Interpretable Room-Level Human Presence Detection Using Ambient Sensors in Smart Homes

Sehrish Rafique[1,2], Patrick Holthaus[1], Gu Fang[2], and Farshid Amirabdollahian[1]

[1] Robotics Research Group, University of Hertfordshire, Hatfield, United Kingdom
{s.rafique, p.holthaus, f.amirabdollahian2}@herts.ac.uk
[2] School of Engineering, Design and Built Environment, Western Sydney University, Sydney, Australia
22200820@student.westernsydney.edu.au, g.fang@westernsydney.edu.au

Abstract. Ambient Assisted Living (AAL) systems are becoming increasingly important for providing personalised assistance in smart homes. One key component for such systems is detecting and localising humans in different areas of the home, which can enhance contextual information to provide efficient support to the human user. Recent approaches often lack interpretability and compromise user privacy.
This work introduces an interpretable, room-level human presence detection system that relies solely on low-cost, privacy-conserving ambient sensors typically used in smart homes. We have developed and evaluated a solution for presence detection based on data collected from a single participant in the Robot House, an ambient assisted living space at the University of Hertfordshire. We developed two models to perform this task, a Random Forest (RF) model and a more complex Long Short-Term Memory (LSTM) model across a triad of test scenarios, including full sensor set, sensor dropout and room dropout. We tested the performance of both models using conventional train-test splits and on an entirely unseen data to assess the generalisation. While LSTM achieved comparable results, RF performed better on new, unseen data, with an accuracy of 91.43% vs. 62.69% for RF and LSTM, respectively. The RF also achieved comparative results against two state-of-the-art models, HOOD and CSI-BiLSTM, with the advantages of being easy to interpret and working better in situations where privacy and cost are important. Overall, our work provides the basis for creating a scalable and interpretable solution for finding a person's location in smart homes.

Keywords: Assistive Technologies, Sensors, Smart Homes, Ambient Assisted Living, Virtual Wards

1 Introduction

Smart homes equipped with assistive technologies can support people, especially older adults and individuals with special needs, through automation, enhancing safety, improving energy management, and enabling personalised assisted living.

A key component in this process is human presence detection to inform the system where the user is and thus where to provide assistance [1]. This supports vital capabilities in AAL including personalised assistance, context-driven interventions, and insights about human behaviours. There has been existing work on this interesting and important area. Many approaches today use cameras, Bluetooth devices, Wireless Fidelity (Wi-Fi), and Channel State Information (CSI) signals to detect the person's exact location within the home. On top of that, they have used deep learning methods for processing the data [2], [3], [4], [5], which provide excellent baseline accuracy but are typically difficult for people to interpret and trust and often do not work well in noisy real-world environments.

In contrast, ambient sensors such as motion, status, plug sensors, and pressure mats are widely used in smart homes and AAL systems. Such sensors are usually low-cost, privacy-preserving, and easy to install. However, they are often overlooked for tasks such as human presence detection and localisation tasks [1]. The majority of previous works have utilised them for basic motion detection or only binary presence detection for the entire home (simply identifying if someone is present at home or not, rather than detailed room-level information, i.e., determining which specific room a person is in). There is a gap in the literature regarding their use for multi-room classification and generalisation to unseen data [5], [3].

In this study, we have examined the use of such ambient sensors for indoor human presence detection. We required a high-quality dataset that could accurately capture fine-grained human movements and room transitions in a real smart home. To the best of our knowledge, no publicly available dataset provides the level of detail, annotation quality, and realistic multi-room coverage required for this study. Therefore, we collected the data at the University of Hertfordshire's Robot House [6], a residential home dedicated to studies on human-robot interaction. We have utilised the standard ambient sensors (e.g. motion, contact, pressure, plug) already embedded in the smart home environment. In our data, we ensured that every transition was recorded and labelled with high temporal precision to help in robust training and evaluation under realistic conditions. We created three experimental scenarios of room-level presence detection to see if a simple sensor system can reliably detect room-level presence and work well with unseen data. For these scenarios, we trained two separate models LSTM and RF. We selected LSTM because it is widely used for sequence modelling in sensor data, while RF was chosen for its interpretability and robustness to class imbalance. Together, they let us compare a complex temporal model with a lightweight, interpretable baseline. We evaluated the model's performance in two settings: an 80/20 offline train-test split on the primary data and a separate, unseen data. We found that RF not only performs just as well as LSTM on known data, but it significantly outperforms LSTM in the generalisation scenario when facing unknown data. Moreover, RF is easy to interpret, as some interpretations are mentioned in Section 5. Based on these results, we suggest that RF is a better choice for real-world deployment. Our key contributions show that a lightweight, interpretable RF model using only standard ambient sensors achieved 99.15% offline accuracy and maintained 91.43% accuracy on unseen data, outperforming LSTM. Additionally, we also compared it to the state-of-the-art HOOD [3] and

CSI-BiLSTM [5], showing that our approach does not require any new infrastructure, effectively handles a more fine-grained multi-class problem and still performs well on new data.

The remaining sections of this paper are organised as follows. Section 2 reviews literature on presence detection methods and highlights some of the research gaps that we address in this paper. Section 3 describes our data collection process, experimental environment, sensor infrastructure, and feature extraction. Section 4 describes our evaluation scenarios and metrics. Section 5 describes our recognition pipeline, training process, and evaluation methods. Section 6 presents experimental results, model generalisation on new data, and comparison with state-of-the-art methods. Section 7 summarises our main findings, potential real-world applications, and suggestions for future work.

2 Related Work

This section reviews existing technologies and methods for indoor human presence detection, highlighting their strengths, limitations, and the gap to be addressed. Real-time human presence detection has received significant attention for a decade. For example, one approach uses infrared sensors with an adaptive k-nearest neighbour (KNN) model for real-time human presence detection [7]. They demonstrated effectiveness in controlled smart home scenarios; however, their method faced challenges in handling false alarms and adapting to changing environments, particularly in real-time scenarios. CO2 sensors have also been utilised for indoor presence detection [8]. However, specialised equipment and sensitivity to external factors such as ventilation and pet activity affect the performance. The Inertial Measurement Unit (IMU) sensors embedded in smartphones and smartwatches have been used to enhance human presence detection, localisation, human activity recognition, and study biomechanics and movement patterns [9]. All of these methods have been tested in controlled laboratory environments with static and offline training methods. In a notable work, Bidirectional Long Short-Term Memory BiLSTM networks utilising Wi-Fi. CSI have been used for real-time indoor human presence detection [5]. They demonstrated impressive results with real-time tests, achieving an accuracy of 91.8% on training and testing data, and 88.67% on unseen data collected under different room conditions. Another notable real-time detection system, HOOD, has been introduced to detect human presence and unusual (out-of-distribution) samples in real time [3]. Their notable contribution includes combining human presence detection with OOD detection to classify scenes as ”presence” (human detected) or ”no presence” (any non-human objects present). However, the focus was on static human activities, such as standing or sitting and not on dynamic activities like moving around the house. HOOD achieved an Area Under the Receiver Operating Characteristic Curve (AUROC) of 95.71% for static activities and 93.02% for very static, outperforming existing radar-based human presence detection methods in all scenarios, with real-time capability validated in both offline and online experiments. Such systems provide motivation and potential to build real-time applications in dynamic environments for human presence detection.

The literature shows that the majority of human presence and localisation systems are tested in controlled environments and trained on offline, pre-recorded datasets, limiting the validity of results in real-world applications [10], [11], [12]. In real-world deployment, variations in room layouts, furniture placement, and environmental conditions can negatively impact system accuracy [10]. Privacy concerns further add to the complication of techniques that rely on cameras, smartphone sensors, or other intrusive data collection methods [13]. There is still a need for real-time, lightweight, fine-grained, room-level location classification, without invading user privacy. To address these challenges, our system demonstrates how standard ambient sensors and interpretable models can deliver privacy-preserving, scalable, room-level presence detection that generalises well to new data as well.

3 Data Collection

Since no available smart home dataset offered detailed information for multi-room human presence detection with precise room transition labelling, we created a new data with fully labelled room transitions under real-world conditions. This section describes our data collection process, sensors, the floor plan of Robot House [6], ground truth annotation and feature extraction techniques.

3.1 Sensors and Layout

This experiment was conducted with a single participant in the University of Hertfordshire's Robot House, where each room is equipped with ambient sensors. These sensors include motion sensors, pressure mats, contact sensors (doors, drawers, and toilet lid), plug monitors (microwave, kettle, and coffee machine), and other environmental sensors (brightness and water flow). We used all of these sensors in our data, since using only a single time-series motion sensor is insufficient for accurate room-level presence detection. Motion sensors only give on/off signals, so they often miss activities like sitting or sleeping. They also struggle to distinguish between movement in nearby rooms, which creates confusion. Together, the combination of these sensors provides a fuller picture by detecting both movement and still activities. They also add context from how objects are used, and reduce room overlap, leading to more accurate and detailed presence detection. Room labels include kitchen, bedroom, bathroom, corridor, upstairs, and living room (divided into sofa area and dining area). Sensor distribution is shown in Table 1, and the layout of the Robot House is given in Figure 1.

3.2 Ground Truth

Room-level ground truths were manually annotated using synchronised video recordings. Videos were used solely for annotation, not for model input. We chose manual annotation because we wanted to be as precise as possible, especially when marking transitions between rooms. Automated methods were not a good fit for this study since they are prone to errors and cannot be fully trusted.

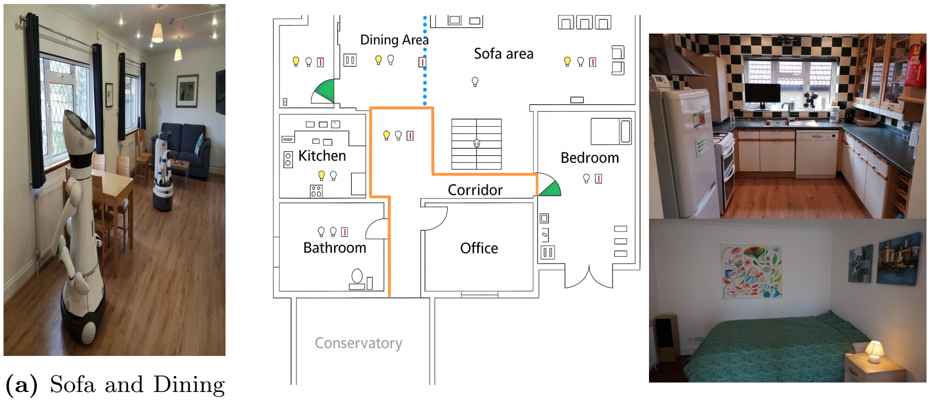

(a) Sofa and Dining Area

(b) Robot House Layout

(c) Kitchen and Bedroom

Figure 1: Real views of key areas in the Robot House and the corresponding layout plan. Sensors are shown as colored icons placed at their respective locations. The borders between different areas are marked with colored lines, orange lines mark corridor boundaries, and the blue dotted line shows the border between the dining and sofa areas.

Table 1: Sensor distribution by location

Location	**Sensors**
Corridor	Motion and brightness
Bathroom	Motion and brightness, contact (door, toilet lid)
Bedroom	Motion and brightness, pressure mat (bed), contact (door)
Dining Area	Motion and brightness
Kitchen	Motion and brightness, power (microwave, kettle), contact (8 cupboards, fridge)
Sofa Area	Motion and brightness, pressure mat (5 seats), power (TV)
Upstairs	No sensors

3.3 Feature Extraction

Min-max scaling was used to adjust non-binary numerical values to the 0 to 1 range, so they are easy to interpret and avoid large-value dominance. Missing values are filled using *Forward fill* (fill()), which fills the missing values with the last known non-missing value. The total number of inputs was 76 and included both sensor status and values.

4 Evaluation Scenarios and Criteria

This section describes the experimental scenarios designed to test our model's performance under various conditions. We tested and evaluated models across three key conditions:

1. **Scenario 1:** This scenario was to check robustness against the data with underrepresented classes and contains all ambient sensor data from the ground floor as well as the staircase, with 76 inputs and 7 location outputs.
2. **Scenario 2:** This scenario is designed to check the change in classification performance when underrepresented rooms (corridor, upstairs) are excluded. Therefore, we used reduced data that excludes the corridor and upstairs, comprising 76 inputs and 5 location outputs.
3. **Scenario 3:** This was to check how performance holds when only minimal sensor data is used. We evaluated it by using only motion sensor data. This includes 12 inputs and 6 location outputs. Since the staircase does not have a motion sensor, it was not included.

Both RF and LSTM were evaluated using accuracy, precision, recall, F1-score, weighted and macro average of these evaluation measures, AUROC, and confusion matrices.

5 Methodology

This section details the structure training process of our RF and LSTM models for classifying room locations from sensor data. The dataset was split into training and testing sets using an 80/20 split. A separate, completely new dataset was also used to check the model's performance on unseen data.

5.1 RF Architecture and Training

RF was used to classify the location label. Since our data was imbalanced, it was expected that RF would handle. Additionally, it is easy to interpret through feature importance analysis [14].

In the parameter setting, we used 100 decision trees, with each tree limited to a maximum depth of 10, with a minimum of 5 samples required to split an internal node and each leaf node required at least 2 samples. Additionally, each tree selected a random subset of features at each split, called feature bagging. Since the total number of features is 76, the model randomly sampled a subset of approximately 8 features ($\sqrt{76}$) for every split. The LabelEncoder was applied to room labels. Each tree independently predicted an output, and the final classification was determined by majority voting from all trees. Mathematically, the final output can be written as

$$\hat{y} = \text{mode}\left(h_1(x), h_2(x), \ldots, h_T(x)\right) \tag{1}$$

Here, $\hat{y}$ is the final predicted label. Each $h_t(x)$ is the output of the t-th decision tree for the input x, and T is the total number of trees. The function mode$(\cdot)$ returns the most common prediction among all trees. The RF model was trained for 2.78 seconds with a data size of 65,920 samples. It took approximately 14.14 seconds in total (including cross-validation and inference time), on an Intel Core i7 (x86_64) CPU setup.

5.2 LSTM Architecture and Training

We implemented LSTM classifier using TensorFlow.keras to model temporal patterns and long-term dependencies [15], and to compare against the simpler but interpretable RF. We first standardised and reshaped the data into a 3D format (samples, timesteps, features). The model has two LSTM layers with 64 and 32 units (Neurons), followed by a dropout layer and a softmax output layer for multi-class classification. The final layer predicted the room using softmax. We used categorical cross-entropy loss with the Adam optimiser. Numerical features were scaled with StandardScaler, and location labels were one-hot encoded after LabelEncoding. The LSTM model was trained for 20 epochs with a batch size of 32 and a data size of 65,920 samples, resulting in a total training time of approximately 2 minutes and 2 seconds on an Intel Core i7 (x86_64) CPU setup.

$$\begin{aligned} \text{Time per epoch:} \quad & \sim 6\text{s (excluding the first epochwhich took} \sim 8\text{s due to} \\ & \text{startup overhead)} \\ \text{Total time:} \quad & \approx 8\,\text{s} + (19 \times 6\,\text{s}) \\ & = 122\,\text{s} \approx 2\,\text{minutes and } 2\,\text{seconds} \end{aligned} \tag{2}$$

6 Results and Discussion

This section compares the performance of RF and LSTM models on standard and unseen data, as described in Section 4.

6.1 Comparative Analysis of RF and LSTM on Standard Train/Test Split

We evaluated RF and LSTM on three distinct use cases as mentioned in Section 4, each differing in sensor selection and location strategies. Detailed performance metrics are shown in Table 2.

Scenario 1 – All Sensors and All Locations: **RF** achieved 95% accuracy, but failed to detect "Upstairs" (F1-score = 0). It also confused the corridor with the kitchen and dining room. In contrast, **LSTM** achieved better results, with an overall 99% accuracy and higher F1 scores for all rooms, including underrepresented classes like "Upstairs", showing better handling of underrepresented classes.

Scenario 2 – All Sensors and Excluded Sparse Locations: With sparse locations excluded, **RF** achieved strong class separation, with 99.15% accuracy, and macro and weighted F1 were 0.99. There were very few misclassifications, mostly between areas that were close to each other (like the Sofa and dining areas). **LSTM** also performed similarly well, with over 99% accuracy. It indicates that when the class instances are balanced and well-covered by sensors, both models can work very well.

Scenario 3 – Motion Sensors Only with One room out: **RF** accuracy dropped to 78.98%, with high confusion between shared spaces, like the corridor and dining. **LSTM** performed better with 87.0% accuracy, showing improved separation between overlapping classes. Both models performed poorly in the bathroom, potentially due to limited sensor input and slight irregularities in motion patterns.

6.2 Comparative Analysis of RF and LSTM for Generalisation on Unseen Data

We further evaluated both models on fully unseen data to assess the generalisation beyond the training distribution and to observe robustness under semi-real-world deployment conditions. Results are shown in Figure 2 and 3 and Table 3.

Scenario 1 – All Sensors and All Locations: **RF** reached 91.4% overall accuracy and a Weighted F1 score of 0.92, showing a strong generalisation, particularly for bedroom, kitchen, and sofa area classes; however, it struggled with unbalanced data, entirely failing for the Upstairs class, with F1 = 0.00. **LSTM** showed a significant drop, achieving only 62.7% overall accuracy, with a weighted F1 score of 0.64. It overfits the bathroom with a recall = 1.00, but precision is 0.18, and it also fails to generalise to the dining area and kitchen. Overall, RF shows far superior generalisation, LSTM suffers from overfitting and is unreliable under temporal variation.

Scenario 2 – All Sensors and Excluded Sparse Locations: **RF** again performed well with 90.5% accuracy and a Weighted F1 score of 0.89, struggling with the corridor class only. **LSTM** dropped to 53.5% accuracy and a Weighted F1 score of 0.54, and the sofa area and the corridor were completely misclassified. Again, RF showed more stable generalisation.

Scenario 3 – Motion Sensors Only with One room out: **RF** still worked decently with overall 82.0% accuracy and a Weighted F1 score of 0.80. The strongest class was the kitchen, and the bathroom and the corridor were misclassified the most. **LSTM** dropped to accuracy of 42.6% and a Weighted F1 score 0.43. It overpredicted the corridor and dining area, showing poor separation. Table 3 shows a comparison of the accuracy and weighted F-1 score of both models on unseen data. It can be concluded from generalisation testing that in a sensor-constrained setting, RF remains usable and has clear class separation, while LSTM performance collapses completely. RF performance was slightly reduced from training, but it remained resilient and balanced. Figure 2 and 3 show confusion matrices of generalisation for both models across all scenarios.

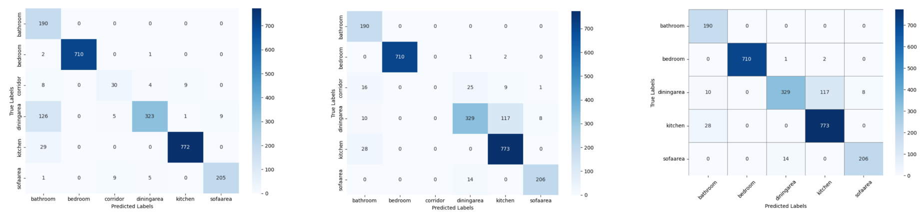

(a) RF Scenario 1 **(b)** RF Scenario 2 **(c)** RF Scenario 3

Figure 2: Confusion Matrix: RF for Generalisation on Unseen Data

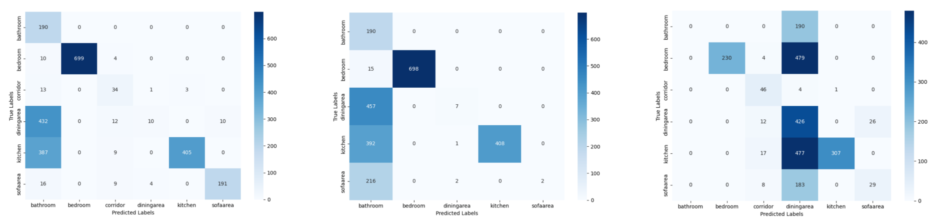

(a) LSTM Scenario 1 **(b)** LSTM Scenario 2 **(c)** LSTM Scenario 3

Figure 3: Confusion Matrix: LSTM for Generalisation on Unseen Data

Table 2: Comparison of RF and LSTM on Testing and Training Data

Model	Scenario	Accuracy	Weighted F1 Score
RF	Scenario 1	95.00%	0.94
	Scenario 2	99.15%	0.99
	Scenario 3	78.98%	0.79
LSTM	Scenario 1	**99.00%**	**0.99**
	Scenario 2	**99.25%**	0.99
	Scenario 3	**87.00%**	**0.87**

Table 3: Model performance for generalisation on unseen data

Model	Scenario	Accuracy	Weighted F1 Score
RF	Scenario 1	91.43%	0.9189
	Scenario 2	90.53%	0.8939
	Scenario 3	82.04%	0.7955
LSTM	Scenario 1	62.69%	0.6351
	Scenario 2	53.51%	0.5370
	Scenario 3	42.56%	0.4303

6.3 Comparative Analysis of RF with Previous Studies

We compared our results against two recent robust presence detection frameworks: CSI-BiLSTM [5] and HOOD [3]. The comparison focuses on generalisation

Table 4: Comparison of Proposed Work with CSI-BiLSTM [5]

Aspect	Proposed Work (Ambient Sensors)	CSI-BiLSTM (Wi-Fi CSI)
Testing Accuracy (20% test split from training data)	Scenario 1: 95.00%, Scenario 2: 99.15%	91.68%
Unseen Data Accuracy (Generalisation)	Scenario 1: 91.43%, Scenario 2: 90.53%	88.67%
Detection Task	Room-level multi-class classification	Binary presence detection
Sensor Type	Ambient sensors (motion, contact, pressure, plug)	Wi-Fi CSI (2.4 GHz, ESP32)
Output Type	Room label (5–7 classes)	Presence (binary)
Hardware Requirements	None beyond ambient sensor network	ESP32 Wi-Fi modules, Raspberry Pi
Interpretability	High (feature importance, decision trees)	Moderate (sequence deep learning)
Privacy Level	High (non-intrusive, no radio frequency or video)	Moderate (radio frequency-based sensing)
Use Case Suitability	AAL, smart homes, fine-grained context	Indoor occupancy tracking, non-intrusive

to unseen data, which is crucial for the real-world deployment of such a system. Tables 4 and 5 give a summary of the comparison, and it can be seen that our RF approach offers better accuracy while being more efficient and lightweight. The CSI-BiLSTM model uses bidirectional LSTM networks to classify the presence or absence of Wi-Fi CSI time series. It achieved 91.68% accuracy on 20% test segment of data from the overall training and testing data, and 88.67% on an unseen file containing data under different room conditions. But it supports only binary presence detection and requires ESP32 hardware, which is becoming harder to find on newer Wi-Fi chips. Our model obtained 95.00% and 91.43% accuracy on 20% test segment of the training data and the unseen file, respectively, for scenario 1. For scenario 2, it achieved 99.15% and 90.53% accuracy, respectively, for the same setting of testing and unseen files.

The HOOD system [3] uses FMCW radar for binary presence detection and sets a high benchmark with an average of 93.71% for static (standing humans) and 93.02% for very static (lying and sitting humans) activities. The constraints of special FMCW radar hardware and built-in signal processing make it harder to use in low-cost or limited-resource smart homes. In contrast, our chosen RF model performance achieved comparable AUROC while handling more granular room-level classification and dynamic activities, as can be seen in Tables 2 and 3. Details of per-class AUROC are provided in Table 6. Our model generalises well on unseen data with 91.43% accuracy and up to 0.99% AUROC, requires no additional hardware, and supports interpretable, fine-grained room-level classification, making it a good baseline for further testing for deployment in smart homes.

Table 6: AUROC per class (one-vs-rest) for Scenario 1 and Scenario 2. ”–” indicates the class was not evaluated in that scenario.

Table 5: Comparison of Proposed Work with HOOD [3]

Aspect	**Proposed Work (Ambient Sensors)**	**HOOD (FMCW Radar)**
Detection Task	Room-level multi-class classification (5–7 classes)	Room-level binary classification (presence vs. absence)
Sensor Type	Ambient sensors (motion, contact, pressure, plug)	FMCW radar
Activity Types in Dataset	Mixture of static and dynamic activities (e.g., cooking, walking, sitting, resting, watching TV, using toilet)	Static and very static activities only
Output Type	Room label per instance (multi-class)	Binary presence detection per room
AUROC (Train-Test Split)	Scenario 1: 0.9991 (macro avg) Scenario 2: 0.9998 (macro avg) *(Includes both static and dynamic activities)*	Static: 95.71 Very Static: 93.02 *(Binary classification on static activities only)*
AUROC (Unseen File)	Scenario 1: 0.9981 (macro avg) Scenario 2: 0.9984 (macro avg)	Not reported
Hardware Requirements	No additional hardware; uses existing ambient sensors	Requires FMCW radar deployment and multiple units for coverage
Interpretability	High (feature importance, decision trees from RF)	Low (deep neural networks, limited transparency)
Privacy Level	High (non-intrusive, no video or radio frequency imaging)	Moderate (senses micro-movements via radar waves)
Use Case Suitability	Highly suitable for AAL, smart homes, scalable deployments	Effective in static detection, but harder to scale due to hardware cost and setup complexity

Class	**Scenario 1**	**Scenario 2**
Bathroom	0.9992	1.0000
Bedroom	0.9998	1.0000
Corridor	0.9990	–
Diningarea	0.9952	0.9993
Kitchen	0.9985	0.9999
Sofaarea	0.9986	0.9997
Upstairs	0.9895	–

7 Conclusion and Future Work

This study shows that standard ambient sensors in smart homes can support an interpretable room-level presence detection. Our RF classifier effectively classified locations with 99.15% accuracy offline and 91.43% on unseen data, outperforming LSTM and exceeding state-of-the-art work like HOOD and CSI-BiLSTM, all without requiring additional hardware. These results confirm that our lightweight, interpretable and privacy-aware approach holds potential for real-world deployment for indoor human presence detection. However, there are still some limitations and areas for improvement. To improve generalisation, future work will explore methods like synthetic data generation or transfer learning to include all underrepresented classes ("upstairs" and "corridor" in our case). Additionally, the model needs to be tested in different homes and furniture layouts, as well as with multiple participants, to evaluate its adaptability to diverse home environments and conditions. Overall, this work contributes as a scalable and interpretable baseline for room-level presence detection in smart homes and sets the stage for more generalisable and accessible AAL solutions.

References

1. S. Rafique, F. Amirabdollahian, G. Fang, and P. Holthaus, "Human presence detection to support contextual awareness in ambient assisted living scenarios," in *2024 IEEE International Conference on Metrology for eXtended Reality, Artificial Intelligence and Neural Engineering (MetroXRAINE)*, 2024, pp. 1189–1194.
2. A. Alazeb, U. Azmat, N. Al Mudawi, A. Alshahrani, S. S. Alotaibi, N. A. Almujally, and A. Jalal, "Intelligent localization and deep human activity recognition through IoT devices," *Sensors*, vol. 23, no. 17, p. 7363, 2023.
3. S. M. Kahya, M. S. Yavuz, and E. Steinbach, "HOOD: Real-time human presence and out-of-distribution detection using FMCW radar," *IEEE Transactions on Radar Systems*, 2024.
4. G. Zhang, D. Zhang, Y. He, J. Chen, F. Zhou, and Y. Chen, "Multi-person passive WiFi indoor localization with intelligent reflecting surface," *IEEE Transactions on Wireless Communications*, 2023.
5. A. K. Panda, M. S. Manikandan, and B. Ramkumar, "Bidirectional LSTM based sequence classification approach for human presence detection using channel state information," in *2023 International Conference on Microwave, Optical, and Communication Engineering (ICMOCE)*, 2023, pp. 1–6.
6. University of Hertfordshire, "Robot House," https://robothouse.herts.ac.uk/, 2025, accessed: Jun. 12, 2025.
7. T. Guettari, J. Boudy, B. E. Benkelfat, G. Chollet, J. L. Baldinger, P. Dore, and D. Istrate, "Thermal signal analysis in smart home environment for detecting a human presence," in *2014 1st International Conference on Advanced Technologies for Signal and Image Processing (ATSIP)*, 2014, pp. 334–339.
8. S. Wilhelm, D. Jakob, and D. Ahrens, "Human presence detection by monitoring the indoor CO_2 concentration," in *Proceedings of Mensch und Computer 2020*, 2020, pp. 199–203.
9. M. Baruah and S. Suresh, "Wireless Measurement System for Real-Time Kinematic Analysis of Human Motion," in *2024 Control Instrumentation System Conference (CISCON)*. IEEE, 2024, pp. 1–6.
10. L. C. Shum, R. Faieghi, T. Borsook, T. Faruk, S. Kassam, H. Nabavi, S. Spasojevic, J. Tung, S. S. Khan, and A. Iaboni, "Indoor Location Data for Tracking Human Behaviours: A Scoping Review," *Sensors*, vol. 22, no. 3, p. 1220, 2022.
11. Y. Zhang, X. Wang, J. Wen, and X. Zhu, "WiFi-based non-contact human presence detection technology," *Scientific Reports*, vol. 14, no. 1, p. 3605, 2024.
12. S. S. Jayaweera, M. Z. Ozturk, B. Wang, and K. J. R. Liu, "Robust Deep Learning Based Residential Occupancy Detection With WiFi," in *Proceedings of the 23rd ACM Conference on Embedded Networked Sensor Systems*, 2025, pp. 662–663.
13. O. Kerdjidj, Y. Himeur, S. S. Sohail, A. Amira, F. Fadli, S. Attala, W. Mansoor, A. Copiaco, A. Gawanmeh, S. Miniaoui *et al.*, "Uncovering the potential of indoor localization: Role of deep and transfer learning," *IEEE Access*, 2024.
14. M. Haddouchi and A. Berrado, "A survey and taxonomy of methods interpreting random forest models," *arXiv preprint arXiv:2407.12759*, 2024.
15. E. Koşar and B. Barshan, "A new CNN-LSTM architecture for activity recognition employing wearable motion sensor data: Enabling diverse feature extraction," *Engineering Applications of Artificial Intelligence*, vol. 124, p. 106529, 2023.