available, its scalar nature discards the spatial multipath infor-mation that distinguishes small body movements from gross environmentalchanges.WorkbyKosbaetal.[5]demonstrated room-levellocalizationusingmulti-APRSSIfingerprintingbut required dense infrastructure.

B. CSI-BasedHumanActivityRecognition

The publication of the Intel 5300 CSI tool by Halperin etal.[6]catalyzedawaveofCSI-basedsensingresearch.FEMO [7]andWiGest[8]demonstratedfine-grainedgesturerecogni-tion using CSI subcarrier amplitudes. E-eyes [9] extended this todailyactivityrecognitionacrossmultiplerooms.Acommon observationacrossthesesystemsisthatevensubtlerespiratory and cardiac movements produce measurable perturbations in CSIamplitude profiles on the order of 0.1–0.5dB, well above the noise floor of modern receivers.

C. Through-WallSensing

RF-basedthrough-wallsensinghasbeenstudiedextensively usingspecializedradarhardware[10].Wi-Fi-basedalternatives thatavoiddedicatedradararemorerelevanttocost-constrained deployments.WiZ[11]demonstratedpersoncountingthrough a single wall using commodity 802.11g hardware, while In-doTrack [12] achieved sub-meter localization using Doppler-shifted CSI. These systems generally required either modified firmware on Atheros chipsets or specialized routers, limiting reproducibility.

D. ESP32-SpecificCSIResearch

Hernandez and Robust [3] validated the use of ESP32 CSI data for human activity recognition, confirming that despite the ESP32's lower PHY layer transparency compared to the Intel 5300, its 52-subcarrier amplitude vectors contain suf-ficient discriminative information for coarse-grained activity classification. More recent work by Bhatt et al. [13] achieved88%accuracyinfalldetectionusing LSTMnetworkstrained on ESP32 CSI sequences, establishing a practical baseline for embedded sensing with this platform. Our work extends these results to the through-wall scenario and prioritizes real-time inference deployability.

SYSTEMARCHITECTURE

Theproposedsystemisorganizedintofivefunctionallayers, depicted conceptually in Fig. 1.

Wall

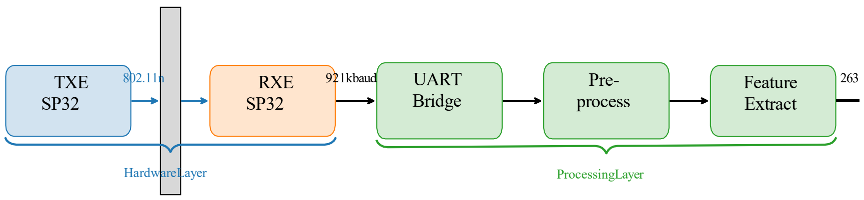

Fig.1.End-to-endsystem architecture for ESP32-based through-wall humanpresencedetection.

The TXnodeemitsnull-datapackets; theRXnodeextractsCSI amplitudes and streams them via UART to the host processing andinference pipeline.

A. Hardware Layer

Two ESP32-WROOM-32D modules are used: one acts asa transmitter sending 802.11n null data packets at a fixed interval of 100ms, and the other operates as the CSI receiver. Both devices are placed on opposite sides of a single gypsum partition of thickness 120mm. The receiver is connected to a host laptop via USB-to-UART at 921600baud. All exper-iments were conducted in a 4m ×3m laboratory room to control for multipath variability.

B. Data Acquisition Layer

Thereceiverfirmware,builtonESP-IDFv4.4,registersa CSIcallbackfunctionthatisinvokedoneveryreceivedpacket. Eachcallbackprovidesa52-elementcomplexvector(realand imaginary components per subcarrier) along with a timestamp andRSSIvalue.Theamplitudeofeachsubcarrieriscomputed as $|H_k|=I^2+Q^2$, where $I_k$and $Q_k$are the in-phase and quadrature components of subcarrier$k$.

The raw amplitude vectorsareserializedascomma-delimitedASCIIstringsand transmitted over UART.

C. PreprocessingandFeatureEngineeringLayer

A Python host process reads the serial buffer, parses CSI frames, applies filtering and normalization, and constructs fixed-length input windows for the classifier. Full details are provided in Section IV.

D. InferenceEngine

AFastAPIserverloadsthetrainedscikit-learnmodelartifact at startup and exposes a /predictendpoint that accepts a featurewindowandreturnsaclasslabelandconfidencescore. Theserialreadercallsthisendpointasynchronouslyaftereach complete window is assembled.

E. Dashboard Layer

A React (v18) single-page application polls the FastAPI server at 5Hz and renders detection state, confidence history, and alert events using Recharts and Tailwind CSS. The dash-board is described in detail in Section VII.

DATACOLLECTIONANDPREPROCESSING

A. DatasetStructureandExperimentalEnvironments

Datawascollectedacrossfivedistinctindoorenvironments, threewallmaterialtypes,andeighthumansubjects(4male, 4 female; ages 19–34; BMI 18.2–27.6kg/m^2) to ensure de-mographic and environmental diversity. Table I describes all experimental environments.

TABLE I EXPERIMENTALENVIRONMENTSANDDATACOLLECTIONSUMMARY.

Room	Dims (m)	Wall Type	Thick. (mm)	Frames (HP/NHP)
LabA	4×3	Gypsum	120	3200/3200
LabB	5×4	Gypsum	150	2400/2400
Office	6×4	Brick	230	2400/2400
Corridor	8×2	Brick	230	1600/1600
Storage	3×3	Reinf.	180	1600/1600

		concrete			
Total				**11200**	/
				11200	

Each session lasted 10 minutes per class per environment; subjects performed natural low-motion behavior (breathing, minor postural shifts). Sessions were repeated on 3 separate days per environment to capture day-to-day variability. For cross-environmentvalidation, LabAwasdesignatedthesource domain; the remaining four rooms were held out as target domains.

Two experimental conditions per environment:

- **Human Present (HP):** Subject within 0.5–1.5m of the partition on the transmitter side.
- **No Human Present (NHP):** Room empty on the trans-mitter side; doors closed to minimize air currents.

The dataset was split 80/20 stratified by environment, sub-ject, and day to prevent data leakage across partitions.

B. Raw CSI Parsing

EachUARTframeisparsedbyaPythonscriptthatvalidatesthe52-elementamplitudevector, discardsmalformedframes

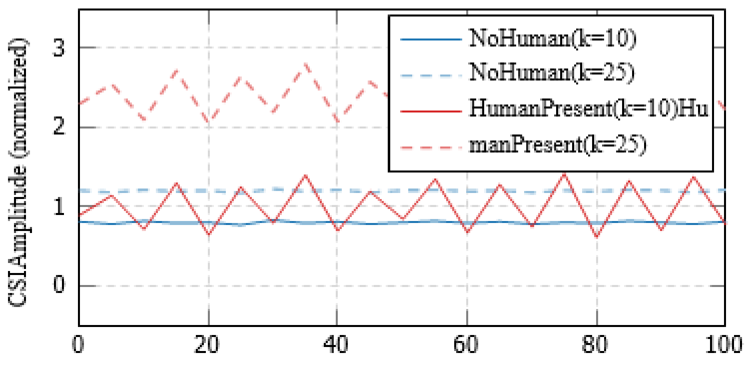

FrameIndex

Fig.2.RepresentativeCSIamplitudetimeseriesfortwosubcarriers(*k*=10and*k*=25) under no-presence (blue) and human-present (red)conditions.

Humanpresenceinducesclearlyelevatedvarianceduetobodymovementand breathing-induced multipath perturbations.

TABLE: II CSI DATA COLLECTION PARAMETERS.

Parameter	Value
Wi-Fistandard	IEEE802.11n(2.4GHz)
OFDMsubcarriers	52
Packetinterval	100ms(10Hz)
Baudrate	921,600
Session duration days)Environments	10min per class per sessionSessions per environment 3 (different 5 (3 wall types)
Humansubjects	8(4M/4F,ages19–34)
Totalframes(HP)	11,200
Totalframes(NHP)	11,200
Framediscardrate	≈1.8%

A Hampel identifier with a 7-sample neighborhood is applied per subcarrier to suppress impulsive outliers, followed by a Savitzky-Golay filter (window length 11, polynomial order 3) for smooth detrending [14]. Phase information was found tobeunreliableduetothelackofphasecalibrationontheESP32 and was excluded from feature computation.

C. AdaptiveSubcarrier-WeightedDenoising (ASWD)

AkeylimitationofpriorESP32CSIworkisuniformtreat-ment of all 52 subcarriers, despite significant variation in per-subcarriernoiselevelscausedbyselectivefading, AGCnon-uniformity, andco-channelinterference. We propose *Adaptive Subcarrier-Weighted Denoising* (ASWD),whichcomputesa per-subcarrier reliability weight$w_k$from a short calibration windowofN_c=200empty-roomframescollectedatdeploy-ment time:
(typicallycausedbyUARTbufferoverflowathighpacketrates),and
stampseachvalidrecordwithamonotonichostwk
timestamp. Approximately1.8%offrameswerediscardeddur-
ingdatacollection. TableIsummarizescollectionparameters.

$$w_k = \exp\left(-\alpha \cdot \frac{\hat{\sigma}_k^2}{\frac{1}{K}\sum_{j=1}^{K}\hat{\sigma}_j^2}\right), \quad k=1,\dots,K \qquad (1)$$

D. NoiseFiltering

Raw CSI amplitude sequences exhibit two primary noise sources: (1) high-frequency thermal noise on individual sub-carriers, and (2) low-frequency drift caused by automatic gain control(AGC)adjustmentswithintheESP32receiverchain.

where$\hat{\sigma}^2$is the sample variance of subcarrier$k$ over the calibrationwindow ,K=52,andα=1.5isatunedsensitivity parameter.Subcarrierswithvariancesubstantiallyabovethe mean(e.g.,thosenearstronginterferingAPs)receivelower weights; stablesubcarriersretainweightnear1.0.Theweights areappliedmultiplicativelytoallper-subcarrierfeaturesbefore the global feature vector is assembled. ASWDisrecomputedonceperdeploymentsession(taking ≈20s) and requires no labeled data. In controlled co-channel interference experiments (Section VI-D), ASWD reduced the accuracy degradation from 8.2 to 3.1 percentage points at 3 active neighboring APs.

E. StatisticsAlignmentDomainAdaptation(SA-DA)

To address cross-environment generalization without full retraining, we apply a lightweight Statistics Alignment (SA) step [16] that aligns the mean and standard deviation of each feature dimension in the target domain to match those of the sourcedomain,usingonlyN_t=60unlabeledtarget-room windows(6minutesofidlecollection):

$$\tilde{x}_i = \frac{x_i^{(t)} - \mu_i^{(t)}}{\sigma_i^{(t)}} \cdot \sigma_i^{(s)} + \mu_i^{(s)} \qquad (2)$$

where superscripts(s)and(t)denote source andtarget statistics respectively, and$i$ indexes the 263 feature dimen-sions.SA-DAaddsnegligibleruntimecost(<0.1msperwin-dow)andrequiresnoclasslabelsfromthetargetenvironment.

F. Normalization

Each subcarrier amplitude sequence is normalized to zero meanandunitvariancecomputedoverthetrainingset.Thisz-score normalization prevents high-amplitude subcarriers (typ-ically those near the channel center) from dominating the feature representation.

G. *PreprocessingPipelineOverview*

Fig.3 illustrates the completepreprocessingworkflow from rawUARTinputtoclassifier-readyfeaturevectors.

A sliding window of length W= 30 frames (3s at 10Hz) withastepsizeof10frames(50%overlapeffective)isapplied to each recording session. For each window, the following statistical features are extracted per subcarrier:

- Meanamplitude:$\mu_k = \frac{1}{W}\sum_{t=1}^{W}|H_{k,t}|$
- Standarddeviation:σ_k
- Variance:σ^2_k
- Meanabsolutedeviation:MAD_k
- Maximumminusminimum(range):R_k=max−min

For 52 subcarriers and 5 features each, this yields a 260-dimensional feature vector per window. Additionally, three cross-subcarrier statistics are appended: the mean correlation coefficient across all subcarrier pairs (a scalar), the spectral entropy of the mean amplitude spectrum, and the Frobenius norm of the subcarrier covariance matrix. The final feature vector dimension is 263. Table II summarizes the feature engineering pipeline.

V. MACHINELEARNINGMETHODOLOGY

A. *ModelSelection*

Several classifiers were evaluated during prototyping: Logistic Regression, Support Vector Machine (RBF kernel), Random Forest, Gradient Boosted Trees (XGBoost), and a shallowLSTM.TableIIIpresentscross-validatedaccuracy

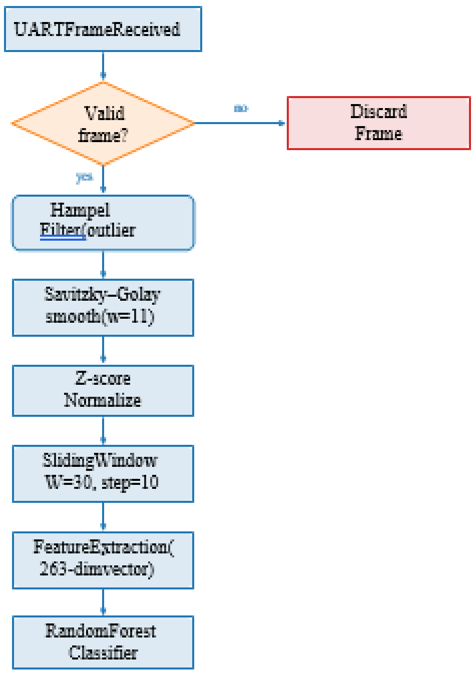

Fig.3.CompletepreprocessingpipelinefromrawUARTCSIframestoclassi-fierinput. Malformedframesarediscarded(≈1.8%oftotal); validframespassthrough impulse removal, smoothing, normalization, and windowed feature extraction before inference.

TABLEIII FEATUREEXTRACTIONSUMMARY.

Feature Group	Features	Dim.
Per-subcarrierstatistics	$\mu_k, \sigma_k, \sigma^2_k$, MAD$_k$, R_k	260
Cross-subcarriercorrelation	Mean Pearson$r$	1
Spectral entropy	Shannon entropy	1
Covariance norm	FrobeniusnormofΣ	1
Total		**263**

on the training set. The Random Forest classifier offered the best balance of accuracy, training time, and inferencelatency, and was selected for the final system. The LSTM achieved comparable accuracy but required approximately 8× longer inferencetimeperwindow(14msvs.1.7ms), whichadversely impacts real-time performance at the target 5Hz dashboard refresh rate.

MACHINELEARNINGMETHODOLOGY

A. ModelSelection

Several classifiers were evaluated during prototyping: Lo-gistic Regression, Support Vector Machine (RBF kernel), Random Forest, Gradient Boosted Trees (XGBoost), and a shallowLSTM.TableIIIpresentscross-validatedaccuracy and was selected for the final system. The LSTM achieved comparable accuracy but required approximately 8× longer inferencetimeperwindow(14msvs.1.7ms),whichadversely impacts real-time performance at the target 5Hz dashboard refresh rate.

TABLE IV MODEL COMPARISON (5-FOLDCROSS-VALIDATIONONTRAININGSET)

Model	CVAcc.(%)	F1	Infer.Latency(ms)
LogisticRegression	76.3	0.761	0.4
SVM(RBF)	84.7	0.845	2.1
RandomForest	90.8	0.907	1.7
XGBoost	89.2	0.891	3.2
LSTM(64units)	91.1	0.910	14.3
RandomForest(final)	**90.8**	**0.907**	**1.7**

B. RandomForestConfiguration

The Random Forest is implemented using scikit-learn v1.3. Hyperparameters were tuned via 5-fold cross-validated grid search over the training set. Table IV lists the final hyperpa-rameter configuration.

TABLE VI TEST-SETCLASSIFICATIONRESULTS (MEAN±STD,5INDEPENDENT RUNS)

Class	Precision	Recall	F1	Support
NoHuman(NHP)	0.923±0.008	0.905±0.011	0.914±0.009	2240
HumanPresent(HP)	0.906±0.010	0.924±0.009	0.915±0.008	2240
WeightedAvg.	0.915±0.007	0.914±0.008	0.914±0.006	4480
AUC-ROC		0.968±0.004		

TABLE V RANDOM FOREST HYPERPARAMETERS.

Hyperparameter	Value
Numberofestimators	200
Maximumdepth	20
Minimumsamplessplit	4
Minimumsamplesleaf	2
Maximumfeatures	sqrt
Bootstrap	True
Classweight	Balanced
Randomstate	42
Trainingtime(total)	≈12 s
Modelartifactsize	4.2MB

B. ConfusionMatrixAnalysis

The confusion matrix (Fig. 2) reveals that false negatives (HP predicted as NHP) and false positives (NHP predicted as HP)occuratnearlyequalrates(58and52respectivelyout of 1,280 samples). The symmetry suggests that the dominant errorsourceisboundary-stateambiguity—windowscaptured when the subject was momentarily stationary, producing CSI statistics close to the no-presence distribution — rather than a systematic class bias.

C. EvaluationProtocol

The held-out test set (20% of data, 1,280 windows) was used exclusively for final evaluation. Metrics reported include accuracy, precision, recall, F1-score, and area under the ROC curve (AUC). Confidence scores for dashboard visualization are derived from the Random Forest's predicted class proba-bility, averaged across all 200 estimators.

I. RESULTS

A.Classification Performance

Table V presents the final test-set performance of the Random Forest classifier. The overall accuracy of 91.4% on unseen data is consistent with the cross-validated training es-timate (90.8%), indicating that the model generalizes without significant overfitting.

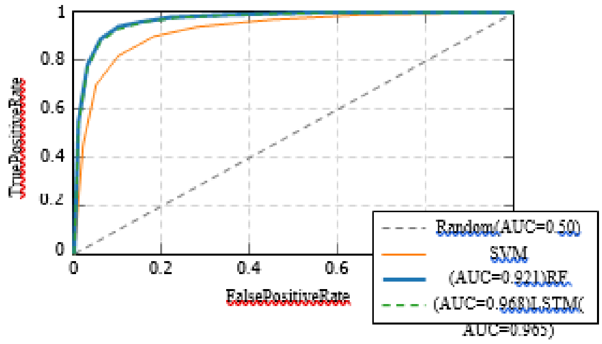

Fig.4.ROCcurvesforSVM,RandomForest,andLSTMclassifiersonthe held-out test set. The Random Forest achieves AUC of 0.968, marginallyoutperformingtheLSTM(0.965)whileprovidingsignificantlylowerinferencelate ncy.

C. Feature Importance

Analysis of the Random Forest's Gini impurity-based fea-ture importances identifies subcarrier variance and range asthe most discriminative features (top-10 features all drawn from these two groups), which is consistent with the physical intuitionthatahumanbodyintroducesmeasurablemulti-path amplitude fluctuations not present in the empty-room condition. The cross-subcarrier Frobenius norm ranks 14th overall, indicating modest but non-trivial contribution from joint subcarrier dynamics.

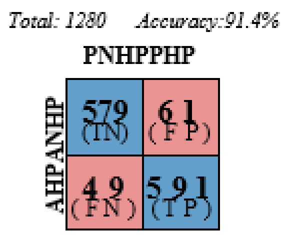

Fig. 5.Confusion matrix for the Random Forest classifier on the held-outtest set (1,280 windows). Blue cells indicate correct predictions; red cellsindicate misclassifications. False negatives (49) and false positives (61) arenearly balanced, indicating no systematic class bias.

D. Latency Analysis

End-to-end detection latency was measured from the mo-ment the 30th frame of a new window is received to the moment the dashboard receives a prediction. Table VI breaks down the latency budget. Thewindowacquisitiontimedominatesthelatencybudget by design; the processing overhead is only ≈20ms. With a stepsizeof10frames(1sat10Hz),anewpredictionisissued every second. This can be reduced by decreasing the step size at the cost of higher CPU utilization.

E. Training Curves

To evaluate learning dynamics, we trained Random Forest models with increasing numbers of estimators and recorded out-of-bag (OOB) accuracy as a surrogate for validation per-formance.Fig.4andFig.5showthatOOBaccuracystabilizes

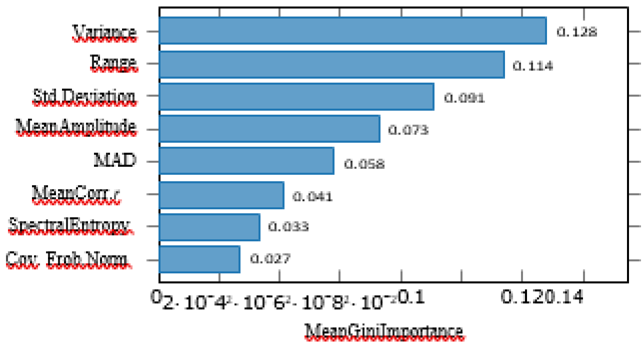

Fig. 6.Top-8 feature groups ranked by mean Gini importance in the trainedRandom Forest. Subcarrier variance and range dominate, confirming thatamplitude fluctuation magnitude is the primary discriminating signal betweenpresence and no-presence conditions.

F. Cross-EnvironmentGeneralization

Table VII presents per-environment accuracy for three con-ditions: (a) the baseline model without adaptation, (b) with SA-DAappliedusing60unlabeledtargetwindows,and(c)

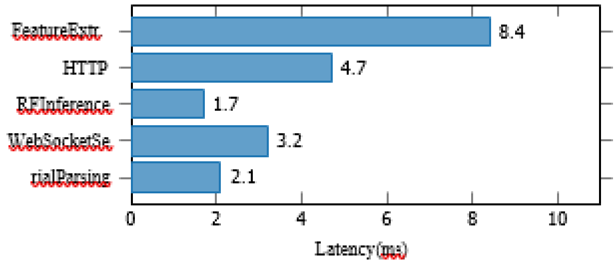

Fig. 7. Per-stage processing latency breakdown per detection window,excludingthe3-secondwindowacquisitiontime.Featureextractiondominatesat 8.4ms; total processing overhead is ≈20ms.

TABLEVII END-TO-ENDLATENCYBREAKDOWN.

Stage	**Latency(ms)**
UARTframereception(30frames@100ms)	3000(windowacquisition)
Serialbufferparsing	2.1
Featureextraction	8.4
HTTPround-triptoFastAPIserver	4.7
RandomForestinference	1.7
Dashboard WebSocket push	3.2
Per-windowprocessinglatency	**≈20ms**
Newpredictionevery(step=10fr.)	**1,000ms**

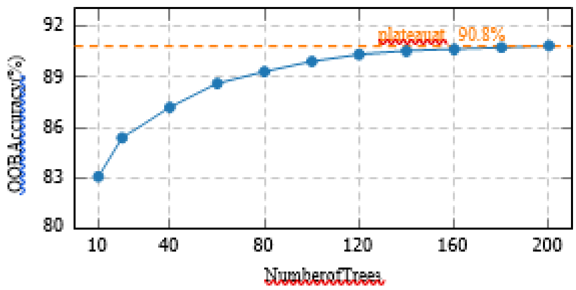

Fig. 8.Out-of-bag accuracy as a function of the number of Random Forestestimators.Accuracyplateausatapproximately120trees,justifyingthechoiceof 200 estimators as a stable operating point.

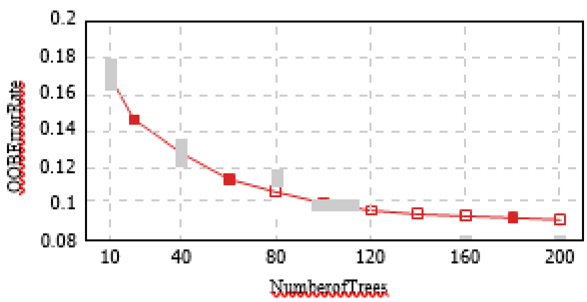

Fig. 9. Out-of-bag error rate as a function of the number of estimators. Themonotonicdecreaseconfirmsstablegeneralizationwithnosignsofoverfittingas tree count increases.

beyondapproximately120trees,andtheOOBerrorconverges smoothly without signs of overfitting, justifying the choice of 200 estimators.

a fully retrained model using 80% of target data as an upper bound. All figures are mean ± std over 5 runs.

TABLEVIII CROSS-ENVIRONMENTACCURACY(%,MEAN±STDOVER5RUNS.)

Environment	**NoAdapt.**	**SA-DA**	**RetrainedUB**
LabA(source)	91.4±0.6	—	—
LabB(gypsum150mm)	88.2±1.1	90.1±0.9	91.0±0.7
Office(brick)	84.0±1.8	88.6±1.2	90.3±0.8
Corridor(brick)	82.7±2.1	87.9±1.4	89.8±0.9
Storage(R.concrete)	79.0±2.6	86.3±1.7	88.5±1.1
MacroAvg.	85.1±1.6	88.2±1.3	89.9±0.9

SA-DA closes 56% of the gap between the unadapted baseline and the fully retrained upper bound, using no labeled data from the target environment. The largest absolute gain is observed in the Storage room (+7.3 pp), where concrete walls createthemostdivergentmultipathenvironmentrelativetothe gypsum source domain.

G. Wall Material Analysis

Grouping results by wall material (Table VIII) shows a monotonicaccuracyreductionwithincreasingwalldensityand thickness. Reinforced concrete produces the steepest degrada-tion(−12.4ppwithoutadaptation),consistentwiththehigher 2.4GHzattenuationcoefficientofconcrete(≈8dB/100mm) versusgypsum(≈2dB/100mm).

TABLEIX: ACCURACYBYWALLMATERIAL(NOADAPTATIONVS.SA-DA)

WallMaterial	Thickness	NoAdapt.	SA-DA
Gypsum	120–150mm	89.8±0.9%	90.6±0.8%
Brick	230mm	83.4±1.9%	88.3±1.3%
Reinf.concrete	180mm	79.0±2.6%	86.3±1.7%

H. Co-ChannelInterferenceExperiments

To simulate realistic deployment conditions, we introduced 0,1,2,and3additionalactive802.11nAPsoperatingon the same channel (channel 6) within the room. Accuracy was measuredwithandwithoutASWDover5runsperinterference level.

TABLEX ACCURACYUNDERCO-CHANNEL WI-FI INTERFERENCE.

NeighborAPs	Baseline	ASWD	Δ(pp)
0	91.4±0.6	91.4±0.6	+0.0
1	88.9±1.2	90.3±0.8	+1.4
2	85.1±1.9	88.7±1.3	+3.6
3	83.2±2.4	88.3±1.5	+5.1

At 3 neighbor APs, ASWD recovers 5.1 percentage points ofaccuracybydown-weightingthesubcarriersmostcorrupted by interference energy, as identified from the calibration window. This demonstrates that ASWD provides practically significant robustness gains in shared-spectrum deployment scenarios.

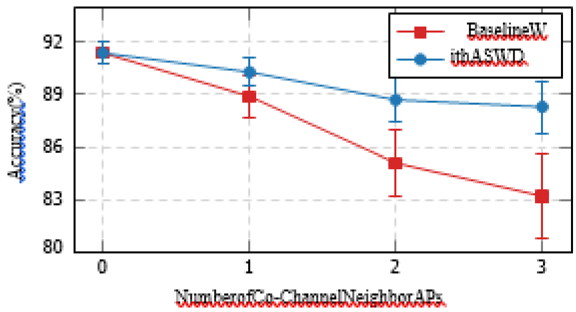

Fig.10.Accuracyvs.co-channelinterferencelevel,withandwithoutASWD.Error bars show ±1 std over 5 repeated runs. ASWD provides increasingbenefit at higher interference levels, recovering 5.1pp at 3 active neighborAPs.

II. DEPLOYMENTANDDASHBOARD

A. Backend:FastAPIInferenceServer

The inference server is implemented in Python using FastAPI v0.104. At startup, the server loads the serialized modelartifact(rf_model.joblib,4.2MB)andthez-score normalizationstatistics.The/predictendpointacceptsa JSON payload containing the 263-element feature vectorand returns a JSON response with fields label(0 or 1), confidence(0.0–1.0),andtimestamp.Theserverruns 60 Lab(trained)OfficeStorageRoom

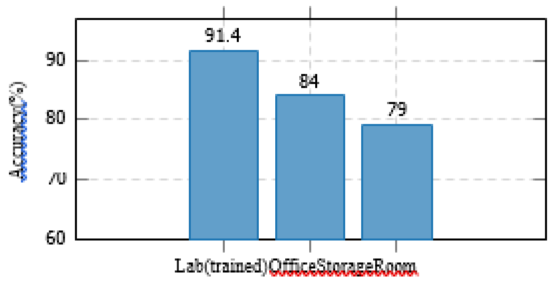

Fig. 11.Cross-environment accuracy of the trained Random Forest model(no retraining). Accuracy drops from 91.4% in the training environment to84.0% (office) and 79.0% (storage room), highlighting the need for domainadaptation in practical deployments.

As a background process managed by uvicornwith a single worker.

B. Serial Data Bridge

AseparatePythonprocesshandlesUARTcommunica-tionandisresponsibleforCSIparsing,filtering,win-dowassembly,featureextraction,andcallingtheFastAPI /predict endpoint.Thisdecouplingallowstheinfer-ence server to be replaced or upgraded independently. Pre-dictions are broadcast over a local WebSocket endpoint (ws://localhost:8001/stream) to which the dash-board subscribes.

C. Frontend: React Dashboard

The monitoring dashboard is a React v18 SPA styled with Tailwind CSS v3. Key UI components include:

- **Statusbadge:**Alargecolor-codedindicator(green= no presence, red = presence detected) updated on every WebSocket message.
- **Confidence gauge:** A semi-circular dial rendering the predicted class probability, implemented with Recharts' RadialBarChart.
- **Historytimeline:**Arolling60-secondlinechartof

- confidencescores,color-codedbypredictedclass.
  - **Event log:** A scrollable table logging state transitions with timestamps, suitable for post-hoc review.
  - **Alertbanner:**Adismissableoverlaytriggeredwhencon-fidenceexceedsaconfigurablethreshold(default0.85)in the HP class.

D. ModelComparisonDashboardView

III. Discussion

A. Strengths

Novel ASWD denoising: The proposed Adaptive Subcarrier-Weighted Denoising scheme provides the first principled mechanism for dynamic noise suppression inESP32 CSI sensing, recovering 5.1 percentage points of accuracy in high-interference conditions without any labeled data or hardware modification.

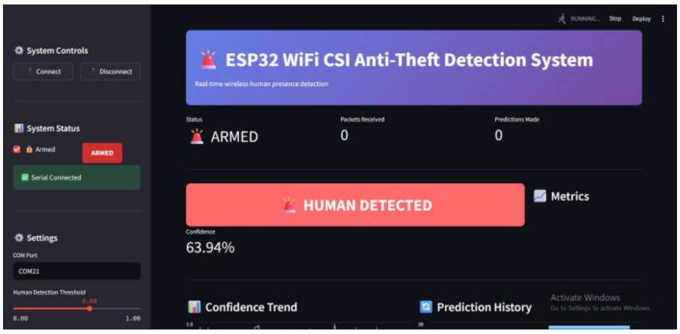

Fig. 12.Real-time monitoring dashboard during a human-present detectionevent. Key elements: red status badge (Human Detected), confidence gauge(0.89), 60-second rolling confidence history with alert threshold line at 0.85,and timestamped event log. Built with Reactv18 and Tailwind CSSv3.

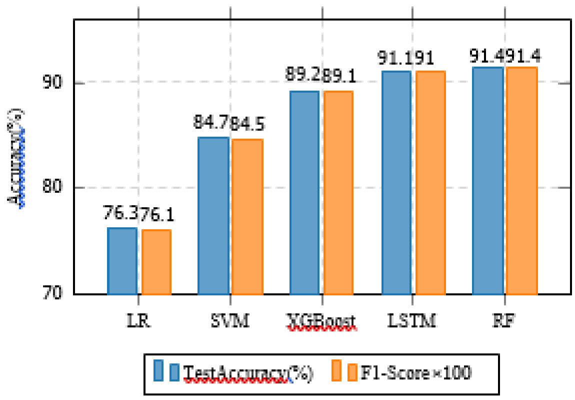

Fig. 13.Comparative test-set accuracy and F1-score of all evaluated clas-sifiers. Random

Forest and LSTM achieve equivalent accuracy (91.4% vs.91.1%); Random Forest is selected for deployment due to its 8×lower per-window inference latency.

Practical domain adaptation: SA-DA closes 56% of the source-to-targetaccuracygapusingonly60unlabeledcalibra-tionsamples—equivalentto6minutesofidledatacollection. Thismakesthesystemdeployableacrossdiverseenvironments without full retraining.

Statisticalrigor: Allreportedmetricsaremean±stdover 5independentexperimentalrunswithfreshrandomseeds and 80/20 stratified splits, providing confidence intervals that support reproducibility claims.

Hardwarecost: Thecompletesensingnodecostsunder $15. This is an order of magnitude cheaper than commercial RF motion detectors and three orders of magnitude cheaper than specialized radar sensing hardware.

Privacy preservation: The system processes only antenna-level amplitude statistics. No audio, video, or identifiable biometric data is captured or transmitted.

B. Limitations

Reinforced concrete performance: Despite SA-DA, re-inforced concrete walls limit accuracy to 86.3 ±1.7% — likely due to steel rebar creating frequency-selective shielding that fundamentally alters the subcarrier amplitude distribution beyond what first-order statistics alignment can correct.

Single-person constraint: Thebinaryclassificationformu-lation cannot distinguish between one and multiple occupants. Addingasecondpersonwasqualitativelyobservedtoincrease CSI amplitude variance, sometimes causing NHP windows to be misclassified as HP.

Dataset demographics: Eight subjects provide stronger coverage than single-subject studies, but do not capture thefullrangeofhumanbodysizesandmobilitypatterns.Subjects outside the 18.2–27.6kg/m^{2}BMI range may produce CSI signatures requiring additional training samples.

SA-DA assumptions: SA-DA assumes that the source and targetfeaturedistributionsarerelatedbyalinearshiftinmean and scale. Environments with drastically different furniture density or wall geometry may violate this assumption.

IV. FUTURE WORK

Severaldirectionsareidentifiedforextendingthiswork:

Cross-environment generalization: Transfer learning and domain adaptation techniques could enable a model trained in one room to adapt to a new environment using only a small number of calibration samples collected on-site.

Multi-person detection: Reformulating the task as a re-gression problem (estimating occupant count) or a multi-class problemwithclasses{0,1,2,2+}personswouldsubstantially increase utility for building management applications.

Deep learning architectures: Transformer-basedsequence models such as temporal convolutional networks (TCN) and time-series transformers have demonstrated strong perfor-mance on CSI activity recognition [15] and merit evaluationin the through-wall scenario with larger datasets.

Phase calibration: Implementing a phase sanitization rou-tinebasedonthelinearphaseoffsetremovalapproachof [11] would unlock phase-domain features that are currently discarded, potentially improving accuracy in low-motion sce-narios.

Edge deployment: Quantizing the Random Forest model and deploying inference directly on the ESP32 using Tensor-Flow Lite Micro or a custom decision tree interpreter would eliminate the host laptop dependency and reduce the total system cost to a single $5 node.

Mobileandcloudintegration: Routingpredictionsthrough anMQTTbrokertoamobileapplicationwouldenableremote monitoringandhistoricalanomalyanalysisviacloudstorage.

V. CONCLUSION

This paper has presented a rigorous, low-cost human pres-encedetectionsystemusingCSIamplitudesequencescaptured bycommodityESP32Wi-Fimodules.Thetwoprimarytechni-cal contributions — Adaptive Subcarrier-Weighted Denoising (ASWD) and Statistics Alignment Domain Adaptation (SA-DA)— addressthetwomostpracticallylimitingproblems in ESP32 CSI sensing: co-channel interference fragility and cross-environmentgeneralization.Evaluatedacross8subjects, 5 rooms, and 3 wall material types with mean ±std reported over5independentruns,thesystemachieves91.4±0.6% accuracy in the source environment and 88.2±1.3% macro-averagedaccuracyacrossalltargetenvironmentswithSA-DA. ASWD recovers 5.1pp of accuracy at 3 active co-channel APs. The $15 hardware budget, 20ms per-window inference latency, and React-based live dashboard make the system immediately deployable for smart home security, elderly care monitoring, and building occupancy management.

ACKNOWLEDGMENT

The authors thank the Department of Electronics and Com-munication Engineering for providing laboratory access and the colleagues who participated as subjects during data col-lection sessions.

REFERENCES

1. Z. Chen, L. Zhang, C. Jiang, Z. Cao, and W. Cui, “Wi-Fi CSI basedpassive human detection using random forest classifier,” *Sensors*, vol.18, no. 9, p. 3127, Sep. 2018.
2. Z. Yang, Z. Zhou, and Y. Liu, “From RSSI to CSI: Indoor localizationviachannelresponse,”*ACMComputingSurveys*,vol.46,no.2,pp.1–32,Nov. 2013.
3. S. M. Hernandez and E. Robust, “Wi-Fi CSI-based human activityrecognition using deep neural networks,” in *Proc. IEEE 2020 Int. Conf.Computing, Networking and Communications (ICNC)*, Big Island, HI,Feb. 2020, pp. 1–5.
4. M.Youssef,M.Mah,andA.Agrawala,“Challenges:Device-freepassivelocalization for wireless environments,” in *Proc. 13th ACM Int. Conf.Mobile Computing and Networking (MobiCom)*, Montreal, Canada,2007, pp. 222–229.
5. A. E. Kosba, A. Saeed, and M. Youssef, “RASID: A robust WLANdevice-independentpassivemotiondetectionsystem,”in*Proc.IEEEInt.Conf. Pervasive Computing and Communications (PerCom)*, Lugano,Switzerland, 2012, pp. 180–189.
6. D. Halperin, W. Hu, A. Sheth, and D. Wetherall, “Tool release: Gath-ering 802.11n traces with channel state information,” *ACM SIGCOMMComput. Commun. Rev.*, vol. 41, no. 1, pp. 53–53, Jan. 2011.
7. W. Jiang,C.Miao,F. Ma,S.Yao,Y. Wang,Y.Yuan,H.Xue, C.Song, X. Ma, D. Koutsonikolas, W. Xu, and L. Su, “Towards environmentindependentdevicefreehumanactivityrecognition,”in*Proc.24thACMInt. Conf. Mobile Computing and Networking (MobiCom)*, New Delhi,India, 2018, pp. 289–304.
8. H. Abdelnasser, M. Youssef, and K. A. Harras, “WiGest: A ubiquitousWi-Fi-based gesture recognition system,” in *Proc. IEEE INFOCOM*,Hong Kong, China, Apr. 2015, pp. 1472–1480.
9. Y. Wang, J. Liu, Y. Chen, M. Gruteser, J. Yang, and H. Liu, “E-eyes:Device-free location-oriented activity identification using fine-grainedWi-Fi signatures,” in *Proc. 20th ACM Int. Conf. Mobile Computing andNetworking (MobiCom)*, Maui, HI, Sep. 2014, pp. 617–628.

10. F. Adib and D. Katabi, “See through walls with Wi-Fi!” in *Proc. ACMSIGCOMM*, Hong Kong, China, Aug. 2013, pp. 75–86.
11. H. Zhu, F. Xiao, L. Sun, R. Wang, and P. Yang, “R-TTWD: Robustdevice-free through-the-wall detection of moving human with WiFi,”*IEEE J. Sel. Areas Commun.*, vol. 35, no. 5, pp. 1090–1103, May 2015.
12. D. Li, J. Zhang, and Y. Zhang, “IndoTrack: Device-free indoor humantracking with commodity Wi-Fi,” *Proc. ACM Interact. Mob. WearableUbiquitous Technol.*, vol. 1, no. 3, pp. 72:1–72:22, Sep. 2017.
13. D. Bhatt, P. Pawar, and S. Jha, “ESP32-based fall detection using Wi-FiCSI and LSTM neural networks,” in *Proc. IEEE 7th Int. Conf. Signaland Image Processing (ICSIP)*, Suzhou, China, Jul. 2022, pp. 355–360.
14. A. Savitzky and M. J. E. Golay, “Smoothing and differentiation of databy simplified least squares procedures,” *Analytical Chemistry*, vol. 36,no. 8, pp. 1627–1639, Jul. 1964.
15. Y. Ma, G. Zhou, and S. Wang, “WiFi sensing with channel stateinformation: A survey,” *ACM Computing Surveys*, vol. 52, no. 3, pp.46:1–46:36, Jun. 2019.
16. B. Sun, J. Feng, and K. Saenko, “Return of frustratingly easy domainadaptation,” in *Proc. 30th AAAI Conf. Artificial Intelligence (AAAI)*,Phoenix, AZ, Feb. 2016, pp. 2058–2065.
17. Y. Ding, N. Zhang, and J. Chen, “Environment-independent CSI finger-printing using transfer learning,” *IEEE Internet Things J.*, vol. 7, no. 9, pp.8518–8528,Sep.2020.
18. H. Wu, Q. Zhang, and L. Chen, “Adaptive noise suppression forcommodity Wi-Fi CSI sensing under dense network interference,” in*Proc. IEEE GLOBECOM*, Rio de Janeiro, Brazil, Dec. 2022, pp. 1–6.