OXFORD JOURNAL
OF STUDENT
SCHOLARSHIP

Privacy-Preserving Doorway People Counting Using Embedded PIR–Ultrasonic Sensor Fusion

Aryan Biradar
aryanbiradar111@gmail.com

ABSTRACT

Occupancy sensing is crucial for smart building automation and energy management systems. Conventional vision-based and cloud-dependent systems raise privacy concerns and often require significant computational resources. This work presents a low-cost, privacy-preserving, offline edge-computing system for people counting. The proposed system combines two passive infrared (PIR) sensors and one ultrasonic sensor integrated with an Arduino-based platform to detect directional movement through a doorway and classify events such as entry, exit, or no-count. Event-level features, including PIR activation sequence and ultrasonic blocked duration, were extracted from the sensor signals and used to train a decision-tree classifier. The learned decision logic was then deployed directly on the embedded system for real-time classification. Experimental evaluation was conducted using controlled single-subject, single-occupant doorway movement scenarios, including valid entry, valid exit, false trigger, and linger-then-retreat behaviors. The live implementation achieved an overall accuracy of 85.86% and a weighted F1-score of 86.33% under these conditions, demonstrating that reliable occupancy-related event classification can be performed without cameras, cloud services, or high-cost infrastructure. These results show the potential of lightweight sensor-fusion approaches for privacy-preserving occupancy monitoring in smart buildings and emergency-response applications.

1. INTRODUCTION

Occupancy count sensing, the process of detecting and determining the number of people within an indoor space, has the potential to significantly improve building safety and optimize emergency response strategies [1]. In critical scenarios such as fire evacuations, real-time data on occupancy density can allow first responders to prioritize search-and-rescue operations and manage rescues more effectively [2][3]. The severity of this problem becomes clear in large buildings, where poor visibility, panic, and limited time can make it extremely difficult to know how many occupants remain inside the building. Without accurate occupancy information, emergency teams may waste valuable time searching for low priority areas while occupants in more critical zones remain at risk [4]. Beyond emergencies, accurate occupancy is also important for improving energy efficiency, ventilation control, and space management in smart buildings [5][6][7][8]. Thus, developing a reliable, low cost, and privacy preserving occupancy sensing system is both a significant engineering challenge and an important real-world need [9].

In this application, passive infrared sensors and ultrasonic sensors are used to capture physical signals related to human presence such as motion heat and reflected waves, which can then be interpreted to estimate how many people are present in a room. Conventional occupancy sensing methods often rely on a single sensor type, such as passive infrared sensors, ultrasonic sensors, or radiofrequency-based systems, each of which detects occupancy in a different way and carries its own limitations [10][11]. PIR sensors detect changes in infrared radiation caused by movement, making them inexpensive and privacy-preserving but ineffective for identifying stationary occupants, while ultrasonic sensors emit sound waves and measure their reflections more continuously but often with reduced precision in complex environments [12][13]. Because raw sensor outputs are frequently noisy, limited, or ambiguous, machine learning has become increasingly important as a method of interpreting sensor data more intelligently than fixed rule-based systems [13]. Rather than depending on simple thresholds or conventional direct detection methods, machine learning models such as Decision Trees, which utilize hierarchical rule-based splits in sensor data to classify patterns, can identify patterns across multiple sensor inputs and improve the system's ability to infer occupancy under more realistic conditions [14]. Despite the clear humanitarian and logistical advantages of occupancy sensing, the transition from theoretical models to widespread real-world deployment has been hindered by a persistent technological setback. As indoor environments become increasingly more advanced, the research community has identified a significant gap between high-precision industrial tracking and the practical, non-invasive requirements of residential and professional safety [15]. Current methodologies are often prone to one of the following core setbacks:

Camera-based occupancy sensing methods, including deep learning and computer vision systems, offer high spatial resolution and strong occupant-tracking performance, but they introduce substantial privacy concerns in indoor environments [16][17]. Vision-based systems commonly detect and track faces, heads, body contours, and motion from images or video, which can expose personal identity and behavioral information in a way many occupants consider intrusive [18]. Because of these privacy risks, camera-centered approaches are often difficult to justify in residential, office, and other non-public indoor settings where non-invasive monitoring is preferred. This privacy limitation is widely recognized in reviews of vision-based occupant sensing and smart-building privacy research. Radio-frequency approaches such as those utilized in [19][20] can achieve highly accurate indoor tracking, but they often require more expensive hardware and deployment than simpler sensing alternatives. Reviews comparing indoor localization technologies note that UWB provides strong accuracy, yet this performance comes with higher hardware cost, and energy demands than systems based on Wi-Fi or BLE. Such cost barriers make these systems less practical for low-budget buildings, large-scale retrofits, or broadly deployable life-safety applications where affordability and scalability are essential design requirements. As a result, the literature continues to emphasize the necessity of lower-cost sensing solutions that preserve useful occupancy information without requiring industrial-grade infrastructure. Wi-Fi-based Channel State Information methods reduce the need for dedicated sensing hardware by leveraging existing wireless infrastructure, but this advantage also makes them dependent on the continued availability and stability of that infrastructure [20][21]. Multiple papers describe CSI-based sensing as operating through Wi-Fi access points, commodity routers, or existing local wireless networks rather than as a fully standalone sensing system. Because these methods rely on active communication hardware and powered network devices, the health of the building's wireless environment determines their robustness. In emergency conditions where local network service or power distribution is disrupted, this dependence may reduce the reliability of Wi-Fi-based occupancy estimation for real-time life-safety use. The literature therefore suggests an important tradeoff: CSI sensing is attractive for low-cost deployment, but infrastructure dependence can limit its suitability for fail-safe emergency occupancy monitoring.

Beyond the barriers of privacy, cost, and connectivity, low-cost sensor arrays face a persistent challenge in interpreting ambiguous doorway events [22]. In an emergency evacuation, individuals rarely move with a standard, predictable gait. Occupants may be carrying large objects, dragging luggage, or extending their limbs to navigate through debris or low-visibility areas [23][24]. These non-standard movements can create distorted sensor signatures that complicate the interpretation of motion events in low-cost sensing systems [22]. For a rescue team, a false positive or the incorrect interpretation of a doorway event can lead to dangerous misallocation of rescue personnel. Conversely, a false negative is equally catastrophic. Standard threshold-based algorithms are often unable to distinguish between valid crossings and irregular motion behaviors that only partially resemble occupant movement through a monitored passage [25]. This ambiguity remains a leading cause of error in existing low-cost fusion systems [26].

This study addresses the convergence of these gaps by developing a low-cost, offline, and privacy-preserving sensor-fusion system designed specifically for emergency rescue contexts. By integrating Passive Infrared and Ultrasonic sensors into a localized edge-computing platform, the system achieves a non-invasive sensing modality that respects occupancy privacy while maintaining hardware affordability [12]. The primary innovation of this research lies in the application of a Decision Tree algorithm to improve the interpretation of ambiguous doorway-event patterns. Unlike cloud-dependent AI, this lightweight logic runs locally on Arduino-based microcontrollers, ensuring that the system remains operational regardless of network status [27]. The algorithm processes the relationships between ultrasonic blockage duration and the frequency of PIR zone transitions. By analyzing these event-level sensor signatures, the system can distinguish valid entry and exit motions from non-count events such as false triggers and lingering behavior [14]. The effectiveness of this system was evaluated through a series of robustness tests simulating various emergency exit behaviors.

The remainder of this paper is organized as follows: Section 1 offers a comprehensive Introduction and literature overview on the implementation of PIR and ultrasonic sensors in indoor settings. Section 2 details the hardware configuration, the optical masking techniques employed to refine the sensor field, and the parameters of the Decision Tree algorithm. Section 3 presents the results of the robustness tests and the accuracy of the metrics of sensor data. Section 4 discusses the practical implications of these findings for privacy-preserving occupancy sensing and emergency-response support, while Section 5 concludes the paper by outlining potential avenues for future integration into smart-building safety protocols.

2. METHODOLOGY

The proposed system uses two sensing modalities with complementary functions: passive infrared sensing for motion detection and ultrasonic sensing for distance measurement. The HC-SR501 PIR sensors are used to detect the order of motion across two doorway zones rather than the exact position. Each module has three pins—VCC, GND, and OUT—and produces a digital signal that can be read directly by the Arduino. Because a PIR sensor responds to motion-induced thermal change rather than static presence, it is well suited for detecting whether a subject moves from one side of the doorway to the other. Typical HC-SR501 specifications include an operating voltage of 4.5–20 V, a 3.3 V TTL output, and an adjustable detection range of approximately 3–7 m.

The HC-SR04 ultrasonic sensor provides second sensing modality. Unlike the PIR sensor, it actively emits an acoustic pulse and measures the echo's return time to estimate distance. The HC-SR04 includes four pins—VCC, GND, TRIG, and ECHO—and typically operates at 40 kHz with a 5 V supply. Its function in the proposed system is to confirm whether an object physically entered the monitored lane and to estimate how deeply the lane was occupied during an event. Features derived from ultrasonic sensing include minimum event distance, distance drop relative to baseline, and blockage duration.

The combination of PIR and ultrasonic sensing is advantageous because the two sensors compensate for each other's limitations. PIR sensors provide directional information through activation order but do not confirm physical passage occupancy, while ultrasonic sensing provides distance-based confirmation but not directional sequence. Together, they form a compact sensor-fusion framework for doorway people counting.

2.1 Hardware Design and Configuration

Specification	Arduino Uno Rev3
Microcontroller	ATmega328P
Operating Voltage	5 V
Input Voltage	7-12 V
Digital I/O Pins	14 (6 of which provide PWM output)
Analog Input Pins	6
Flash Memory	32 KB
Clock Speed	16 MHz

Table 1. Technical specifications of the Arduino Uno Rev3 used as the central processing unit

Sensor Specifications	HC-SR04 ultrasonic sensor	Sensor Specifications	HC-SR501 PIR sensors
Detection Range	3 cm-450 cm	Detection Range	3-7 m
Frequency	40 kHz	Detection Angle	< 110°
Delay Time	N/A	Delay Time	5-200 s
Consumption	<2 mA	Consumption	<50 µA
Operating Voltage	5 V	Operating Voltage	4.5-20 V
Detection Angle	<15°	Output Voltage	3.3 V
Resolution	0.3 cm	Board Dimensions	32 mm x 24 mm

Table 2: Technical parameters of HC-SR04 ultrasonic sensor and HC-SR501 PIR sensor used for directional motion detection in the proposed doorway-monitoring system

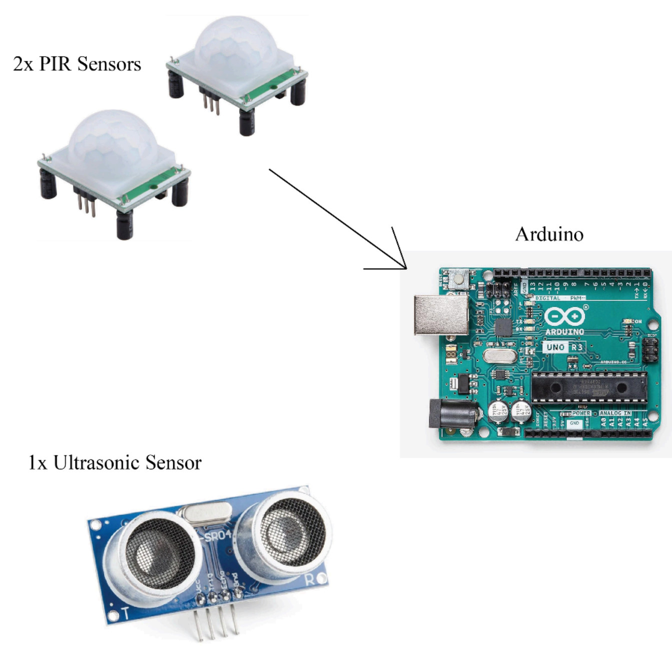

Figure 1. Integration of two HC-SR501 PIR sensors and one HC-SR04 ultrasonic sensor with the Arduino Uno

processing platform.

The proposed doorway-monitoring system was developed as a low-cost, privacy-preserving, and offline-capable sensing platform using an Arduino Uno Rev3, two HC-SR501 passive infrared (PIR) sensors, and one HC-SR04 ultrasonic sensor. The technical specifications of the Arduino, PIR sensor, and ultrasonic sensor are reported in Table 1 and Table 2, respectively, while their conceptual integration is shown in Figure 1.

The Arduino Uno serves as the central processing unit of the system. It powers the sensing components, receives their outputs, extracts event-level features, and executes the logic used to determine whether the room count should increase, decrease, or remain unchanged. All sensing and processing are performed locally on the microcontroller, allowing the system to operate without cameras, cloud services, or internet connectivity.

The two PIR sensors are used to detect motion in adjacent doorway zones. Rather than measuring exact position, they are used to identify the order in which motion occurs across the passage. Each sensor was connected to the Arduino through VCC, GND, and OUT, allowing the microcontroller to monitor both zones independently and determine the first and second activation during a doorway event. The HC-SR04 ultrasonic sensor was mounted centrally to confirm whether an object physically entered the monitored lane. Its VCC and GND pins were connected to the breadboard rails, while the TRIG and ECHO pins were wired to dedicated Arduino digital pins so that distance could be measured during each event.

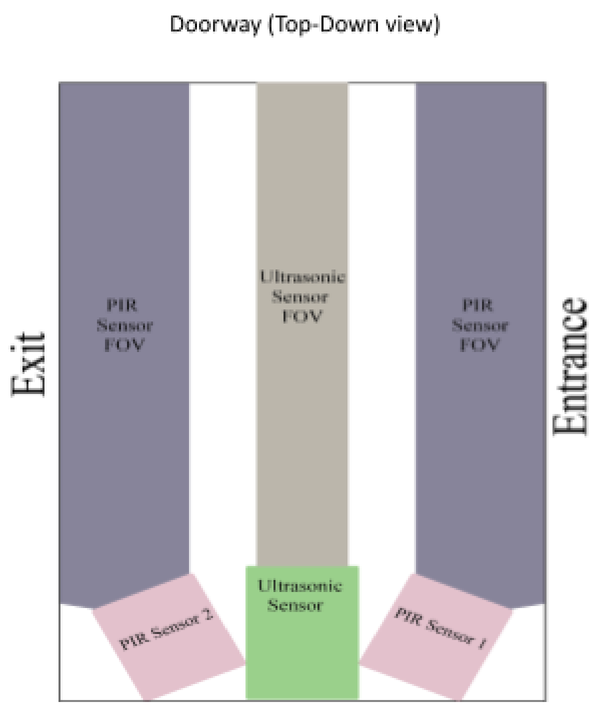

Figure 2. Top-down view of the proposed doorway sensing configuration. The grey regions represent the effective

fields of view of the two PIR sensors, while the beige region represents the sensing path of the centrally positioned ultrasonic sensor.

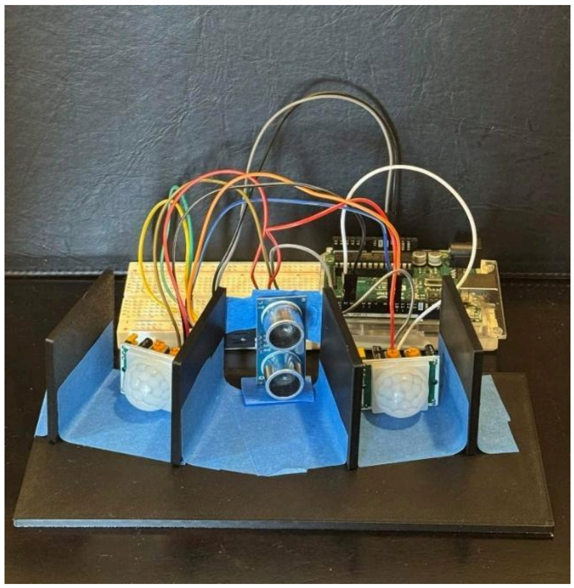

Figure 3. Photograph of the physical prototype used in the experimental setup, showing the two PIR sensors, the centrally mounted ultrasonic sensor, and the barriers used to reduce overlap between adjacent PIR sensing zones.

The physical sensing arrangement is illustrated in Figure 2 and implemented in the prototype shown in Figure 3. The system was designed around a narrow doorway lane monitored by two PIR sensors, and one centrally positioned ultrasonic sensor. The PIR sensors were placed on opposite sides of the passage and angled inward so that each one primarily covered a different portion of the doorway.
This created two distinct sensing zones and enabled motion direction to be inferred from the order of PIR activation. The ultrasonic sensor was aligned with the center of the lane so that it monitored the core passage region rather than either side zone alone.

Because unshielded PIR sensors may have overlapped fields of view, physical barriers were added to reduce ambiguity and better separate the sensing zones. As shown in Figure 2, the PIR sensors were constrained to cover adjacent regions of the doorway, while the ultrasonic sensor monitored the central lane. The physical prototype in Figure 3 shows the barriers used to achieve this configuration. During prototyping, these barriers were made from cardboard so that spacing and field-of-view restrictions could be adjusted easily. Their purpose was to reduce overlap between the PIR sensing areas and improve the consistency of the recorded event patterns. The same geometry was later intended to be transferred to 3D-printed walls without changing the sensing configuration.

The sensing rig was kept in a fixed position during data collection to preserve repeatability. Sensor orientation, spacing, and barrier geometry were maintained as consistently as possible so that differences in recorded patterns reflected movement behavior rather than changes in hardware placement. This was important because the system relies on event-level differences between valid crossings and non-count cases such as false triggers.

Overall, the hardware design combines two PIR sensors for directional motion detection and one centrally mounted ultrasonic sensor for occupancy confirmation. As shown in Figure 1, Figure 2, and Figure 3, this configuration provides a compact sensor-fusion framework in which the PIR sensors supply temporal direction information and the ultrasonic sensor provides distance-based confirmation of passage occupancy. Together, these components form the physical basis of the proposed system's ability to classify doorway events as count in, count out, or no-count.

2.2 Decision Logic Framework

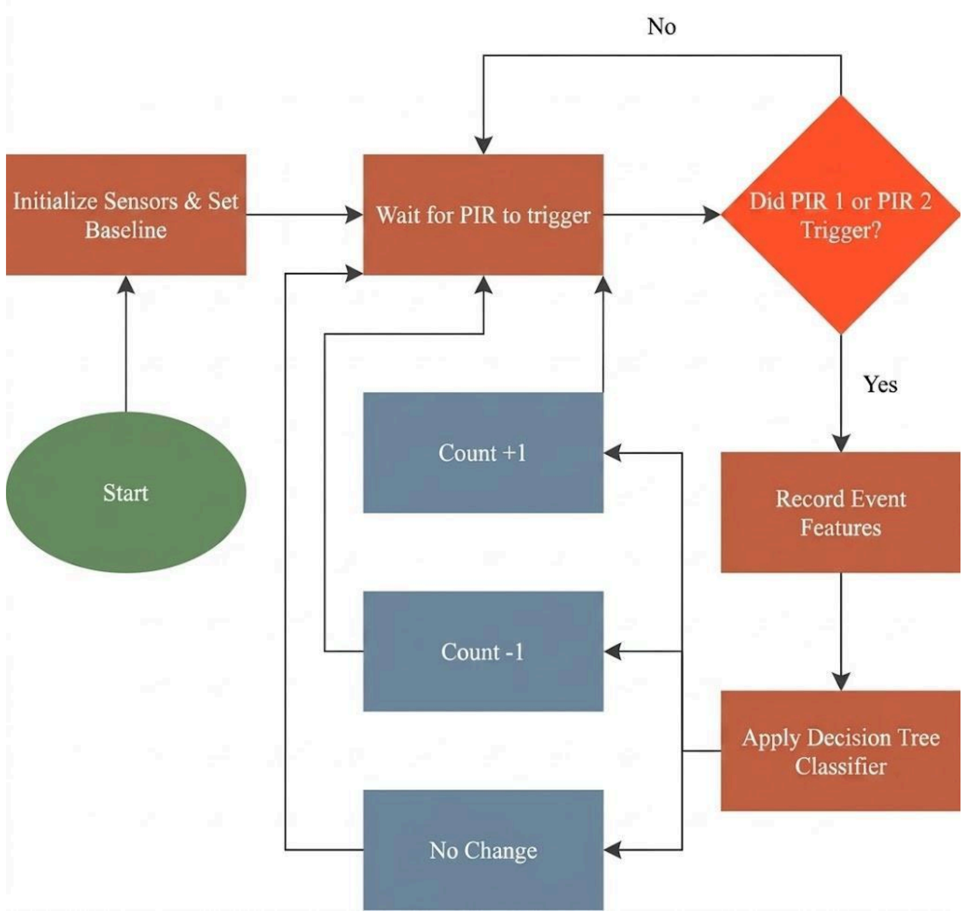

Figure 4. Flowchart of the proposed doorway-event classification algorithm used to detect motion events and determine whether the room count should increase, decrease, or remain unchanged.

The system operates in an event-driven manner: sensors are initialized with baseline ultrasonic distance measurement, then the system enters a waiting state. When motion is detected, event features are recorded (PIR activation order, time gap, active durations, ultrasonic distance measurements, and blockage duration). A decision-tree classifier maps the event to one of three outputs: count_in, count_out, or no_count. The overall workflow is illustrated in Figure 4.

At the beginning of system operation, the sensors are initialized, and a baseline ultrasonic distance is established. This baseline represents the clear-lane reference used to interpret later ultrasonic measurements. After initialization, the system enters a waiting state in which it continuously monitors the two PIR sensors. If neither PIR sensor is triggered, the algorithm loops back to the waiting state. When one PIR sensor detects motion, the system begins processing the event. During the event-recording stage, the Arduino collects the information needed to characterize the doorway motion. These event features include the order of PIR activation, the time gap between PIR detections, the active duration of each PIR sensor, the minimum ultrasonic distance recorded during the event, the distance drop relative to baseline, the duration for which the sensing lane appears blocked, and whether the lane returns to a cleared state afterward. By grouping these measurements into one structured event, the system avoids responding multiple times to a single movement. After the event's features are recorded, the algorithm applies to the decision-tree classifier shown in Figure 4. The classifier maps the event to one of three outputs: a positive count update, a negative count update, or no count change. If the event corresponds to a valid entry, the count increases by one. If it corresponds to a valid exit, the count decreases by one. If the event is incomplete, ambiguous, or not associated with a true crossing, the count remains unchanged. After the update, the system returns to the waiting state and resumes monitoring.

Directional logic: Directional logic: A PIR1-to-PIR2 activation sequence indicates valid entry (count increases by one); a PIR2-to-PIR1 sequence indicates valid exit (count decreases by one). The ultrasonic sensor confirms whether the lane was physically occupied.

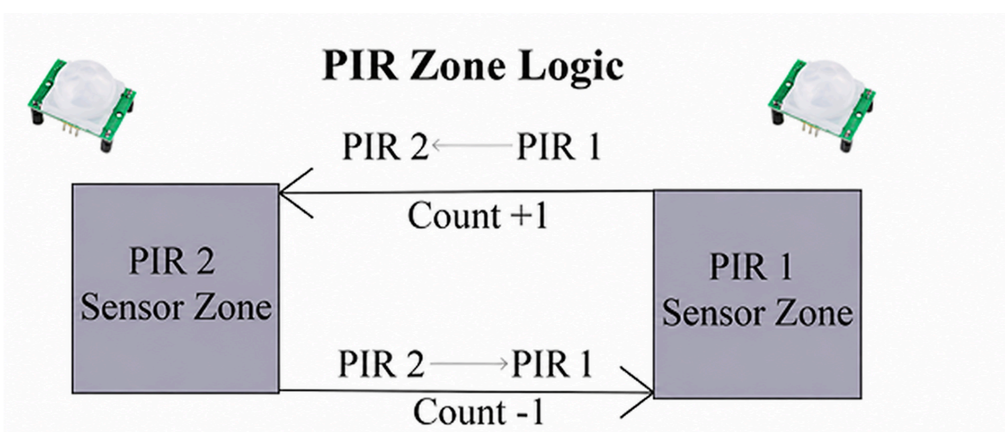

Figure 5. Directional counting logic of the proposed doorway monitoring system based on PIR activation order across the two sensing zones.

2.3 Experimental Movement Scenarios

Scenario	Description	Expected Label	Count Action
One In	Subject crosses doorway into room	count_in	+1
One Out	Subject crosses doorway from room	count_out	−1
False Trigger	Motion detected without valid crossing	no_count	0
Linger Then Retreat	Subject pauses and exits without completing crossing	no_count	0

Table 3. Experimental movement scenarios used to generate the event dataset, including the corresponding event labels and expected room-count actions.

Data collection was performed using one human subject under typical indoor daytime conditions without controlled illumination. The sensing rig remained fixed relative to the doorway. Each trial began with an unoccupied lane so baseline ultrasonic distance could be measured. The system summarized each trial as a single event rather than storing continuous raw data. Operational labels (count_in, count_out, no_count) mapped to required room-count actions.

2.4 Event Feature Extraction

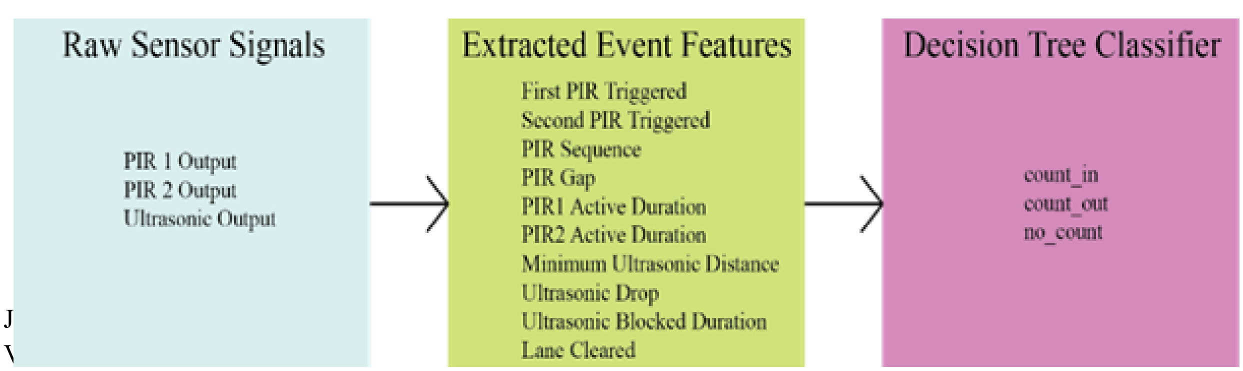

Figure 6. Event feature-extraction pipeline of the proposed system. Raw PIR and ultrasonic sensor signals are transformed into event-level features, which are then used as inputs to the decision tree classifier.

Raw sensor signals are transformed into extracted event-level features: first and second PIR triggers, PIR activation sequence, time gap between activations, active duration of each PIR, minimum ultrasonic distance, distance drop relative to baseline, blockage duration, and lane-cleared status. These engineered features improve interpretability and reduce sensitivity to low-level noise. The feature set forms the input to the decision-tree classifier.

2.5 Training and Deployment

To ensure reproducibility of the proposed classification approach, the decision-tree model was trained using the labeled doorway-event dataset collected from the sensor-fusion system. After data cleaning and removal of invalid or duplicate entries, the final dataset contained 148 labeled events. Each event was represented by extracted features derived from the PIR and ultrasonic sensors, including PIR activation order, PIR activation gap, PIR active durations, ultrasonic baseline distance, ultrasonic minimum distance, ultrasonic distance drop, and ultrasonic blocked duration. The target labels used for classification were count_in, count_out, and no_count. This 148-event dataset was used only for model development, including training and internal test evaluation, and did not include any events from the separate 99-event live evaluation dataset reported in Section 3.

Parameter	Value
Dataset Size	148 Labeled Events
Target Classes	count_in, count_out, no_count
Software Environment	Python, Google Colab
Libraries	Pandas, scikit-learn
Train-test Split	80% / 20%
Classifier Type	Decision Tree
Max Tree Depth	4
Min Samples to Split	5
Min Samples per Leaf	3

Table 4. Decision-tree training configuration and deployment settings used for model development and embedded implementation.

3. RESULTS

The live evaluation set contained 99 events: 15 count_in, 14 count_out, and 70 no_count. These 99 live trials were collected separately after model development and embedded deployment and were fully independent of the 148-event dataset used during training and internal testing. The embedded classifier achieved overall accuracy of 85.86% and weighted F1-score of 86.33%. The no_count class showed strongest performance (precision 0.9385, recall 0.8714, F1 0.9037), which is critical because it prevents false count updates from non-entry behaviors. The count_in class achieved perfect recall (1.0000), meaning all true entry events were detected, but precision was 0.6000 (F1 0.7500). The count_out class showed precision 1.0000 but recall 0.6429 (F1 0.7826). All 15 true count_in events were correctly classified; 9 of 14 count_out events and 61 of 70 no_count were correct.

Parameter	Value/Description
Loop Sampling Delay	40 ms
Maximum Event Duration	4000 ms
Quiet Time for Event Termination	600 ms
Ultrasonic Occupancy Threshold	Distance drop > 2.0 cm below baseline
Baseline Update Condition	Updated only when no event is active and both PIR sensors are inactive
Event Start Condition	Event begins when either PIR1 or PIR2 is triggered
No-Count Blocked-Duration Threshold	≤ 177 ms
Entry Threshold	PIR1 → PIR2 and blocked duration > 177 ms and ≤ 1091 ms
Linger/No-Count Threshold	PIR1 → PIR2 and blocked duration > 1091 ms
Exit Sequence Condition	PIR2 → PIR1
Default Classification for Incomplete Events	no_count

Table 5. Operating parameters and embedded decision thresholds used in the live doorway-event classification system.

As shown in Figure 7, the embedded classifier achieved an 85.86% overall accuracy, surpassing the 80% target. This confirms the viability of lightweight, embedded machine learning for offline, privacy-preserving people counting.

The class-level performance breakdown is as follows:

- **no_count:** Achieved the highest reliability (Precision: 0.9385, Recall: 0.8714, F1: 0.9037). Because this class represents the largest portion of the dataset, its strong performance is crucial for stabilizing the overall system by preventing false count updates from lingering or false triggers.
- **count_in:** Reached perfect recall (1.0000) but lower precision (0.6000), resulting in an F1-score of 0.7500. While the system successfully detected all true entries, it was overly sensitive and occasionally misclassified non-entry behaviors as entries.
- **count_out:** Showed the reverse pattern with perfect precision (1.0000) but lower recall (0.6429), yielding an F1-score of 0.7826. The classifier was highly conservative here; every predicted exit was correct, but it missed several true exit events.

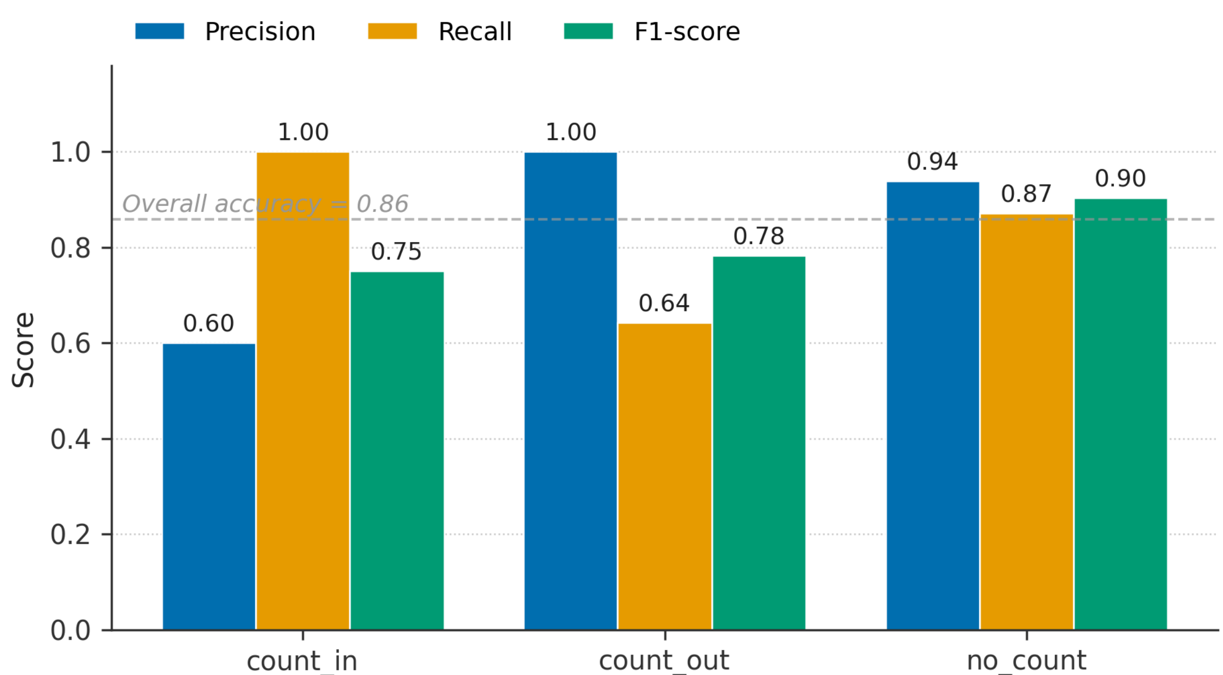

Figure 7. Performance metrics of the live embedded decision-tree classifier on the real-time doorway-event evaluation dataset.

While these metrics provide a compact quantitative summary, the specific distribution of correct and incorrect predictions is further detailed in the confusion matrix (Figure 8).

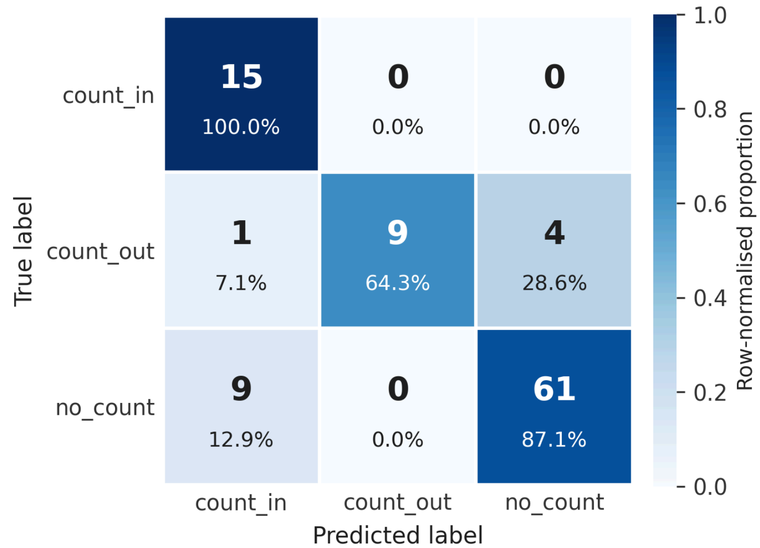

Figure 8. Confusion matrix of the live embedded decision-tree classifier evaluated on the real-time doorway-event dataset.

The main error source was confusion between no_count events and entry-like patterns. The classifier effectively separated short-duration/incomplete events into no_count while preserving directional discrimination for valid entries and exits.

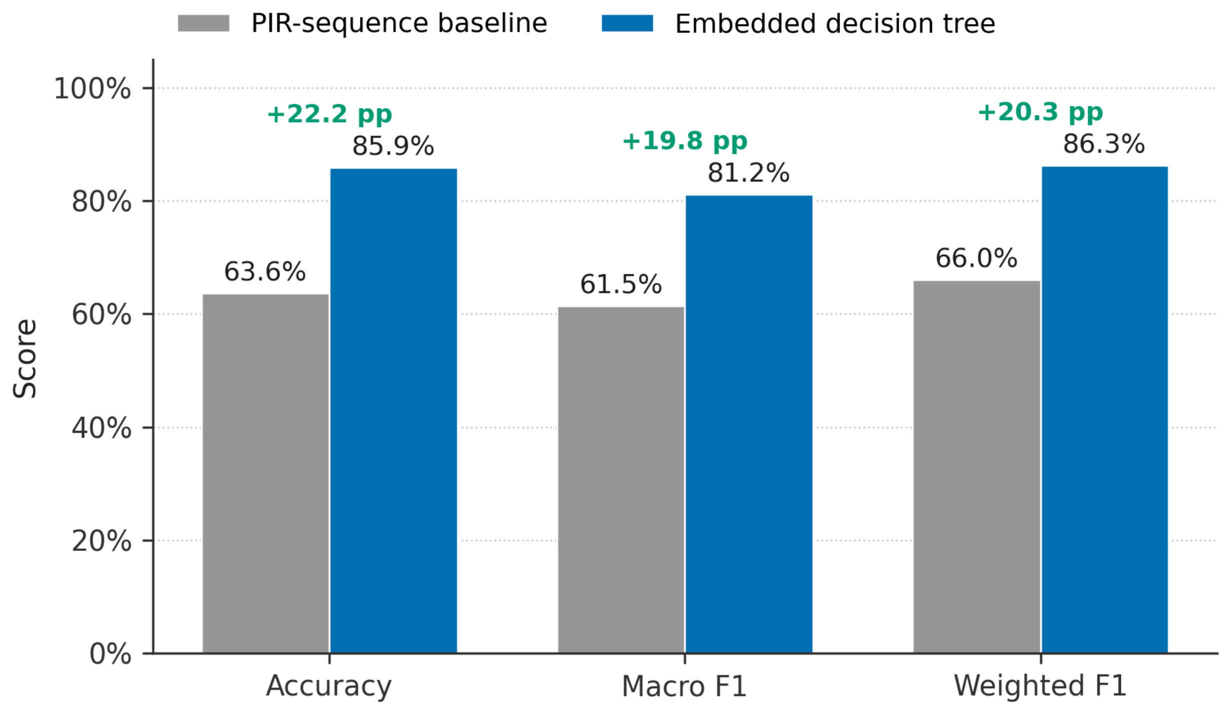

Figure 9. Performance comparison between the PIR-sequence baseline heuristic and the proposed embedded decision-tree classifier on the live dataset. Values are shown as rounded percentages for accuracy, macro F1-score, and weighted F1-score.

A simple rule-based baseline (PIR1→PIR2 = count_in, PIR2→PIR1 = count_out, others = no_count) achieved 63.64% accuracy, 61.45% macro F1, and 65.99% weighted F1 on the live evaluation dataset. By comparison, the embedded decision-tree classifier achieved 85.86% accuracy, 81.21% macro F1, and 86.33% weighted F1. Figure 9 presents these values rounded to one decimal place for visual comparison, showing that the proposed classifier substantially outperformed the simpler PIR-sequence baseline after incorporating ultrasonic blocked-duration thresholds.

Bootstrap resampling was also used to estimate 95% confidence intervals for the live classification metrics in order to quantify uncertainty associated with the finite evaluation dataset. The embedded classifier achieved an overall accuracy of 85.86%, with a 95% confidence interval of 78.79% to 91.92%. The per-class confidence intervals are summarized in Figure 10. As expected, the intervals for count_in and count_out were wider than those for no_count, since the live dataset contained fewer entry and exit samples than no-count events. This reflects the smaller support of those classes rather than instability in the classifier itself.

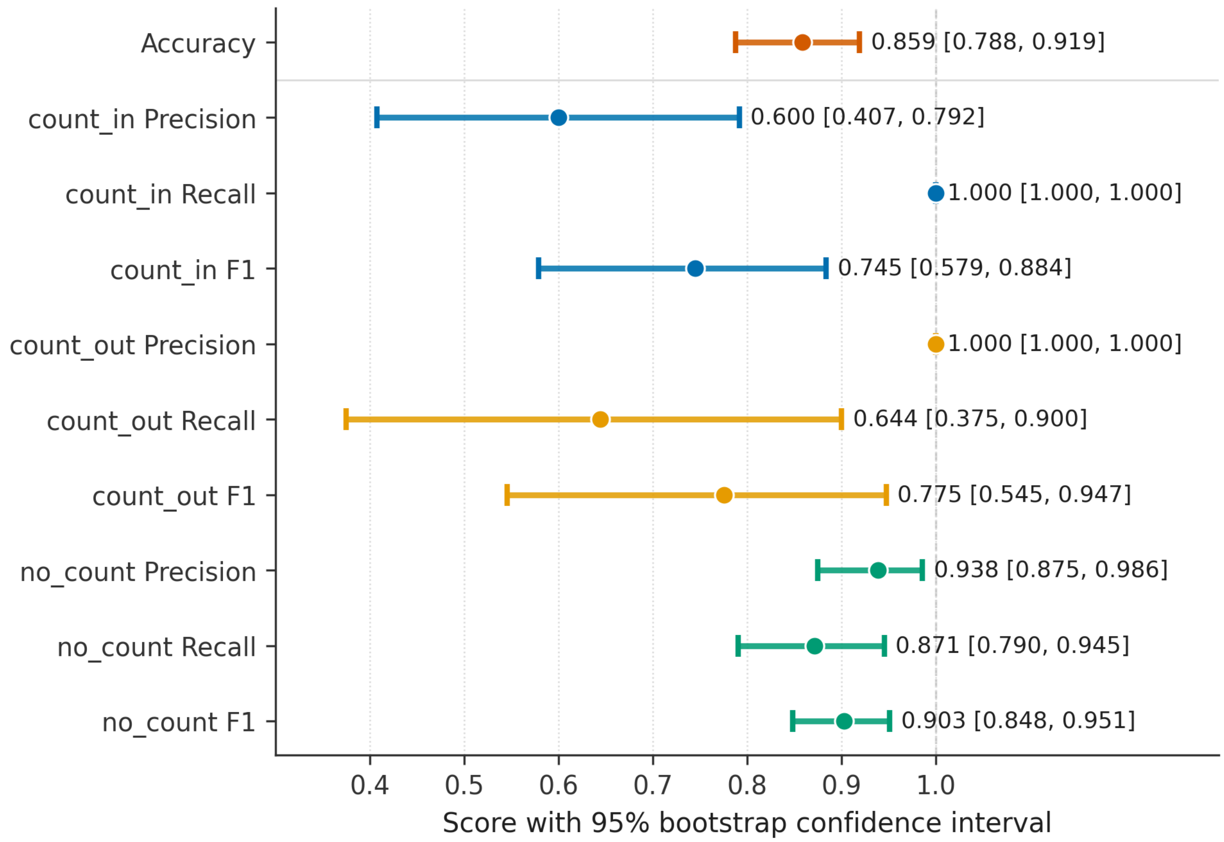

Figure 10. Bootstrap 95% confidence intervals for overall accuracy and per-class live classification metrics.

As shown in Figure 10, the no_count class maintained relatively strong and stable performance, with an F1-score of 0.9030 and a 95% confidence interval of 0.8480 to 0.9510. The count_in and count_out classes showed broader confidence ranges, particularly for precision and recall, but their intervals remained consistent with the class-level behavior already observed in Figure 7 and Figure 8. In particular, the perfect recall interval for count_in and the perfect precision interval for count_out reflect the fact that these outcomes were observed consistently across the bootstrap samples drawn from the evaluation set. Overall, the confidence-interval analysis supports the conclusion that the proposed embedded decision-tree classifier provides strong real-time performance while also making clear where additional data collection could further narrow uncertainty in future work.

4. DISCUSSION

The 85.86 % live accuracy achieved by the embedded classifier demonstrates that doorway-level occupancy sensing is feasible on a low-cost hardware platform without recourse to cameras, cloud services, or network infrastructure. A brief bill of materials for the prototype included one Arduino ($27.60), two HC-SR501 PIR sensors (approximately $5.66 prorated from a three-sensor bundle costing $8.49), one HC-SR04 ultrasonic sensor (approximately $1.80 prorated from a five-sensor bundle costing $9.00), one breadboard (approximately $3.00 prorated from a three-piece pack costing $9.00), a USB cable ($7.22), and jumper wires ($6.98). This corresponds to an approximate prototype material cost of about USD 52 including cable and wiring accessories. The no_count class which dominates normal operation and captures the false triggers and linger behaviors that would otherwise destabilize a running occupancy estimate was identified with an F1-score of 0.90, indicating that the learned tree is particularly effective at the discrimination most consequential for sustained deployment. Entry and exit events were correctly directionally classified in the large majority of cases, with the residual errors concentrated in short or partial crossings where ultrasonic blockage was brief.

The directional behavior of the classifier is well matched to the intended application. Where misclassifications occur, they bias the running count slightly upward rather than downward, which is a preferable failure mode in building-safety contexts because the system errs on the side of assuming that more occupants remain inside rather than fewer. In the live evaluation, this tendency was primarily reflected in the confusion between no_count and count_in, where 9 of 70 true no_count events were predicted as count_in, corresponding to an empirical false-positive rate of approximately 12.9% for no_count events. If the average rate of no_count events in realistic operation is denoted by X events per hour, then this error rate would correspond to an expected upward drift of approximately 0.129X spurious counts per hour in the absence of correction. Because the dominant residual error was upward bias through no_count–count_in confusion rather than bidirectional instability, reconciliation during known empty-room intervals offers a candidate mitigation: if resets occur every T hours, the expected accumulated drift between resets would be bounded to approximately 0.129XT counts. However, because the present study did not evaluate long-duration deployment or measure no_count event frequency under realistic continuous operation, these drift estimates and the adequacy of periodic resets should be interpreted as illustrative rather than fully validated, and future work should empirically characterize event rates over extended deployments.

In broader context, these results are consistent with the accuracy band reported for comparable PIR–ultrasonic systems but achieved here with a learned classifier in place of hand-coded thresholds and with the entire inference pipeline resident on the microcontroller. Together, the privacy-preserving sensing modality, the network-independent operating profile, and the interpretability of the decision-tree logic make the system a credible building block for low-cost occupancy awareness in residential, office, and emergency-response settings. Future work will extend the evaluation to multi-occupant scenarios and longer deployment windows to characterize behavior under more diverse real-world conditions.

5. CONCLUSION

This study demonstrates that occupancy sensing for emergency response and building safety can be achieved with low-cost sensors and lightweight machine learning deployed on embedded platforms. By fusing PIR activation patterns with ultrasonic blockage duration through a decision-tree classifier, the system achieved 85.86% accuracy in controlled single-subject, single-occupant real-time doorway event classification without cameras, cloud connectivity, or expensive infrastructure.

The key contribution is addressing a practical gap: time-critical occupancy estimation that is privacy-preserving, network-independent, and interpretable. In emergency contexts, reliable occupancy awareness directly impacts first-responder resource allocation. The system's transparent decision logic enables field verification and trustworthy operation in high-stakes environments.
While this evaluation focused on single-occupant controlled scenarios, real deployments will require handling multi-occupant flows, diverse gait patterns, and environmental variations. Future work should expand datasets to include realistic evacuation-like movement and explore online learning for site-specific threshold adaptation. Multiple doorway sensors could be networked to build building-wide occupancy models without centralized processing.

This work demonstrates that properly engineered low-cost sensor fusion can meet rigorous performance requirements in safety-critical contexts. In the long term, solutions of this kind could play an important role in the future of smart-building safety. As buildings become more connected and responsive, there is increasing demand for systems that can provide reliable occupancy awareness while preserving privacy and maintaining low deployment cost. A robust offline sensor-fusion system for doorway monitoring could serve as a building block for larger safety networks, supporting evacuation assistance, emergency triage, intelligent ventilation, and occupancy-aware control systems. In this sense, the present study is not only a prototype for a specific people-counting task, but also a step toward more practical, trustworthy, and scalable sensing solutions for real-world indoor environments.

REFERENCES

[1] Z. Chen, C. Jiang, and L. Xie, "Building occupancy estimation and detection: A review," *Energy and Buildings*, vol. 169, pp. 260–270, 2018.

[2] R. Tomastik, S. Narayanan, A. Banaszuk, and S. Meyn, "Model-based real-time estimation of building occupancy during emergency egress," in *Pedestrian and Evacuation Dynamics 2008*, W. Klingsch, C. Rogsch, A. Schadschneider, and M. Schreckenberg, Eds. Berlin: Springer, 2010, pp. 215–224.

[3] A. Mohammadiounotikandi, H. F. Fakhruldeen, M. N. Meqdad, B. F. Ibrahim, N. Jafari Navimipour, and M. Unal, "A fire evacuation and control system in smart buildings based on the Internet of Things

and a hybrid intelligent algorithm," *Fire*, vol. 6, no. 4, art. 171, 2023.

[4] V. Nunavath, A. Prinz, and T. Comes, "Identifying first responders information needs: Supporting search and rescue operations for fire emergency response," *International Journal of Information Systems for Crisis Response and Management*, vol. 8, no. 1, pp. 25–46, 2016.

[5] H. Saha, A. R. Florita, G. P. Henze, and S. Sarkar, "Occupancy sensing in buildings: A review of data analytics approaches," *Energy and Buildings*, vol. 188–189, pp. 278–285, 2019.

[6] P. Anand, D. Cheong, and C. Sekhar, "A review of occupancy-based building energy and IEQ controls and its future post-COVID," *Science of the Total Environment*, vol. 804, art. 150249, 2022.

[7] I. S. Walker and B. Less, "Rethinking occupancy-based ventilation controls," Lawrence Berkeley National Laboratory, Berkeley, CA, Tech. Rep. LBNL-2001199, 2019. [Online]. Available: https://eta-publications.lbl.gov/sites/default/files/osvc_paper_bne_01282019_iain_walker_lbnl201199.pdf

[8] L. Rueda, K. Agbossou, A. Cardenas, N. Henao, and S. Kelouwani, "A comprehensive review of approaches to building occupancy detection," *Building and Environment*, vol. 180, art. 106966, 2020.

[9] Q. H. Abbasi, A. Albarakati, A. Alomainy, M. Imran, A. Zahid, and J. Cooper, "Privacy-preserving non-wearable occupancy monitoring system exploiting Wi-Fi imaging for next-generation body centric communication," *Micromachines*, vol. 11, no. 4, art. 379, 2020.

[10] "Occupancy detection of an office room from light, temperature, humidity and CO_2 measurements using statistical learning models," *Energy and Buildings*, 2015. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378778815001334

[11] A. R. Mena, H. Ceballos, and J. Alvarado-Uribe, "Measuring indoor occupancy through environmental sensors: A systematic review on sensor deployment," *Sensors*, vol. 22, no. 10, art. 3770, 2022.

[12] M. R. Andersen, T. Jacobsen, and J. P. Bergmann, "Passive infrared sensor-based occupancy monitoring in smart buildings: A review of methodologies and machine learning approaches," *Sensors*, vol. 24, no. 5, art. 1533, 2024.

[13] F. Alam, R. Mehmood, I. Katib, N. N. Albogami, and A. Albeshri, "Fundamentals, algorithms, and technologies of occupancy detection for smart buildings using IoT sensors," *Sensors*, vol. 24, no. 7, art. 2123, 2024.

[14] E. Hailemariam, R. Goldstein, R. Attar, and A. Khan, "Real-time occupancy detection using decision trees with multiple sensor types," in *Proc. 2011 Symp. Simulation for Architecture and Urban Design (SimAUD)*, 2011, pp. 141–148.

[15] M. Esrafilian-Najafabadi and F. Haghighat, "Non-intrusive room occupancy prediction performance analysis using different machine learning techniques," *Energies*, vol. 15, no. 23, art.

9231, 2022.

[16] S. Salimi and A. Hammad, “Critical review and research roadmap of office building energy management based on occupancy monitoring,” *Building and Environment*, vol. 203, art. 108089, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/abs/pii/S0360132321004662

[17] A. Heydarian, E. Pellicer-Chená, B. Becerik-Gerber, M. Pollet, M. Berges, and F. Jazizadeh, “Occupant privacy perception, awareness, and preferences in smart office environments,” *Scientific Reports*, vol. 13, no. 1, art. 4073, 2023.

[18] F. Ahmad, A. Rasool, H. Naseer, and S. Athanasopoulos, “A state-of-the-art review on methodologies of occupancy estimating in buildings from 2011 to 2021,” *Electronics*, vol. 11, no. 19, art. 3173, 2022.

[19] M. Kuts, S. Sundberg, J. Kärki, and M. Rossi, “An experimental evaluation of indoor localization in autonomous mobile robots,” *ResearchGate Preprint*, 2025. [Online]. Available: https://www.researchgate.net/publication/390418965_An_Experimental_Evaluation_of_Indoor_Localization_in_Autonomous_Mobile_Robots

[20] “Wi-Fi channel state information for human sensing and indoor occupancy estimation: A review,” *Computer Networks*, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1389128625004943

[21] F. Palmese and A. E. C. Redondi, “Collecting channel state information in Wi-Fi access points for IoT forensics,” arXiv:2305.10554, 2023. [Online]. Available: https://arxiv.org/pdf/2305.10554

[22] H. Rajabi, Z. Hu, X. Ding, S. Pan, W. Du, and A. Cerpa, “MODES: Multi-sensor occupancy data-driven estimation system for smart buildings,” in *Proc. 13th ACM Int. Conf. Future Energy Systems (e-Energy)*, 2022, pp. 50–61.

[23] “Human activity and occupancy detection using deep learning and ambient sensing for smart-building applications,” *Energy and AI*, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2666165925001243

[24] “Indoor occupancy detection and people counting using low-cost sensors for emergency-response applications,” *Results in Engineering*, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/abs/pii/S2214367X24000176

[25] J. A. Sallay, “Automatic People Counting and Matching,” M.S. thesis, Utah State Univ., Logan, UT, USA, 2009. [Online]. Available: https://digitalcommons.usu.edu/cgi/viewcontent.cgi?article=1495&context=etd

[26] S. Arvidsson, M. Gullstrand, B. Sirmacek, and M. Riveiro, “Sensor fusion and convolutional neural networks for indoor occupancy prediction using multiple low-cost low-resolution heat sensor data,” *Sensors*, vol. 21, no. 4, art. 1036, 2021.

[27] M. Yousuf, H. F. Khan, M. Alharbi, and K. Aly, "Tiny machine learning and on-device inference: A survey of applications, challenges, and future directions," *Sensors*, vol. 25, no. 10, art. 3191, 2025.

[28] M. F. Martínez-Martín, R. Sanz, and A. Soriano, "A multi-sensor fusion approach based on PIR and ultrasonic sensors installed on a robot to localise people in indoor environments," *Sensors*, vol. 23, no. 15, art. 6963, 2023.