Infrastructure-Free Indoor Occupancy Estimation via Passive BLE Scanning

Venkata Srikanth Varma Datla
datla@di.uniroma1.it
Department of Computer Science, Sapienza University of Rome
Rome, Italy

Alessandro Aiuti
alessandro.aiuti@uniroma1.it
Department of Planning, Design, Technology of Architecture, Sapienza University of Rome
Rome, Italy

Alba Bisante
bisante@di.uniroma1.it
Department of Computer Science, Sapienza University of Rome
Rome, Italy

Gabriella Trasciatti
trasciatti@di.uniroma1.it
Department of Computer Science, Sapienza University of Rome
Rome, Italy

Stefano Zeppieri
zeppieri@di.uniroma1.it
Department of Computer Science, Sapienza University of Rome
Rome, Italy

Emanuele Panizzi
panizzi@di.uniroma1.it
Department of Computer Science, Sapienza University of Rome
Rome, Italy

Abstract

Accurately estimating indoor occupancy is fundamental to the development of modern smart buildings, which aim to optimize critical parameters such as Heating, Ventilation, and Air Conditioning (HVAC) control, safety, and resource management in real time to reduce energy waste. Traditional sensing approaches, including cameras, Passive Infrared (PIR) sensors, and CO_2 monitors, often encounter high deployment costs, maintenance overhead, and significant privacy concerns, particularly under General Data Protection Regulation (GDPR) regulations. This paper presents the design, implementation, and evaluation of a non-invasive occupancy estimation system that exclusively relies on Bluetooth Low Energy (BLE) scans performed via a mobile device, eliminating the need for prior structural information or dedicated sensing infrastructure. The proposed method analyzes statistical differences between various university environments, such as *Laboratories*, *Classrooms*, and *Corridors*, while integrating variables such as the number of fixed and mobile devices, the average device density per person, and the interference caused by signal bleed-through between adjacent rooms.

Based on a foundational study of device ownership behavior, we develop context-dependent calibration coefficients to address the multi-device phenomenon, in which a single occupant may carry multiple Bluetooth Low Energy emitters. Our system utilizes a three-layer architecture that includes *Passive Bluetooth scanning*, *Signal filtering with night-baseline infrastructure detection*, and *Automatic room-type classification*. This design allows for the dynamic selection of estimation parameters without the need for manual input. Field experiments conducted across various university spaces over a multi-week data collection period demonstrate that our context-aware model significantly reduces estimation error compared to traditional device-counting methods. This approach offers a scalable, cost-effective, and privacy-preserving engineering solution for real-time occupancy monitoring in smart campus environments.

CCS Concepts

• **Networks** → *Wireless access points, base stations and infrastructure*; **Mobile networks**; • **Computer systems organization** → **Sensor networks**; • **Applied computing** → *Smart cities*; • **Hardware** → *Sensor applications and deployments*; • **Computing methodologies** → Classification and regression trees.

Keywords

Indoor Occupancy Estimation, Bluetooth Low Energy (BLE), Passive Human Sensing, Infrastructure-Free Sensing, Context-Aware Modeling, Smart Buildings, Snapshot-Based Estimation

ACM Reference Format:
Venkata Srikanth Varma Datla, Alessandro Aiuti, Alba Bisante, Gabriella Trasciatti, Stefano Zeppieri, and Emanuele Panizzi. 2018. Infrastructure-Free Indoor Occupancy Estimation via Passive BLE Scanning. In *Proceedings of The 18th ACM SIGCHI Symposium on Engineering Interactive Computing Systems (EICS '26).* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3816778

1 Introduction

Accurately estimating indoor space occupancy is a cornerstone of modern smart building development, where operational efficiency depends on the building's ability to sense and react to human presence. In dynamic environments such as university campuses, real-time occupancy data is essential for optimizing energy consumption in HVAC systems and lighting, enhancing safety through improved crowd management, identifying available study spaces to improve the campus experience, and supporting long-term space planning and utilization analysis. However, occupant behavior is often misrepresented in building simulations, resulting in an "energy

EICS '26, Patras, Greece

ACM ISBN 978-1-4503-XXXX-X/2018/06
https://doi.org/10.1145/3816778

performance gap" and suboptimal space allocation. Accurate, real-time Human Occupancy Detection (HOD) is therefore fundamental to demand-driven building management.

The current study focuses on monitoring university spaces, which are dynamic environments characterized by fluctuations in student density, diverse architectural structures, and varied usage patterns. However, the proposed model is designed to be scalable and applicable to a broader range of settings, including corporate offices, co-working spaces, and healthcare facilities, as long as sufficient data is collected.

1.1 Limitations of Traditional Approaches

Traditional methods for detecting occupancy rely on various types of dedicated infrastructure. PIR sensors are widely deployed due to their low cost; however, they frequently struggle to detect stationary individuals and provide only binary presence data, rendering them unsuitable for counting applications. Additionally, PIR sensors are susceptible to false-on errors caused by environmental factors such as air turbulence. CO_2 -based sensing offers a privacy-preserving indirect measurement, but it suffers from slow response times due to the inherent lag in gas dispersion, and its accuracy is highly dependent on the specific ventilation configuration of a room. Camera-based systems can achieve high accuracy through computer vision, yet they raise serious privacy concerns related to facial recognition and biometric data retention, making them unsuitable for semi-public spaces such as university classrooms that must adhere to strict GDPR regulations specifically in Europe. Furthermore, retrofitting every classroom and laboratory with dedicated sensors can be prohibitively expensive for public universities.

These limitations have led researchers to investigate Radio Frequency (RF) analysis as a privacy-preserving and low-cost alternative to traditional sensing methods.

1.2 The BLE Passive Sensing Opportunity

BLE is a short-range wireless communication technology that transmits radio signals in the 2.4 GHz ISM band and is known for its low power consumption. Many personal electronic devices, such as smartphones, smartwatches, wireless headphones, laptops, and tablets, are equipped with BLE modules that periodically send out advertising packets containing a device identifier. These transmissions can be detected by a receiver positioned in a monitored room, allowing for the estimation of occupancy by counting the signals detected within a specific time interval.

Crucially, this approach is *passive* and requires no user cooperation, dedicated application installation, or active pairing with a receiver. Any individual carrying a BLE-enabled device automatically contributes to the sensing process, effectively shifting the complexity from the user to the system. This characteristic makes passive BLE scanning particularly attractive for environments with transient populations, such as university buildings where thousands of students and visitors pass through each day.

1.3 Engineering Challenges

While the passive sensing opportunity is compelling, estimating occupancy from raw BLE signals in complex environments such as university facilities introduces several significant challenges.

(1) **The Multi-Device Phenomenon.** In modern digital ecosystems, a single individual typically carries multiple BLE emitters. A student in a computer laboratory may have a smartphone, laptop, wireless mouse, and headphones all actively advertising, while the same student attending a lecture may only have a smartphone powered on. Previous approaches often employed a static "device-to-person" coefficient, such as assuming 1.2 devices per person, which fails to account for the varying contexts of different room types.
(2) **Fixed Infrastructure Devices.** University environments contain numerous "always-on" BLE emitters, including desktop computers, wireless keyboards, smart projectors, and IoT peripherals, that do not correlate with human occupancy. Our initial measurements revealed that a single laboratory can contain up to 15 fixed devices, creating a permanent "phantom occupancy" offset that must be identified and subtracted.
(3) **Inter-Room Signal Interference.** BLE signals in the 2.4 GHz band can penetrate thin walls and propagate into adjacent spaces. During our data collection, we observed a direct correlation between device-count peaks in a corridor and those in an adjacent classroom, confirming that BLE signals pass through the walls of the entrance area. This spatial bleed-through introduces noise that inflates the device count for any given room.
(4) **Temporal Variability.** Occupancy patterns fluctuate significantly throughout the day and across days of the week. Laboratories exhibit peak attendance around 11:00 and 15:00, coinciding with teaching and research activities, while corridors experience rapid fluctuations tied to class transition periods. An estimation model must account for these temporal dynamics rather than assuming static conditions.

1.4 Proposed Solution

In this paper, we present the design, implementation, and evaluation of a privacy-preserving, infrastructure-free occupancy estimation system that effectively addresses the identified challenges. The proposed solution relies on a non-invasive estimation model that utilizes exclusively BLE scans performed via a mobile device, thereby eliminating the need for prior structural information or expensive dedicated sensors.

From a privacy and GDPR perspective, the system processes BLE identifiers in a transient manner. Each identifier is used solely during a single scan window for de-duplication purposes and is discarded immediately afterward. Modern BLE devices often rotate their advertising identifiers using resolvable or non-resolvable private addresses, which further prevents long-term linkability. In line with GDPR principles of data minimization and privacy by design, the system stores only aggregated, non-identifying statistics, such as device counts per window and per room type. It never retains MAC addresses or any persistent identifiers that could be linked to individuals.

We categorize university rooms into three environment types: *Classrooms*, *Laboratories*, and *Corridors*. We analyze the statistical differences among these types by integrating variables such as the number of fixed and mobile devices, the average device density per person, and the interference caused by proximity between adjacent rooms. Grounded in a formative study of device ownership behavior ($N = 109$), we derive context-dependent calibration coefficients that accurately reflect the empirically observed gap between device ownership and active BLE detectability.

The system features a modular three-layer architecture:

- A **Data Acquisition Layer** that passively scans for BLE advertising packets using a standard smartphone, applying an RSSI (Received Signal Strength Indicator) threshold to limit detection to the physical boundaries of the target room.
- A **Signal Processing Layer** that identifies and filters fixed infrastructure devices through a night-baseline calibration procedure, scanning between 21:00 and 06:00 when the building is known to be empty.
- A **Semantic Layer** that automatically infers the room type based on the number of detected devices and the scan time, dynamically selecting the appropriate estimation parameters without requiring manual input of the room category.

The estimation model generates a probabilistic range rather than a single deterministic count, incorporating a calibrated margin of error to acknowledge the inherent uncertainty of RF-based sensing. An automatic calibration pipeline continuously integrates new data, ensuring the model remains accurate over time as the usage patterns of university environments evolve.

1.5 Contributions

The main contributions of this work are:

(1) A **formative study** ($N = 109$) characterising the distribution of BLE-emitting devices carried by university occupants. This study reveals context-dependent behavioural profiles that challenge the validity of static device-to-person ratios.
(2) The **design and implementation** of a modular, three-layer occupancy estimation architecture that operates on commercial off-the-shelf smartphones without requiring dedicated sensing infrastructure.
(3) A **night-baseline calibration method** for automatically identifying and filtering fixed infrastructure devices, which reduces phantom occupancy offsets.
(4) An **automatic room-type classification mechanism** that dynamically infers the environment category and selects context-appropriate estimation coefficients.
(5) A **multi-week field evaluation** across various university environments, including laboratories, classrooms, and corridors. This evaluation compares the proposed context-aware model against naive and static baselines.

This work contributes to the field of Engineering Interactive Computing Systems (EICS) by developing a reusable service-layer module for interactive buildings. The system is designed to operate on standard off-the-shelf devices and provides straightforward output, including room type and occupancy range. This output can be utilized by various interactive systems, such as demand-driven HVAC control, room availability displays, adaptive signage, and crowd-aware campus services. The semantic layer functions as an abstraction for user interactions, aligning with EICS's focus on engineering interactive computing systems.

The remainder of this paper is organized as follows: Section 2 reviews related work on occupancy sensing, positioning the present contribution in relation to existing BLE-based approaches. Section 3 presents the formative study on device ownership. Section 4 details the system architecture. Section 5 formalizes the estimation model. Section 6 reports the experimental evaluation. Section 7 discusses findings, limitations, and potential directions for future work and concludes the paper.

2 Related Works

Indoor occupancy detection and estimation span a wide range of sensing modalities, deployment assumptions, and modeling strategies. Existing approaches include environmental sensors, vision systems, Wi-Fi/CSI (Channel State Information) and RF-based methods, and BLE solutions, each offering different trade-offs in terms of cost, scalability, privacy, and infrastructure requirements. Among these, passive BLE and RF techniques have emerged as cost-effective and privacy-conscious alternatives to camera-based or sensor-heavy systems, particularly in settings where minimizing infrastructure and protecting user identity are essential. This section reviews state-of-the-art research across beaconed and infrastructure-free BLE sensing, Modeling strategies, Wi-Fi/CSI, and RF-based approaches, and environmental and vision-based methods, with particular attention to challenges such as multi-device ownership, fixed emitters, inter-room signal interference, and temporal variability.

2.1 Environmental and Vision-Based Sensing

Environmental and vision-based sensing approaches are among the most established solutions for indoor occupancy detection. Surveys such as Elkhoukhi et al. [9] provide a broad overview of sensing technologies used in smart buildings, comparing modalities such as PIR, CO_2, and camera systems in terms of cost, latency, and deployment complexity. Their review highlights how environmental sensors can offer continuous monitoring but often require carefully placed hardware and may struggle with fine-grained or real-time responsiveness.

Chaudhari et al. [3] examine IoT-based occupancy detection from a technological and algorithmic perspective, discussing how environmental sensors integrate into building automation systems. Their work emphasizes the diversity of sensing modalities and the need for algorithmic strategies that balance accuracy with deployment constraints. Zeleny et al. [25] focus specifically on room-level occupancy detection and analyze how environmental and signal-based methods can be combined to improve reliability in smart building contexts.

Vision-based systems provide high-resolution occupancy information but raise significant privacy and regulatory concerns. Shahbazian and Trubitsyna [22] survey RF-based and camera-based human sensing, noting that while vision systems can provide detailed behavioral insights, they also raise substantial biometric data and privacy risks. Their analysis underscores why camera-based solutions are often unsuitable for semi-public environments governed by strict data-protection requirements.

In contrast to environmental and vision-based sensing, our proposed system avoids dedicated hardware and camera-based monitoring entirely. By relying solely on ambient BLE advertisements from personal devices, it prioritizes privacy, low deployment cost, and ease of installation, trading fine-grained visual fidelity for a lightweight, infrastructure-free approach better suited to everyday university spaces.

2.2 Wi-Fi, CSI, and RF-Based Sensing

Wi-Fi, CSI, and RF-based sensing approaches represent a major line of research in device-free and device-assisted occupancy detection. Shahbazian and Trubitsyna [22] provide a comprehensive survey of RF-based human sensing, covering both occupancy and activity detection. Their review highlights how RF signals can capture human presence without requiring users to carry devices, but also emphasizes the complexity of RF propagation and the privacy considerations associated with RF-based monitoring.

Deep learning enhanced RF sensing has been explored in several works. Iannizzotto et al. [12] investigate how deep learning can improve passive BLE-based human sensing, focusing on robustness to noise and variability in BLE signal patterns. In a related contribution, Iannizzotto et al. [13] provide a broader perspective on passive human sensing with Bluetooth, discussing the challenges of using BLE signals for occupancy and activity inference and outlining the potential of machine learning to address these limitations.

Hybrid Wi-Fi/BLE and RSSI fingerprinting approaches have also been proposed. Chen et al. [4] explore the collection of high-resolution building occupancy data using Wi-Fi and BLE networks, demonstrating how combining multiple wireless modalities can improve spatial resolution. García-Paterna et al. [11] conduct an empirical study of room-level localization using BLE beacons, showing how BLE RSSI patterns can support indoor positioning. Janczak et al. [14] examine Bluetooth RSSI mean estimators for indoor positioning in evacuation supervision systems, highlighting the potential of BLE for safety-critical applications.

Device-free Wi-Fi RSSI and RF crowd-counting methods further demonstrate the potential of RF signals for occupancy estimation. Zou et al. [26] show that Wi-Fi-enabled IoT devices can be used for device-free occupancy detection and crowd counting, leveraging ambient RF variations caused by human presence. Pratama et al. [20] investigate multi-user, low-intrusive occupancy detection, noting that while RF-based systems can avoid requiring users to carry devices, they may suffer from drift or missed events in transition-based counting.

More recent work explores deep-learning-based CSI processing. Park et al. [19] propose a deep residual CNN for device-free occupancy detection using Wi-Fi CSI, demonstrating how fine-grained channel information can be exploited for high-accuracy sensing. However, CSI-based approaches typically require specialized hardware or multi-antenna access points, limiting their deployability in everyday environments.

Unlike CSI and RF systems, which often depend on specialized hardware, multi-antenna Wi-Fi access points, or dense infrastructure, the present work targets deployments where only periodic BLE scans from a single smartphone are available. This trades spatial resolution for low cost, simplicity, and infrastructure-free operation, making the system suitable for real-world university environments without requiring additional hardware.

2.3 BLE Beaconed Deployments and Beacon-Based Counting

A substantial body of research investigates BLE-based occupancy detection using dedicated beacons or installed receivers. These systems typically rely on fixed BLE transmitters or scanning infrastructure, which simplifies signal collection and improves spatial separation but requires installation and maintenance effort.

Demrozi et al. [8] and Demrozi et al. [7] examine low-cost BLE-based systems for distance estimation, occupancy detection, and counting. Their work demonstrates how dedicated BLE devices can be deployed to capture advertising packets in a controlled manner, enabling more stable signal acquisition and improving the reliability of occupancy estimates. Because the beacons are installed intentionally, these systems can better manage spatial coverage and signal consistency compared to opportunistic sensing.

An empirical study of room-level localization using BLE beacons was coinducted by García-Paterna et al. [11]. Their work highlights how beacon placement and RSSI characteristics can be leveraged to distinguish between adjacent rooms, illustrating the potential of BLE infrastructure for fine-grained indoor positioning. This line of research shows that beacon-based systems can achieve high spatial resolution when the environment is instrumented appropriately.

Mateos Sanchez et al. [17] propose a BLE-based tool for real-time occupancy monitoring in indoor and outdoor spaces. Their system integrates BLE data with open datasets to publish occupancy information in real time, demonstrating how BLE infrastructures can support large-scale monitoring services. This work emphasizes the value of BLE beacons for continuous, real-time occupancy analytics.

The study done by Tekler et al. [24] investigates scalable BLE deployments for identifying occupancy patterns and behavioral profiles in office environments. Their study shows how BLE infrastructure can be used not only for instantaneous occupancy detection but also for long-term pattern analysis, highlighting the versatility of beacon-based systems in workplace analytics.

Stoian et al. [23] explore augmenting BLE fingerprinting using instantaneous frequency. Their work focuses on improving the robustness and accuracy of BLE-based localization techniques by incorporating additional signal features beyond RSSI. This demonstrates ongoing efforts to enhance the precision of BLE infrastructure systems through advanced signal processing.

An intelligent Bluetooth virtual door system for occupancy monitoring was proposed by Koksal et al. [15]. Their approach uses BLE beacons to detect entry and exit events, enabling doorway-level monitoring that can support room-level occupancy estimation. This work illustrates how BLE infrastructure can be adapted for practical, service-oriented applications such as access monitoring and automated occupancy tracking.

Beacon-based systems benefit from controlled deployments and stable signal sources, but they require dedicated hardware installation and maintenance. In contrast, our work is *infrastructure-free*, relying solely on ambient BLE advertisements from devices people already carry. This eliminates installation overhead but introduces challenges such as fixed emitters, multi-device ownership, and inter-room interference, which the system addresses through automated night-baseline filtering and a semantic layer that adapts estimation parameters to room context.

2.4 Mobile and Infrastructure-Free BLE Sensing

Mobile and infrastructure-free BLE sensing approaches rely on smartphones or other user-carried devices rather than installed beacons or fixed receivers. These methods aim to exploit ambient BLE advertisements or smartphone scanning capabilities to infer occupancy without requiring dedicated infrastructure.

The work of Filippoupolitis et al. [10] investigates BLE-based occupancy detection in emergency management scenarios. Their work demonstrates how BLE signals emitted by personal devices can be leveraged to detect the presence of individuals during emergencies, showing the feasibility of smartphone-centric BLE sensing in dynamic and safety-critical environments.

A smartphone-based multi-floor indoor positioning system for occupancy detection was explored by Mashuk et al. [16]. Their study highlights how smartphones can be used to collect BLE signals across multiple floors of a building, illustrating the potential of mobile BLE sensing for vertical localization and multi-level occupancy estimation.

Putra et al. [21] compare the energy consumption of Wi-Fi and BLE in smart building communication. Their findings emphasize the energy-efficiency advantages of BLE over Wi-Fi, which is particularly relevant for mobile deployments where battery consumption is a critical constraint. This comparison underscores why BLE is often preferred for continuous or periodic scanning on smartphones.

A BLE-based occupancy detection method for buildings was proposed by Park et al. [18], demonstrating how BLE signals can be used to infer occupancy levels without requiring dedicated infrastructure. Their work shows that BLE-based sensing can be integrated into building performance analysis workflows, reinforcing its applicability for occupancy monitoring.

Recent workshop contributions by Datla et al. [5, 6] explore smartphone BLE beaconing and context-aware passenger detection in mobility and interactive settings. These studies highlight how smartphones can act not only as scanners but also as BLE advertisers, enabling richer interaction patterns and context-aware detection in mobile environments. Their work also emphasizes practical deployment considerations and the role of BLE in user-centered applications.

Both Mashuk et al. [16] and Putra et al. [21] further highlight the trade-offs associated with mobile BLE sensing, including scanning modes, energy consumption, and practical deployment constraints. These insights reinforce the importance of lightweight, energy-efficient scanning strategies for smartphone-based occupancy detection.

Our proposed system builds on mobile and infrastructure-free BLE sensing by introducing three key extensions: (i) an empirical device-ownership study (N=109) to ground device-to-person ratios in real-world behavior, (ii) an automated night-baseline procedure to identify and subtract fixed emitters, and (iii) a semantic layer that infers room type and selects context-appropriate estimation coefficients. These additions address challenges such as multi-device ownership, fixed-device interference, and environmental variability, enabling more robust occupancy estimation from periodic smartphone BLE scans.

2.5 Modeling Strategies for Occupancy Estimation

Occupancy estimation methods differ widely in how they model temporal dynamics, signal variability, and device behavior. A major distinction in the literature is between transition-based counting approaches and snapshot-based estimation. Transition-based systems attempt to detect entry and exit events over time, but as Pratama et al. [20] note, these methods can suffer from cumulative drift and missed transitions, especially in multi-user environments. Their work on low-intrusive multi-user occupancy detection highlights the challenges of maintaining accuracy when events are inferred sequentially rather than from independent observations.

Snapshot-based approaches avoid drift by treating each scan window as an independent inference problem. Tekler et al. [24] demonstrate how BLE-based systems can identify occupancy patterns and behavioral profiles using periodic snapshots of BLE activity in office spaces. Their work shows that snapshot-based modeling can capture long-term occupancy trends without requiring continuous tracking. Similarly, Chen et al. [4] explore the collection of high-resolution occupancy data using Wi-Fi and BLE networks, illustrating how snapshot-based measurements can support building-scale monitoring when combined with appropriate calibration.

Machine learning and deep learning have been applied to RF and BLE features to improve robustness against noise and environmental variability. Iannizzotto et al. [13] provide a perspective on passive human sensing with Bluetooth, outlining the potential of machine learning to address the inherent instability of BLE signals. Building on this, Iannizzotto et al. [12] investigate deep learning techniques to enhance BLE-based passive human sensing, showing how learned models can better capture complex signal patterns than handcrafted heuristics.

Reinforcement learning has also been explored as a modeling strategy. Billah et al. [1] propose a reinforcement learning approach for RF-based indoor occupancy detection, demonstrating how adaptive models can learn optimal sensing or inference strategies from experience. Their work highlights the potential of RL to handle dynamic environments where signal conditions change over time.

Deep-learning-based CSI processing represents another direction. Park et al. [19] apply a deep residual CNN to Wi-Fi CSI for device-free occupancy detection, showing that fine-grained channel information can be exploited for high-accuracy inference. However, CSI-based approaches typically require specialized hardware or multi-antenna access points, limiting their applicability in everyday deployments.

Our system adopts a lightweight, explainable snapshot-based modeling strategy that uses context-dependent device-to-person coefficients and uncertainty-aware ranges. Unlike CSI or deep-learning based methods, it does not require specialized hardware or computationally intensive models. Instead, it relies on calibrated BLE snapshots, automated baseline filtering, and semantic context inference to produce robust occupancy estimates from a single smartphone scan.

2.6 Practical Deployment, Privacy, and Engineering Studies

Practical deployment studies examine how wireless sensing systems behave under real-world constraints, including energy consumption, scalability, data publication, and integration with building services. Putra et al. [21] compare the energy consumption of Wi-Fi and BLE in smart building communication, showing that BLE offers significantly lower energy usage than Wi-Fi. Their findings highlight why BLE is often preferred for continuous or periodic scanning on mobile devices, where battery constraints are critical. This comparison is particularly relevant for smartphone-based occupancy sensing, where energy efficiency directly affects usability and deployment feasibility.

Mateos Sanchez et al. [17] propose a BLE-based tool for calculating occupancy in indoor and outdoor spaces and publishing results in real time using open data. Their system demonstrates how BLE infrastructures can support continuous occupancy monitoring at scale, and how real-time data pipelines can be integrated into public-facing services. Tekler et al. [24] similarly investigate scalable BLE deployments for identifying occupancy patterns and behavioral profiles in office environments. Their work emphasizes the importance of long-term monitoring and the ability of BLE infrastructures to capture temporal patterns beyond instantaneous occupancy.

Privacy considerations are a recurring theme in RF and BLE sensing. Shahbazian and Trubitsyna [22] highlight privacy and GDPR-related concerns in RF-based human sensing, noting that RF identifiers can be considered personal data and must therefore be handled carefully. Chen et al. [4] also discuss privacy implications when collecting high-resolution occupancy data using Wi-Fi and BLE networks, recommending the use of transient identifiers and aggregated outputs to reduce privacy risks. These studies underscore the need for privacy-preserving design choices in wireless occupancy systems.

Engineering and UX-focused contributions further explore how BLE sensing integrates with interactive systems. Bisante et al. [2] examine the use of artificial intelligence in service-oriented applications such as car-parking systems, illustrating how sensing pipelines can be embedded into interactive user experiences. García-Paterna et al. [11] evaluate a room-level localization system based on BLE beacons, demonstrating how BLE sensing can be incorporated into building services that require spatial awareness and user interaction.

Our proposed system explicitly targets integration as an interaction-facing service. It produces uncertainty-aware occupancy ranges and inferred room types that can be consumed by downstream applications such as dashboards or building automation systems. Furthermore, it adopts transient identifier handling and aggregated storage to align with privacy best practices identified in prior work, while maintaining a lightweight, infrastructure-free design suitable for real-world deployments.

Prior research demonstrates a wide spectrum of sensing and modeling strategies for indoor occupancy detection, each shaped by different assumptions about hardware availability, privacy, and deployment complexity. While environmental sensors, vision systems, Wi-Fi/CSI, and beaconed BLE infrastructures can offer high accuracy, they often require dedicated hardware or raise privacy concerns. Mobile and infrastructure-free BLE approaches reduce these barriers but face challenges such as multi-device ownership, fixed emitters, and signal variability. Our system positions itself within this space by leveraging ambient BLE signals from everyday devices and combining them with baseline filtering and semantic context inference, offering a lightweight, privacy-preserving, and deployable alternative to more infrastructure-heavy solutions.

3 Formative Study: Device Ownership

To engineer an occupancy estimation model that goes beyond naïve packet counting, we first characterized the *digital footprint* of the population in our target deployment context (university buildings). In passive BLE sensing, a fundamental source of error is the gap between (i) the number of BLE-capable devices that people *own/carry* and (ii) the subset of those devices that are *actively detectable* during typical activities (e.g., a laptop in sleep mode may stop advertising while a smartphone continues broadcasting). This distinction directly challenges the validity of adopting a single, static device-to-person ratio across heterogeneous settings.

Survey design and population. We conducted a digital survey with $N = 109$ participants from the university, targeting the primary campus demographics (undergraduate students, researchers, and administrative staff). The survey asked respondents to inventory the electronic devices they carry and to indicate which of them have Bluetooth capability, including *smartphones, smartwatches/fitness trackers, wireless headphones/earbuds, laptops, and tablets.* Beyond device ownership, we also collected information about usage patterns relevant to BLE detectability (e.g., whether Bluetooth is typically enabled, and whether devices are actively used or stowed during lectures and study sessions).

Ownership vs detectability. Results show substantial variability in the number of potential BLE emitters per person. Smartphone ownership is nearly universal (reported at 99%), and many participants reported additional devices such as laptops/tablets (reported at approximately 85%) and wearable/audio devices (reported at approximately 60%). In principle, a single person may therefore account for multiple distinct BLE identities (e.g., phone, watch, laptop, headphones). However, the study highlights that not all owned devices are simultaneously detectable: devices may be powered off, in sleep mode, or otherwise not advertising at scan time. This motivates treating device ownership as an upper bound and explicitly calibrating for *active BLE detectability.*

Context-dependent behavioral profiles and calibration implications. The survey findings support the need for context-dependent *device-per-person* coefficients k, rather than a universal constant. Two behavioral profiles emerge as practically relevant for modeling: (i) a *low-activity* profile where occupants mainly carry and use smartphones while other devices remain inactive, and (ii) a *high-activity* profile where laptops and peripherals are actively used (e.g., during work sessions), increasing the number of simultaneously detectable BLE emitters. In the updated model, these patterns are reflected by adopting a lower k (approximately $k \approx 1.3$) for low-activity contexts and a higher k (often exceeding $k > 2.5$) for high-activity contexts, which motivates the later design choice of dynamically selecting k based on inferred environment characteristics.

Motivation-based device counts. To complement the above ownership characterization with an activity-oriented view, the questionnaire study also asked respondents *why* they were in the building (three options: *attending/following a lecture*, *attending a laboratory*, or *just passing through*) and *how many electronic devices* they had with them at that moment. This framing helps connect device-carrying behavior to typical room usage and supports the idea that different campus contexts can systematically differ in observable device density.

The measured mean number of declared devices per person differs across motivations: participants reported an average of 2.88 devices/person when attending a laboratory, 3.09 devices/person when just passing through, and 2.15 devices/person when attending/following a lecture. : In addition, the distribution of motivations indicates that the largest fraction of respondents reported being in the building to attend a laboratory, followed by attending a lecture, and then passing through. Together, these results reinforce that (a) device counts are not constant across activities and (b) calibration must allow for context variation rather than relying on a single global correction factor.

Takeaway for system design. The formative study establishes two practical requirements for passive BLE occupancy estimation in university environments: (1) the mapping from detected devices to people must account for multi-device behavior that varies by activity/context, and (2) calibration should be grounded in empirically observed device ownership and usage rather than adopting arbitrary static ratios. These findings directly motivate a context-dependent estimation strategy in which k is selected to reflect the expected active-emitter density associated with different campus usage patterns.

4 System Architecture

The proposed occupancy estimation framework is engineered as a modular software system that operates on commercial off-the-shelf (COTS) mobile hardware and converts raw BLE advertising traffic into a human-readable occupancy estimate. The design follows a three-layer architecture: (i) a **Data Acquisition Layer** that performs passive BLE scanning and aggregates observations within fixed scan windows, (ii) a **Signal Processing Layer** that reduces noise and removes persistent non-human emitters via filtering and night-baseline calibration, and (iii) a **Semantic Layer** that infers contextual properties of the environment and produces an occupancy estimate with an uncertainty-aware range.

4.1 Data Acquisition Layer: Passive BLE Scanning

The system passively monitors BLE advertising packets using a standard smartphone in *Central* role, listening for *Peripheral* devices that broadcast generic advertisement payloads. Unlike active Bluetooth use (e.g., pairing), the proposed system does not establish connections; it only captures broadcast signals present in radio proximity. In the reference implementation, the acquisition module is built on iOS using the `CoreBluetooth` framework in order to obtain consistent RSSI readings from the platform's standardized BLE stack.

Each scan runs for a configurable window length (used as $W = 30$ seconds in the reference implementation), balancing temporal resolution and battery usage. During a window, the system aggregates unique identifiers and their received signal strength values into a raw observation set:

$$S_{\text{raw}} = \{(id_1, rssi_1), (id_2, rssi_2), \ldots, (id_n, rssi_n)\}. \tag{1}$$

Here, S_{raw} represents all devices detected within the RF vicinity of the scanning node for that window.

Acquisition modes and data logging. To support both calibration and long-term deployment, the mobile application supports (i) a *single-scan mode* used during initial ground-truth sampling, and (ii) a *continuous/periodic scanning mode* designed for autonomous collection at scale. In the continuous mode, the application performs an automatic scan every 15 minutes (96 scans per day) and aggregates results into a daily CSV file. The daily CSV can be transmitted automatically via a Telegram bot to a dedicated channel (or stored locally if offline), enabling longitudinal analysis without requiring continuous user intervention.

4.2 Signal Processing Layer: Filtering and Baseline Calibration

Raw BLE detections are noisy due to multipath propagation, penetration through thin walls, and environmental dynamics. To isolate signals likely originating inside the target room and to mitigate systematic bias from non-human emitters, the system implements a multi-stage filtering pipeline consisting of RSSI thresholding and infrastructure-device exclusion.

RSSI thresholding (range limiting). To reduce bleed-through from adjacent spaces, the system applies a hard RSSI cutoff before estimation. In the system implementation, empirical testing showed that a threshold around −85 dBm effectively filters out many devices originating from corridors or neighboring rooms while retaining devices within the immediate scanning radius (approximately 5-10 meters in the tested building layouts). In the data-collection protocol, an RSSI threshold of −88 dBm was used to approximate a ∼10-meter effective radius and reduce interference from nearby environments during scanning.

Night-baseline infrastructure detection and exclusion. A major engineering issue in university environments is the presence of *fixed* BLE emitters (e.g., desktop peripherals, printers, projectors, or other always-on devices) that do not correlate with human occupancy and therefore introduce a constant "phantom occupancy" offset. The system addresses this via a calibration mode executed when the room is assumed empty. In this architecture, the calibration scan is performed during night hours (between 21:00 and 06:00), and stable identifiers observed in this empty-state window are stored in a local blocklist; during live operation, any matching IDs are discarded before estimation. In the analysis pipeline, the same assumption (building closure between 21:00-06:00) is used to estimate a room-specific baseline of fixed devices D_{fissi} using the median device count observed in that interval, enabling a net device quantity:

$$D_{\text{net}} = D - D_{\text{fissi}}. \quad (2)$$

This net count stabilizes device-to-people relationships, especially in laboratories with substantial fixed-device presence.

4.3 Semantic Layer: Context Inference and Occupancy Estimation

The semantic layer converts filtered detections into an occupancy estimate by incorporating context-dependent behavior and uncertainty-aware reporting. The key architectural motivation is that the "device-per-person" ratio is not constant across environments (e.g., laboratories often exhibit more concurrently active peripherals than lecture settings), so the system infers environmental context and selects estimation parameters dynamically.

Two complementary inference strategies arise: one based on statistical patterns in device distributions across different room types, and another driven by temporal cues such as scan times and expected usage profiles of university spaces.

- **Composition-based inference (service/manufacturer cues)** Our architecture describes a context inference engine that analyzes the composition of detected devices, using cues such as Service UUIDs (Universally Unique Identifiers) and Manufacturer Data to identify device classes. A higher density of peripheral-like devices (e.g., mice/keyboards) and higher RSSI variance indicates a workspace-like environment (laboratory profile), whereas a more homogeneous device set dominated by smartphones with lower variance indicates a seated audience (lecture/seminar profile).
- **Expectation-distance inference (time-conditioned profiles)** The proposed system specifies a *RoomTypeEstimator* that computes expected device levels for each room type at the scan time, measures the distance between observed counts and expected profiles, selects the room type with minimum distance, and outputs a confidence score for that choice.

Estimation and uncertainty-aware output. After filtering and context inference, the system estimates occupancy by adjusting device counts for fixed infrastructure and dividing by a context-appropriate device-per-person coefficient. In the formulation, the estimate is derived from the number of unique, non-blocklisted devices and a context-selected coefficient:

$$P_{\text{est}} = \frac{|S_{\text{filtered}}| - N_{\text{static}}}{k_{\text{context}}}. \quad (3)$$

Here, $|S_{\text{filtered}}|$ is the count of unique devices remaining after RSSI thresholding and blocklist removal, N_{static} denotes the estimated/identified fixed-device contribution, and k_{context} is selected based on inferred environment characteristics.

Rather than outputting a single deterministic integer, the architecture reports an interval to reflect RF uncertainty. In the design, a confidence/margin is computed using the standard deviation of RSSI values, where higher variance indicates instability (e.g., movement/obstruction) and therefore increases the uncertainty band:

$$\text{Margin} = P_{\text{est}} \pm \left(\alpha \cdot \frac{\sigma}{\sqrt{n}}\right), \quad (4)$$

with σ computed over RSSI readings within the scan window and n the number of observations. :contentReference[oaicite:16]index=16 :contentReference[oaicite:17]index=17 The proposed system similarly specifies a *PeopleEstimator* that produces an occupancy range using net devices (total minus fixed), applies calibrated margins derived from observed MAE, optionally corrects via time-of-day patterns, and enforces a minimum-range rule (the minimum estimated occupancy is always 1).

End-to-end data flow. Overall, the architecture executes the following pipeline each scan window: (1) collect S_{raw} via passive BLE scanning on a smartphone, (2) apply RSSI thresholding to limit spatial bleed-through, (3) remove fixed emitters using a night-baseline blocklist and/or baseline subtraction, (4) infer contextual room characteristics (room type) to select estimation parameters, and (5) compute and report an occupancy interval rather than a single point estimate.

4.4 Deployment & Workflow

This work adopts a modular system architecture that transforms raw BLE advertising traffic into a reliable occupancy estimate using only commercial-off-the-shelf smartphones. The architecture is organized into three layers: Data Acquisition, Signal Processing, and Semantic Inference, each responsible for a distinct stage of the pipeline. This structure reflects the practical constraints mentioned previously, where smartphones operate as passive scanners and personal devices act as advertisers, emitting packets whose presence and RSSI values serve as the basis for the estimation process.

The Data Acquisition Layer performs passive BLE scanning in fixed 30-second windows, aggregating unique identifiers and their RSSI values. This design filters asynchronous streams and RSSI-based filtering to isolate devices within the target area. The system supports both single-scan operation for ground-truth collection and periodic scanning for long-term deployments (a scan every 15 minutes), enabling continuous monitoring without requiring user intervention.

The Signal Processing Layer addresses the inherent noise of indoor BLE environments. Because BLE signals are attenuated, subject to multipath, and affected by interference from adjacent rooms, the system applies RSSI thresholding to limit the effective detection range. It also incorporates a night-baseline procedure to identify fixed devices, such as computers, peripherals, or always-on equipment, that do not correspond to human presence. Subtracting these fixed emitters stabilizes the relationship between detected devices and actual occupancy, particularly in laboratory settings where the number of static devices is high.

The Semantic Layer converts the filtered detections into an occupancy estimate by incorporating context-dependent behavior. Different environments exhibit distinct device-to-person ratios and interference patterns. The system therefore infers the room type based on the composition and variability of detected devices and selects appropriate estimation parameters. Rather than producing a single deterministic value, the system outputs an uncertainty-aware range that reflects the variability of RSSI measurements and the statistical nature of indoor radio propagation.

Overall, the architecture provides an end-to-end pipeline that begins with passive BLE scanning on a smartphone and ends with

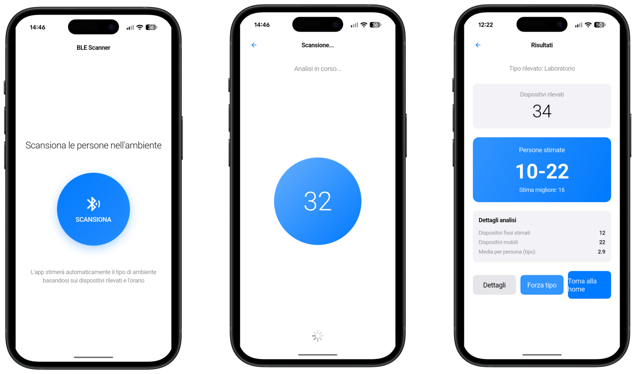

(a) Initial scan interface.

(b) Scanning in progress.

(c) Estimated occupancy results.

Figure 1: BLE scanning application.

a context-aware occupancy estimate. Each scan window undergoes range limiting, fixed-device removal, room-type inference, and calibrated estimation, enabling the system to operate with minimal manual room classification and without dedicated sensing infrastructure.

In Figure 1, the three screenshots illustrate the workflow of the mobile application used for passive BLE scanning and occupancy estimation. In image (a), the home screen of the app displays a central button labeled *Scan*, accompanied by the instruction at the top, "Scan the people in the environment," and a note at the bottom stating, "The app automatically estimates the environment type based on detected devices and time."

When the user taps the *Scan* button, the app initiates a BLE scanning session, as shown in image (b). The number displayed in the center represents the count of detected BLE devices during the ongoing scan, while the text "Analysis in progress..." indicates that the system is collecting and processing advertising packets within the current scan window.

Once the scanning process completes, the user is directed to the results screen shown in image (c), where the app presents the inferred environment type (e.g., Laboratory), the total number of detected devices, and the estimated occupancy range. The analysis section within this interface further distinguishes between fixed and mobile devices and reports the average device-to-person ratio used for the estimation. This example workflow demonstrates the end-to-end operation of the proposed system, from passive data acquisition to semantic inference and probabilistic occupancy estimation, executed entirely on a standard smartphone without dedicated infrastructure.

5 Occupancy Estimation Model

This section formalizes the proposed occupancy estimation model, which converts raw BLE detections observed within a scan window into a probabilistic estimate of human occupancy. The model is designed to correct three systematic distortions in passive BLE sensing: (i) the *multi-device phenomenon* (multiple BLE emitters per person), (ii) *fixed infrastructure emitters* (persistent devices unrelated to human presence), and (iii) *inter-room interference* (signals originating outside the target room). The output is reported as an occupancy *range* rather than a single deterministic value to reflect the inherent variability of RF sensing.

5.1 Mathematical Formulation

Let W denote a scan window (e.g., 30 seconds) and let $S_{\text{total}}(W)$ be the set of all unique device identifiers detected during W. The raw device count is:

$$N_{\text{raw}} = |S_{\text{total}}(W)|\,. \tag{5}$$

Since N_{raw} includes personal devices, fixed emitters, and noise from adjacent spaces, the estimated human occupancy P_{est} is computed as:

$$P_{\text{est}} = \frac{N_{\text{raw}} - (N_{\text{static}} + N_{\text{noise}})}{k_{\text{context}}}. \tag{6}$$

The three correction terms are defined as follows:

- N_{static}: the number of fixed infrastructure devices (e.g., always-on peripherals) that must be removed to avoid permanent "phantom occupancy" offsets.
- N_{noise}: the portion of detections attributed to inter-room bleed-through and low-confidence signals, largely mitigated by RSSI-based range limiting.
- k_{context}: a context-dependent "devices-per-person" coefficient that corrects for multi-device users by adapting to the inferred environment.

5.2 Noise Filtering and Effective Spatial Boundaries

Passive BLE signals can traverse walls and propagate through corridors, producing detections that do not correspond to people inside the target space. To reduce this inter-room interference, the model applies an RSSI threshold to limit detections to a practical sensing radius. Let τ_{rssi} be the minimum accepted RSSI value (e.g., around −85 dBm in the deployment setting). The filtered detection set is:

$$S_{\text{filt}}(W) = \{d \in S_{\text{total}}(W) \mid \text{RSSI}(d) \geq \tau_{\text{rssi}}\}. \tag{7}$$

This filtering step reduces N_{noise} by discarding devices detected at very weak signal levels, which are more likely to originate from adjacent rooms. In practice, N_{noise} can be interpreted as the detections removed by thresholding:

$$N_{\text{noise}} = |S_{\text{total}}(W)| - |S_{\text{filt}}(W)|\,. \tag{8}$$

5.3 Infrastructure Subtraction via Empty-State Blocklisting

A core source of overestimation in university environments is the presence of fixed BLE emitters (e.g., desktop-related peripherals, smart boards, or other always-on devices). These emitters are largely stable across time and do not correlate with occupancy changes. The model addresses this through an *empty-state calibration* procedure.

During a time interval when the room is assumed empty (e.g., night hours), the system performs one or more scans and identifies stable, high-signal emitters. Let $\mathcal{B}$ be a blocklist containing identifiers classified as infrastructure devices. During live operation:

$$N_{\text{static}} = \sum_{d \in S_{\text{total}}(W)} \mathbb{I}(d \in \mathcal{B}), \tag{9}$$

where $\mathbb{I}(\cdot)$ is the indicator function.

Operationally, the blocklist is populated by flagging devices that appear consistently during the empty state with strong RSSI (e.g., RSSI > -70 dBm), since these are likely to be stationary devices physically located in the room. After blocklist removal, the effective in-room, non-infrastructure device count becomes:

$$N_{\text{eff}} = |S_{\text{filt}}(W)| - N_{\text{static}}. \tag{10}$$

5.4 Context-Aware Device-to-Person Coefficient

The most influential factor in device-based occupancy estimation is the relationship between detected devices and actual people. This ratio varies by environment because occupant behavior changes with context:

- **Low-activity contexts (Classroom mode).** Occupants are typically seated; secondary devices (e.g., laptops) may be inactive, and the dominant emitter is often the smartphone. Empirically, the device-per-person ratio is relatively low (approximately $k \approx 1.2$-1.4).
- **High-activity contexts (Laboratory mode).** Occupants are actively working; laptops, mice/keyboards, and audio peripherals are more likely to be active simultaneously, increasing detectable emitters per person (approximately $k \approx 2.0$-2.8).

To adapt to this variation without requiring manual room labels, the model selects k_{context} using environment cues extracted from scan data. A compact feature vector is defined as:

$$\mathbf{F} = \left[\sigma_{\text{rssi}}, \rho_{\text{peripheral}}\right], \tag{11}$$

where σ_{rssi} is the standard deviation of RSSI values observed in the scan window and $\rho_{\text{peripheral}}$ is the density of peripheral-class signals (e.g., device classes inferred from observed service identifiers).

A decision rule selects the coefficient:

$$k_{\text{context}} = \begin{cases} 1.3 & \text{if } \rho_{\text{peripheral}} < \tau_{\text{lab}} \quad \text{(Classroom mode)} \\ 2.5 & \text{if } \rho_{\text{peripheral}} \geq \tau_{\text{lab}} \quad \text{(Laboratory mode)} \end{cases} \tag{12}$$

If the context classification is uncertain (e.g., intermediate peripheral density), the model falls back to a conservative weighted average (e.g., $k = 1.8$) and marks the output as lower confidence.

5.5 Final Estimate and Confidence Interval

Using the effective device count and the selected coefficient, the point estimate is:

$$P_{\text{est}} = \frac{N_{\text{eff}}}{k_{\text{context}}} = \frac{|S_{\text{filt}}(W)| - N_{\text{static}}}{k_{\text{context}}}. \tag{13}$$

Because RF propagation is stochastic (multipath, body shadowing, packet collisions), a single integer estimate can be misleading. Therefore, the model reports an occupancy *range*:

$$[P_{\min}, P_{\max}] = [P_{\text{est}} - \delta,\ P_{\text{est}} + \delta], \tag{14}$$

where δ is a variance-dependent margin computed from signal instability. Using RSSI dispersion as a proxy for instability:

$$\delta = \alpha \cdot \log\left(1 + \sigma_{\text{rssi}}\right), \tag{15}$$

with α calibrated empirically from observed estimation error under representative deployment conditions.

Finally, the reported estimate is rounded and bounded to enforce valid occupancy outputs (e.g., non-negative values), and can optionally enforce a minimum reported occupancy in operational settings when required by the application logic.

5.6 Summary of the Estimation Pipeline

For each scan window W, the model executes:

(1) Collect all detections and compute $S_{\text{total}}(W)$ and N_{raw}.
(2) Apply RSSI thresholding to obtain $S_{\text{filt}}(W)$ and reduce inter-room interference.
(3) Subtract fixed infrastructure devices using the empty-state blocklist to compute N_{eff}.
(4) Infer context features and select k_{context}.
(5) Compute P_{est} and output an uncertainty-aware range $[P_{\min}, P_{\max}]$.

6 Experimental Evaluation

This section presents the empirical validation of the proposed occupancy estimation system. The evaluation is structured to assess (i) the accuracy of the context-aware estimation model with respect to manually collected ground truth data, (ii) the robustness of the night-baseline infrastructure subtraction strategy, and (iii) the stability of the system under continuous multi-week deployment in real university environments.

6.1 Evaluation Setup

Environments. Experiments were conducted across representative university spaces categorized into three environment types:

- **Laboratories**, characterized by the presence of multiple fixed desktop stations and peripherals;
- **Classrooms**, characterized by scheduled lectures and minimal fixed BLE emitters;
- **Corridors**, characterized by transient flows and potential inter-room interference.

In the initial calibration phase, three reference spaces were used for detailed ground-truth analysis:

- Computer Laboratory (laboratory profile),
- Classroom (classroom profile),
- Hall (corridor profile).

Ground Truth Collection. To construct a reliable reference dataset, manual BLE scans were performed using the smartphone application in single-scan mode. During these scans, the operator inserted the real number of occupants present at the moment of measurement. Scans were performed three times per day (09-11, 12-14, 15-17) until at least 70 samples were collected per environment. Only devices with RSSI greater than -88 dBm were considered, limiting the effective sensing radius to approximately 10 meters and reducing interference from adjacent spaces.

Manual scans were labeled with the true occupancy value ($P_{\text{real}} \geq 0$) and used exclusively for calibration and validation of the estimation model.

Continuous Data Collection. After the calibration phase, the system was deployed in continuous scanning mode. In this configuration, the application performed an automatic BLE scan every 15 minutes (96 scans per day). The resulting daily measurements were aggregated into CSV files and exported automatically. Periodic scans do not include ground-truth occupancy values and are used to:

- analyze temporal patterns,
- estimate fixed-device baselines during empty-state hours (21:00-06:00),
- validate the long-term stability of the estimation model.

6.2 Characterization of Environment Profiles

Laboratory Environment. Laboratories exhibited the most complex signal structure due to the presence of fixed infrastructure devices. In the Gamification Lab, night-time measurements revealed a stable baseline of approximately 15-17 persistent devices. During peak hours (around 11:00 and 15:00), total detected devices frequently exceeded 60 and, on examination days, occasionally surpassed 100.

Scatter analysis of manual scans showed that the difference between total devices and real occupants increases with occupancy, reflecting the multi-device phenomenon (e.g., laptops, Bluetooth mice, headphones). After subtracting the night baseline, the relationship between net devices (D_{net}) and people became more stable, supporting the use of a context-specific coefficient ($k \approx 2.5$ in high-activity contexts).

Classroom Environment. In the classroom, no persistent fixed-device baseline was detected during night hours. The difference between devices and people remained relatively stable (approximately 40-44 devices across time slots in the manual calibration dataset), suggesting a quasi-linear relationship between device count and occupancy. This behavior supports a lower device-per-person coefficient (approximately $k \approx 1.3$) for low-activity, lecture-based contexts.

Temporal analysis showed occupancy peaks aligned with lecture schedules and minimal activity outside teaching hours.

Corridor Environment. The Hall exhibited high variability and pronounced inter-room interference. Device counts were significantly inflated during adjacent classroom sessions, particularly during 12-14 time slots. Correlated peaks between the Hall and the previously discusses nearby classroom, confirmed BLE signal bleed-through through thin walls. In this environment, raw device counts cannot be directly interpreted without aggressive RSSI filtering and contextual awareness.

6.3 Baseline Subtraction Effectiveness

To evaluate the effectiveness of the night-baseline calibration method, we compared raw device counts (D) and net device counts ($D_{\text{net}} = D - D_{\text{fissi}}$) in laboratory environments.

The use of the median device count between 21:00-06:00 as D_{fissi} significantly reduced the permanent offset introduced by infrastructure devices. In the Gamification Lab, subtracting the baseline reduced systematic overestimation and stabilized the device-to-person relationship across time slots.

Without baseline subtraction, occupancy would be overestimated by approximately the magnitude of the fixed-device count (15-17 units in this case), demonstrating the necessity of this calibration step for infrastructure-rich environments.

6.4 Model Accuracy Against Ground Truth

Accuracy evaluation was performed using only manual scans, where real occupancy values were available. For each scan window, the following were computed:

- Raw device-based estimate (naïve counting without correction),
- Static-coefficient estimate (single global k),
- Context-aware estimate (dynamic k_{context} with baseline subtraction).

Error metrics included:

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^{N} \left| P_{\text{est}}^{(i)} - P_{\text{real}}^{(i)} \right| \tag{16}$$

$$\text{MAPE} = \frac{100}{N} \sum_{i=1}^{N} \left| \frac{P_{\text{est}}^{(i)} - P_{\text{real}}^{(i)}}{P_{\text{real}}^{(i)}} \right| \tag{17}$$

The context-aware model consistently reduced estimation error compared to both naïve device counting and static coefficient approaches, particularly in laboratories where multi-device behavior and infrastructure devices are most pronounced.

6.5 Temporal Stability and Multi-Week Deployment

Continuous scanning over multiple weeks enabled evaluation of temporal robustness. The following observations were made:

- Laboratories show recurring daily peaks around 11:00 and 15:00, confirming repeatable occupancy patterns.
- Classrooms exhibit strong schedule-driven peaks with near-zero activity outside teaching hours.
- Corridors display highly dynamic patterns with transient spikes during class transitions.

Figure 2 shows the average daily device count, obtained by averaging the measured device counts over multiple weeks to represent a typical day.

The automatic calibration pipeline periodically integrates new data, recomputing baseline medians and refining estimation parameters. This ensures that the model remains stable even when device usage patterns evolve over time (e.g., increased wearable adoption or infrastructure changes).

6.6 Summary of Evaluation Findings

Table 1 summarizes the error metrics obtained for each room type at different stages of the occupancy estimation pipeline. The first row of each room type presents occupancy estimates obtained using only fixed-device subtraction. The second row additionally applies a

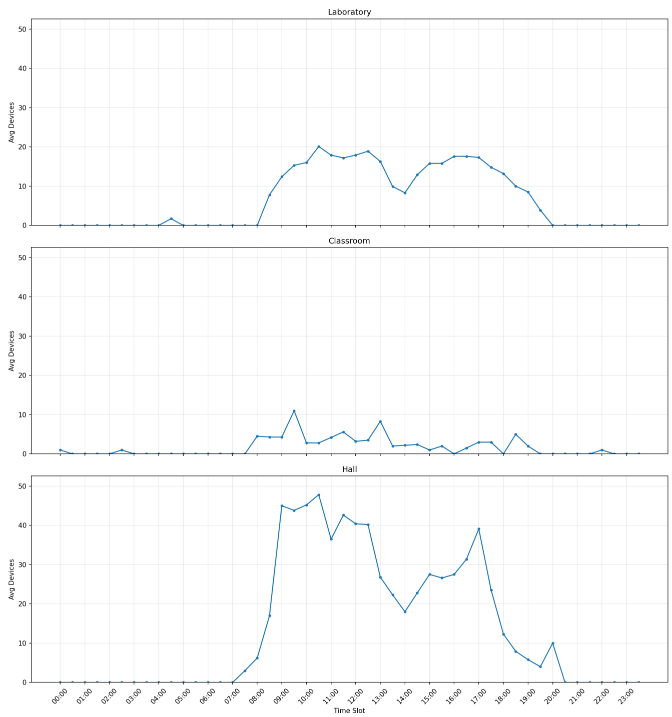

Figure 2: Temporal variation in the number of detected devices across different environments during a single day.

fixed device-per-person coefficient. Finally, the third row shows the results produced by the complete pipeline, which incorporates all processing stages and context-aware adjustments. Moreover, Figure 3 compares the occupancy estimation results for different room types across successive stages of the proposed processing pipeline and at different time slots, with the presented data corresponding to averages over selected time slots; only those time slots with more than three detected devices (N>3), excluding fixed devices, are considered, as they provide the most significant results.

The experimental evaluation demonstrates that:

- Night-baseline subtraction is essential in laboratory environments with significant fixed-device presence.
- Static device-to-person coefficients are insufficient across heterogeneous environments.
- Context-aware coefficient selection improves estimation accuracy and reduces systematic bias.
- Snapshot-based estimation avoids cumulative drift observed in transition-based systems.
- The system operates reliably on commodity smartphones without requiring structural information or dedicated sensing hardware.

Overall, the results confirm that passive BLE scanning, when combined with calibrated infrastructure subtraction and context-aware modeling, provides a robust and scalable solution for real-time indoor occupancy estimation in university environments.

Table 1: Error metrics for the occupancy estimation pipeline across different room types and processing stages.

Room Type	Method	MAE	MAPE (%)
Classroom	Scanned devices only	44.62	129.6
	Fixed subtraction + flat DPP	13.15	28.2
	Full pipeline	8.56	20.5
Corridor	Scanned devices only	41.69	312.0
	Fixed subtraction + flat DPP	16.50	82.2
	Full pipeline	12.96	74.2
Laboratory	Scanned devices only	12.73	178.0
	Fixed subtraction + flat DPP	0.37	8.6
	Full pipeline	0.37	8.6

7 Conclusion and Future Work

This paper presented an infrastructure-free approach for indoor occupancy estimation based on passive BLE scanning using commodity smartphones. The core contribution is a deployable engineering pipeline that transforms raw BLE advertising observations into occupancy estimates by explicitly addressing the main failure modes that limit device-counting approaches in real buildings: (i) the multi-device phenomenon, where a single occupant may contribute multiple BLE emitters, (ii) persistent fixed infrastructure devices that create a constant phantom offset, and (iii) inter-room signal bleed-through caused by RF propagation in dense indoor layouts.

The proposed solution operationalizes these challenges through a modular three-layer architecture. A data acquisition layer performs periodic scan windows with an RSSI threshold to approximate the physical boundaries of a room and reduce external interference. A signal processing layer identifies and removes stable fixed emitters using an empty-state (night) baseline strategy, yielding a net device count that better reflects human presence. Finally, a semantic layer infers environmental context and dynamically selects context-appropriate parameters, producing an uncertainty-aware occupancy range rather than a single deterministic count. Field deployment across heterogeneous university environments (laboratories, classrooms, and corridors) demonstrates that context-aware estimation and baseline subtraction consistently reduce the error of naive device-based approaches, particularly in infrastructure-rich spaces where fixed emitters and peripheral density would otherwise dominate the signal.

7.1 Limitations

While passive BLE scanning enables a low-cost and privacy-preserving sensing alternative to camera-based or infrastructure-heavy solutions, it remains subject to limitations inherent to RF-based sensing and to modern device ecosystems. First, BLE detectability is not guaranteed: some users may keep Bluetooth disabled, and certain devices reduce or stop advertising due to operating system policies, sleep states, or power-saving modes. Second, RSSI is an imperfect proxy for distance and is affected by multipath propagation, body shadowing, and packet collisions, which can cause temporal instability even under constant occupancy. Third, inter-room interference cannot be fully eliminated by thresholding alone in buildings

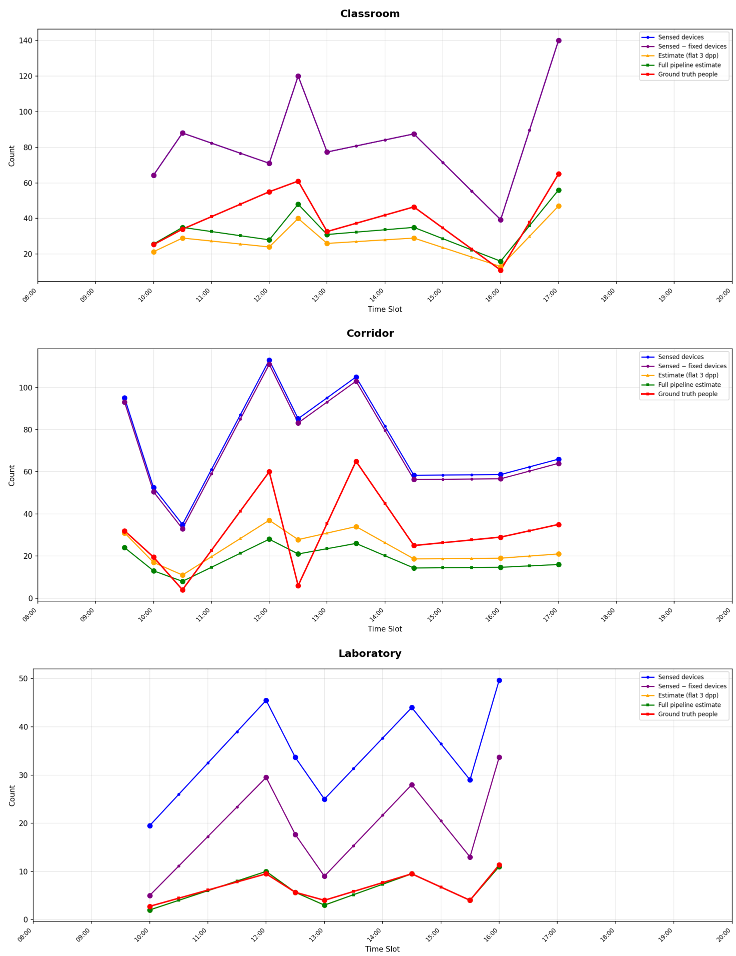

Figure 3: Comparison of occupancy estimation performance across different room types. From left to right: classroom, laboratory, and corridor environments. Within each figure, the results are shown for successive stages of the processing pipeline: fixed-device subtraction only, fixed-device subtraction with a fixed device-per-person coefficient, and the complete context-aware processing pipeline. The presented values correspond to averages over the selected high-activity time slots.

with thin walls, open corridors, or closely adjacent rooms. Finally, the device-to-person mapping remains a model-based approximation and may require recalibration when the population's device habits change (e.g., increasing wearable adoption) or when room configurations are modified (e.g., added always-on equipment).

Our deployment relied on an empirically tuned RSSI threshold to limit detections to the physical boundaries of each room. While this approach effectively reduced inter-room interference in the university building studied, we acknowledge that a single threshold may not generalize to more complex architectural layouts or materials. During data collection, we observed clear bleed-through patterns, particularly between the Hall and the adjacent classroom, highlighting the sensitivity of BLE propagation to wall thickness and spatial configuration. Future work will focus on exploring adaptive or multi-threshold strategies. In this approach, we could utilize a crowd-sourced version of our proposed system to gather data from multiple mobile devices scanning different adjacent rooms. This data could help us better define spatial boundaries and reduce cross-room interference in larger and more diverse buildings.

The system currently models three types of environments: Classrooms, Laboratories, and Corridors. However, the semantic layer is not limited to these categories. It is designed to be extensible, allowing for the inclusion of additional room types such as offices, meeting rooms, or waiting areas. This is achieved by collecting calibration data to derive the necessary device-to-person coefficients and fixed-device baselines. This approach is applicable not only in university settings but also in other indoor environments. However, introducing a new category requires a dedicated data collection phase. The accuracy of the resulting model depends strongly on the representativeness of the collected samples. While generalization to unseen room types is possible, it is contingent on further calibration efforts.

7.2 Future Work

Several directions can strengthen robustness, scalability, and integration into real building workflows:

Beyond Single-Point BLE Scanning. A natural extension is to deploy multiple scanning mobile devices across adjacent spaces and jointly infer room-level occupancy. Coordinated sensing gathered from these mobile devices could reduce cross-room bleed-through by leveraging spatial diversity and consistency checks (e.g., assigning detections to the most plausible room based on relative signal strength patterns). This also enables coverage of larger areas and reduces reliance on a single point of observation.

Adaptive calibration and drift monitoring. Although the current approach supports continuous data collection and periodic baseline estimation, future work can formalize automatic drift detection: identifying when the fixed-device baseline changes, when device-per-person behavior shifts, or when estimation confidence degrades. This could trigger targeted recalibration scans (e.g., additional empty-state sampling) or conservative fallback parameterization until stability is restored.

Richer context inference and confidence modeling. The context inference component can be expanded to recognize additional environment states beyond coarse room categories (e.g., exam sessions vs. lectures, lab work vs. idle). Similarly, uncertainty quantification can be improved by combining signal variance with additional stability indicators (e.g., fraction of newly appearing identifiers, persistence across consecutive windows), producing more informative confidence intervals for downstream decision-making.

Lightweight machine-learning extensions. The current system intentionally avoids heavy machine-learning pipelines and instead relies on calibrated, context-dependent coefficients within a snapshot-based occupancy estimation model. Future work could explore lightweight supervised or semi-supervised models to complement this estimator, particularly for adapting to new buildings, populations, or temporal usage patterns. Models such as decision trees, random forests, or shallow neural networks could learn context-specific features (e.g., RSSI variance, identifier persistence, temporal stability indicators) without requiring large labeled datasets. These approaches would preserve the infrastructure-free nature of the system while enabling automatic adaptation to diverse architectural layouts and occupant behaviors, extending the generalizability demonstrated in our earlier BLE-based work.

Integration with building operations and interaction flows. Future work can strengthen the interactive computing systems perspective by integrating the estimator into building dashboards and resource-management workflows (e.g., occupancy-aware HVAC control, room availability displays, or space planning analytics). This includes designing interaction mechanisms for administrators to validate estimates, annotate anomalies, and manage privacy-preserving aggregation policies.

Generalization across buildings and populations. Finally, broader validation across different architectural layouts and occupant populations would improve generalizability. This includes exploring how threshold selection, baseline calibration periods, and context-dependent coefficients transfer across campuses, offices, and other semi-public indoor settings, and identifying which elements can be standardized versus learned per site. The presented system shows that passive BLE scanning, when engineered with calibration, context-aware modeling, and uncertainty-aware reporting, can provide an effective infrastructure-free solution for indoor occupancy estimation. The outlined future directions aim to further improve reliability in challenging environments and to support seamless integration into real-world smart building operations.

References

[1] Md Fazlay Rabbi Masum Billah, Nurani Saoda, Jiechao Gao, and Bradford Campbell. 2021. BLE can see: a reinforcement learning approach for RF-based indoor occupancy detection. In *Proceedings of the 20th international conference on information processing in sensor networks (co-located with CPS-IoT week 2021)*. Association for Computing Machinery, New York, NY, USA, 132–147. doi:10.1145/3412382.3458262

[2] Alba Bisante, Venkata Srikanth Varma Datla, and Gabriella Trasciatti. 2024. An Approach to Leverage Artificial Intelligence for Car-Parking Related. In *Engineering Interactive Computer Systems. EICS 2023 International Workshops and Doctoral Consortium: Swansea, UK, June 26-27, 2023, Selected Papers*, Vol. 14517. Springer Nature, Springer, Cham, Cham, Switzerland, 63.

[3] Pratiksha Chaudhari, Yang Xiao, Mark Ming-Cheng Cheng, and Tieshan Li. 2024. Fundamentals, algorithms, and technologies of occupancy detection for smart buildings using IOT sensors. *Sensors* 24, 7 (2024), 2123.

[4] Jiayu Chen, Hainan Chen, and Xiaowei Luo. 2019. Collecting building occupancy data of high resolution based on WiFi and BLE network. *Automation in Construction* 102 (2019), 183–194.

[5] Venkata Srikanth Varma Datla, Alessandro Aiuti, Alba Bisante, Gabriella Trasciatti, Stefano Zeppieri, and Emanuele Panizzi. 2025. Detecting Human Presence via Smartphone BLE Beaconing: Preliminary Investigations. In *Proceedings of the 24th International Conference on Mobile and Ubiquitous Multimedia*. Association for Computing Machinery, New York, NY, USA, 483–485.

[6] Venkata Srikanth Varma Datla, Stefano Zeppieri, Alessandro Aiuti, Alba Bisante, Gabriella Trasciatti, and Emanuele Panizzi. 2025. Towards Context-Aware UX in Automated Mobility: BLE Based Passenger Detection via Smartphones. In *Proceedings of the 16th Biannual Conference of the Italian SIGCHI Chapter*. Association for Computing Machinery, New York, NY, USA, 1–7.

[7] Florenc Demrozi, Fabio Chiarani, and Graziano Pravadelli. 2021. A low-cost BLE-based distance estimation, occupancy detection and counting system. In *2021 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, New York, NY, USA, 1430–1433.

[8] Florenc Demrozi, Cristian Turetta, Fabio Chiarani, Philipp H Kindt, and Graziano Pravadelli. 2021. Estimating indoor occupancy through low-cost BLE devices. *IEEE Sensors Journal* 21, 15 (2021), 17053–17063. doi:10.1109/JSEN.2021.3080632

[9] H Elkhoukhi, M Bakhouya, Abdellatif Elmouatamid, D El Ouadghiri, and J Gaber. 2025. A review of occupancy sensing technologies and approaches in smart buildings. *International Journal of RF Technologies* 15, 1 (2025), 19–48.

[10] Avgoustinos Filippoupolitis, William Oliff, and George Loukas. 2016. Bluetooth low energy based occupancy detection for emergency management. In *2016 15th international conference on ubiquitous computing and communications and 2016 International Symposium on Cyberspace and Security (IUCC-CSS)*. IEEE, New York, NY, USA, 31–38.

[11] Pedro J García-Paterna, Alejandro S Martínez-Sala, and Juan Carlos Sánchez-Aarnoutse. 2021. Empirical study of a room-level localization system based on bluetooth low energy beacons. *Sensors* 21, 11 (2021), 3665.

[12] Giancarlo Iannizzotto, Lucia Lo Bello, and Andrea Nucita. 2023. Improving BLE-based passive human sensing with deep learning. *Sensors* 23, 5 (2023), 2581.

[13] Giancarlo Iannizzotto, Miryam Milici, Andrea Nucita, and Lucia Lo Bello. 2022. A perspective on passive human sensing with bluetooth. *Sensors* 22, 9 (2022), 3523.

[14] Dariusz Janczak, Wojciech Walendziuk, Maciej Sadowski, Andrzej Zankiewicz, Krzysztof Konopko, Maciej Slowik, and Malgorzata Gulewicz. 2024. Indoor positioning based on Bluetooth RSSI mean estimator for the evacuation supervision system application. *IEEE Access* 12 (2024), 55843–55858.

[15] Nasrettin Koksal, AbdulRahman Ghannoum, William Melek, and Patricia Nieva. 2025. Occupancy Monitoring Using BLE Beacons: Intelligent Bluetooth Virtual Door System. *Sensors* 25, 9 (2025), 2638.

[16] Md Shadab Mashuk, James Pinchin, Peer-Olaf Siebers, and Terry Moore. 2018. A smart phone based multi-floor indoor positioning system for occupancy detection. In *2018 IEEE/ION Position, Location and Navigation Symposium (PLANS)*. IEEE, New York, NY, USA, 216–227.

[17] Montserrat Mateos Sanchez, Roberto Berjon Gallinas, M Encarnacion Beato Gutierrez, and Ana M Fermoso Garcia. 2020. A tool to calculate the level of occupancy in indoor and outdoor spaces using BLE and open data to be published in real-time. *Sensors* 20, 14 (2020), 3916.

[18] June Young Park, Thomas Dougherty, and Zoltan Nagy. 2018. A bluetooth based occupancy detection for buildings. In *Proceedings of of Building Performance Analysis Conference and SimBuild. IBPSA*. International Building Performance Simulation Association, International, 807–814.

[19] Seonghyeon Park, Jaehan Joo, and Suk Chan Kim. 2025. Device-Free Occupancy Detection via Wi-Fi CSI and Deep Residual CNN. In *2025 Sixteenth International Conference on Ubiquitous and Future Networks (ICUFN)*. IEEE, New York, NY, USA, 193–195. doi:10.1109/ICUFN65838.2025.11169962

[20] Azkario Rizky Pratama, Widyawan Widyawan, Alexander Lazovik, and Marco Aiello. 2018. Multi-user low intrusive occupancy detection. *Sensors* 18, 3 (2018), 796.

[21] Guntur Dharma Putra, Azkario Rizky Pratama, Alexander Lazovik, and Marco Aiello. 2017. Comparison of energy consumption in Wi-Fi and bluetooth communication in a Smart Building. In *2017 IEEE 7th Annual Computing and Communication Workshop and Conference (CCWC)*. IEEE, New York, NY, USA, 1–6.

[22] Reza Shahbazian and Irina Trubitsyna. 2023. Human sensing by using radio frequency signals: A survey on occupancy and activity detection. *IEEE Access* 11 (2023), 40878–40904.

[23] George-Alexandru Stoian, Thiemo Voigt, and Christian Rohner. 2025. Augmenting BLE Fingerprinting Using Instantaneous Frequency. In *18th ACM Conference on Security and Privacy in Wireless and Mobile Networks*. Association for Computing Machinery, New York, NY, USA, 274–279.

[24] Zeynep Duygu Tekler, Raymond Low, Burak Gunay, Rune Korsholm Andersen, and Lucienne Blessing. 2020. A scalable Bluetooth Low Energy approach to identify occupancy patterns and profiles in office spaces. *Building and Environment* 171 (2020), 106681.

[25] Ondrej Zeleny, Tomas Fryza, Tomas Bravenec, Shoaib Azizi, and Gireesh Nair. 2024. Detection of room occupancy in smart buildings. *Radioengineering* 33, 3 (2024), 432–441.

[26] Han Zou, Yuxun Zhou, Jianfei Yang, and Costas J Spanos. 2018. Device-free occupancy detection and crowd counting in smart buildings with WiFi-enabled IoT. *Energy and Buildings* 174 (2018), 309–322. doi:10.1016/j.enbuild.2018.06.040

Received 14 February 2026; revised 12 March 2009; accepted 5 June 2009