ISCTE IUL
REPOSITÓRIO
INSTITUTO UNIVERSITÁRIO DE LISBOA

Repositório ISCTE-IUL

Deposited in *Repositório ISCTE-IUL*:
2026-09-18

Deposited version:
Accepted Version

Peer-review status of attached file:
Peer-reviewed

Citation for published item:
Marinheiro, R. N., Brito e Abreu, F., Vieira, T. & Martins, M. (2025). Adaptive crowd sensing with privacy-preserving WiFi fingerprinting. 2025 IEEE International Conference on Smart Internet of Things (SmartIoT). IEEE. https://doi.org/10.1109/SmartIoT66867.2025.00053

Further information on publisher's website:
10.1109/SmartIoT66867.2025.00053

Publisher's copyright statement:
This is the peer reviewed version of the following article: Marinheiro, R. N., Brito e Abreu, F., Vieira, T. & Martins, M. (2025). Adaptive crowd sensing with privacy-preserving WiFi fingerprinting. 2025 IEEE International Conference on Smart Internet of Things (SmartIoT). IEEE. https://doi.org/10.1109/SmartIoT66867.2025.00053, which has been published in final form at https://dx.doi.org/10.1109/SmartIoT66867.2025.00053. This article may be used for non-commercial purposes in accordance with the Publisher's Terms and Conditions for self-archiving.

Use policy

Creative Commons CC BY 4.0
The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a link is made to the metadata record in the Repository
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Serviços de Informação e Documentação, Instituto Universitário de Lisboa (ISCTE-IUL)
Av. das Forças Armadas, Edifício II, 1649-026 Lisboa Portugal
Phone: +(351) 217 903 024 | e-mail: administrador.repositorio@iscte-iul.pt
https://repositorio.iscte-iul.pt

Adaptive Crowd Sensing with Privacy-Preserving WiFi Fingerprinting

Rui Neto Marinheiro*‡, Fernando Brito e Abreu†‡, Tiago Vieira†‡, Miguel Martins†‡
*Instituto de Telecomunicações, Portugal
†Information Sciences, Technologies and Architecture Research Center (ISTAR)
‡Iscte – Instituto Universitário de Lisboa, Portugal
{rui.marinheiro, fba, tfava1, miguel_vale}@iscte-iul.pt

***Abstract*—This paper presents ongoing work in the context of a recently funded research project, MoniCrowd, aimed at advancing crowd monitoring in dynamic urban settings, particularly during temporary public events. The proposed system adopts an adaptive architecture based on passive Wi-Fi probe request detection, comprising portable sensors with multi-radio access connectivity and a rule-based fingerprinting method for anonymous and reliable device counting. Sensors autonomously select the best uplink from available connectivity options, enabling operation in suboptimal locations. A novel tool, the Information Elements Automatic Analyser (IEAA), enhances fingerprint robustness through fine-grained feature selection. A new dataset, collected under controlled Faraday cage conditions, supports this development. Preliminary field deployment results show strong correlation with manual counts, validating the approach under real-world conditions. The proposed fingerprinting method also achieved top accuracy in the international CONFRONT challenge. To optimise sensor performance in diverse environments, a UAV-assisted calibration tool is under development; its design and preliminary sensitivity mapping results are presented. Altogether, this work lays the foundation for scalable, autonomous crowd sensing solutions that can be rapidly deployed in dynamic urban contexts without requiring specialised local expertise.**

***Index Terms*—Crowd Monitoring, Wi-Fi Probe Requests, Privacy-Preserving Fingerprinting, Adaptive Sensor Deployment, Urban Sensing, Temporary Events**

I. Introduction

International tourist arrivals have increased dramatically, from 808.6 million in 2005 to an estimated 1.46 billion in 2024, according to UN Tourism [1]. Despite a sharp drop during the COVID-19 pandemic, global tourism has now rebounded to its previous peak levels (see Fig. 1), renewing concerns over overcrowding and sustainability.

Overcrowding in public spaces, especially in tourist or event-based settings, presents increasing challenges for urban authorities. Crowding situations degrade visitors' quality of experience, reducing their feeling of safety, affecting the use of basic services, while reducing destinations' authenticity [2]. Overcrowding also deteriorates the lives of residents, due to an increase in noise, less effective urban cleaning, and cultural clashes when visitors fail to respect local customs, traditions, and privacy, sometimes leading to the former expressing negative feelings against the latter [3]. Last, but not least, the environmental sustainability, structures, and cultural heritage of overcrowded destinations are also jeopardized, leading to a loss of authenticity [4]. To mitigate these concerns, smart

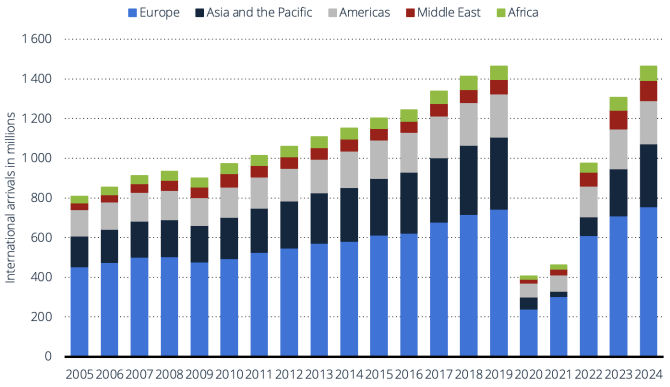


Fig. 1. International tourist arrivals worldwide from 2005 to 2024, by region

tourism tools (STTs) have been developed by this team for planning, monitoring, and analyzing collected crowding data, in the scope of the RESETTING [5, 6] project.

Regarding crowding monitoring, traditional methods such as surveillance cameras, acoustic sensing, or telecom data suffer from either high infrastructure cost, privacy concerns, or limited real-time insight. However, a low-cost and scalable architecture composed of crowding sensors has been achieved [7]. The number of people in the sensor's vicinity is estimated in real-time using a GDPR-compliant, passive wireless sensing approach that captures trace elements generated by mobile devices during Wi-Fi activity. This method counts devices to infer crowd density while addressing challenges such as MAC address randomization.

Sensors may communicate data using one of the preconfigured uplink options to mitigate the network limitations present at the sensor location. This solution has been thoroughly tested at different target areas, such as our University Campus. However, deployments were planned for long-term use, and the placement of sensors has been static with predictable connectivity. Nevertheless, local public administration authorities, like municipalities, face an additional major problem in their capacity to monitor overcrowding, predominantly during temporary special events that are planned, such as festivals, concerts, or anticipated to occur, like public demonstrations.

To address these challenges, a new research project, MoniCrowd, has been financed [8], with completion expected

in 2026. This initiative builds upon the current system to incorporate capabilities that enable flexible and temporary deployments. Given the unpredictable and short-term nature of these events, the project aims to deliver a monitoring solution that can be rapidly deployed or repositioned by local authorities as needed, without requiring specialised technical expertise from on-site personnel.

To fulfill these goals, MoniCrowd introduces an adaptive system architecture that integrates self-configuring sensors capable of operating in suboptimal locations with limited connectivity. The system presented in this paper includes the sensor design, incorporating a data-driven rule-based fingerprinting approach for anonymous device counting, and the backend system for data aggregation, remote control, and visualization. The sensors use the best available uplink option among available radio access technologies and operators, ensuring reliable communication under diverse conditions. Recognising that such locations may also lead to overlapping coverage or blind spots, the project incorporates a UAV-assisted calibration tool, whose design is presented in this work. Although still under development, manually surveyed sensitivity maps have been obtained to explore this approach. These maps are used to determine optimal sensor configurations, which are later updated over the air to maximise detection accuracy and minimise overcounting. An initial field deployment is also presented, validating the system's performance in a real-world scenario at a high-traffic tourist site.

In parallel, the project addresses key limitations of fingerprinting techniques that rely on outdated or limited datasets. This paper presents a new dataset currently being collected under controlled conditions using a custom-designed Faraday cage and introduces the Information Elements Automatic Analyser (IEAA), a tool developed to enhance fingerprinting robustness through fine-grained feature selection. While the current fingerprinting approach is based on a rule-based method validated in the CONFRONT challenge [9], future work will explore semi-supervised learning models, including Siamese networks, to further enhance adaptability to evolving device behaviours. The complete system will be validated in upcoming pilot deployments during temporary public events organised by local authorities. To ensure reusability, MoniCrowd will also deliver a comprehensive toolkit with technical documentation, tutorials, and training resources, enabling adoption by municipalities. To support broader adoption, the paper also provides a cost estimation analysis for deployments at different scales. The final expected result is that this adaptive system for crowd monitoring may be easily adopted by other public administration authorities, not only municipalities but also law enforcement or civil protection authorities [8].

The remainder of this paper is structured as follows. Section II reviews related work on Wi-Fi fingerprinting and crowd monitoring, while Section III details the system architecture, including the design of sensors and backend infrastructure. Section IV describes the implementation details and evaluates the system through controlled experiments and a real-world deployment, including preliminary results for calibration and a cost analysis. Finally, Section V concludes the paper and outlines directions for future work.

II. Related Work and Background

Several approaches may be considered for crowd detection, such as image capturing, sound capturing, social networks, mobile operators' data, and wireless spectrum analysis [10].

Capturing sound to detect crowds [11] leads to low precision and detection ranges. Image and video processing have also been used, where cameras provide visual information from a location, and images are processed through learning techniques [12]. However, for these approaches, processing is computationally demanding and lacks scalability, regarding implementation and communication costs. Other approaches relied on the geo-location information of content that is published on social media networks, such as Instagram, X, or Weibo, for the same purpose [13]. However, this approach depends on the usage of these platforms, thus making it less representative of the actual crowding. Moreover, all the mentioned approaches raise privacy issues. Alternatively, some authors have used massive volumes of data available from mobile operators [5, 14]. Nevertheless, these methods focus on retrieving patterns and predictions from past gathered data, making it inconvenient to generate relevant metrics in real time from those massive datasets. There is a need for solutions that are more capable of detecting crowds in real-time, while maintaining privacy.

For the above reasons, research on crowd monitoring has been gradually tending towards sensor-based monitoring approaches, supported by edge processing, fostered by the need to develop less-intrusive and privacy-preserving crowd-monitoring techniques that comply with restrictive privacy guidelines such as GDPR [15].

Sensing wireless communication traces from mobile devices has been a widely used approach for detecting crowds effectively, as most users carry a mobile phone. It has been assumed that the number of mobile devices in a given area is directly correlated with the number of people in the same area. This assumption is corroborated by De Meersman [16], who demonstrated that mobile phone data is a valuable source for statistical counting of people. Earlier approaches relied on counting the number of unique Media Access Control (MAC) addresses in messages emitted by mobile devices.

However, due to privacy concerns, since 2014, mobile device manufacturers have started to implement address randomization, making it more challenging to accurately estimate crowding. As such, research has advanced towards the exploration of other properties and fields of devices' messages. Several strategies have already been adopted for this purpose. Some authors employed a fingerprinting technique based on information elements (IE) fields to generate a unique identifier of each device [17, 18], while others combine this with other properties from device's messages, such as sequence numbers, burst size, time between messages in a burst, in a clustering algorithm, where each cluster should represent a unique device [19–22] or a device model [23]. However, most

processing has been performed in the cloud, making it difficult to comply with GDPR.

Earlier work [10] by this research team complies with users' privacy rights, performing passive non-intrusive detection, by implementing only edge processing. More recent efforts [5, 7, 24] continue along this line, using a public dataset [25] that was enriched with additional data collected under controlled laboratory conditions, as detailed in this paper. This enriched dataset enabled the development of the current data-driven, rule-based Wi-Fi fingerprinting algorithm, which went on to win the international CONFRONT challenge.

In contrast, camera- and acoustic-based systems, though capable of providing contextual information, face high deployment costs, heavy computational demands, and serious privacy concerns. Social media and mobile operator data allow large-scale analyses but are constrained by dependency on user activity or delayed availability, preventing real-time use. By comparison, our approach directly addresses these limitations by ensuring privacy by design, supporting real-time responsiveness, and enabling scalable deployments suitable for temporary public events under strict GDPR requirements.

III. System Overview

The proposed architecture for adaptive crowd monitoring is designed to address the challenges faced by public authorities during temporary events in urban settings, such as festivals or public demonstrations. It comprises three primary subsystems: the **CrowdingSensorUnit**, the **CrowdingServerUnit**, and the novel **AdaptiveCalibrationUnit**, each responsible for key stages in the sensing, calibration, and data management pipeline (see Fig. 2). This work extends a previous architecture [7] that introduced a Wi-Fi-based crowd monitoring solution with edge processing for anonymous fingerprinting.

The *CrowdingSensorUnit* represents the edge layer of the system, responsible for real-time, privacy-preserving detection of crowd presence. Each sensor is equipped with a passive Wi-Fi sniffer that captures probe requests from nearby mobile devices. These are processed locally using a lightweight *WiFiCrowdingDetector*, which performs device fingerprinting based on Information Elements (IEs). The local *SensorDBMS* stores configuration and detection records. The *DetectionEngine* counts distinct devices over configurable sliding intervals, while the *DataRetentionManager* enforces retention policies aligned with GDPR compliance. The *CommunicationManager* ensures robust uplink and downlink communication, enabling Multi-RAN connectivity and supporting both data transmission and over-the-air updates for sensor configuration, such as sensitivity adjustments. The *ControlManager* handles sensor control and configuration, including applying calibration updates received remotely.

The *AdaptiveCalibrationUnit* introduces a novel, automated sensor calibration methodology tailored for temporary and dynamic deployments. Central to this unit is the *MissionControl* module, which orchestrates a UAV-assisted calibration process. During installation, a drone equipped with a talkative device follows a predefined flight path to emit controlled probe requests across a grid of positions. These emissions are captured by the deployed sensors to generate spatial sensitivity maps. The *GeoTrack* module records high-precision UAV positioning using Real-Time Kinematic (RTK) corrections. The signal data and positioning are then processed by the *CalibrationEngine* in the *CrowdingServerUnit*, which implements an optimization function to derive ideal sensor parameters. These are applied remotely via the downlink channel.

The *CrowdingServerUnit* serves as the system's backbone, supporting two main verticals: data ingestion and visualization, and calibration and control. Incoming detection data are collected by the *CrowdingCollector* and stored in the *TimeSeriesDBMS*. Visualization and decision support are facilitated by the *VisualizationPlatform* and *AlertingEngine*, offering real-time insights. The *CalibrationEngine* processes calibration data and issues configuration updates, supported both by metadata from the *AdaptiveCalibrationUnit* and sensor data maintained in the *ServerConfigDBMS*. The *MessageServer* and *ProxyServer* enable Multi-RAN communication with sensors.

In contrast to the previous version of the architecture [7], the design proposed in this work introduces two key innovations: (i) the *AdaptiveCalibrationUnit*, which leverages UAV-assisted signal injection and RTK-based positioning to enable automated, location-aware sensor calibration; and (ii) integrated Multi-RAN communication capabilities, supporting flexible uplink and downlink connectivity and remote over-the-air configuration. These enhancements significantly increase the adaptability and scalability of the system for temporary and mobile deployments.

IV. Implementation and Deployment Results

The current implementation of the proposed crowd monitoring system and results from its initial deployment will be presented. It begins with the development of the sensor units and a data-driven, rule-based fingerprinting method for anonymous device identification. It then outlines the process of collecting high-quality training data in a controlled environment to support robust fingerprinting, including the validation of its use. The next subsection presents the central server and its role in data aggregation, remote control, and real-time visualization, supported by field deployment results. This deployment validates the system's accuracy in estimating crowd sizes under real-world conditions. The section continues with a description of an adaptive calibration strategy in development, which uses UAV-based signal scans to adjust the sensor configuration. It concludes with a cost estimation analysis for different deployment scales.

Sensor Implementation and Fingerprinting

The primary function of the Crowding Sensor Units (Sensor) is to estimate crowd density in specific urban locations by passively detecting Wi-Fi probe requests broadcast by nearby mobile devices. By identifying and counting distinct devices within a defined time window, each Sensor contributes to the real-time estimation of the number of individuals present in the monitored area. To ensure accurate counting in the presence

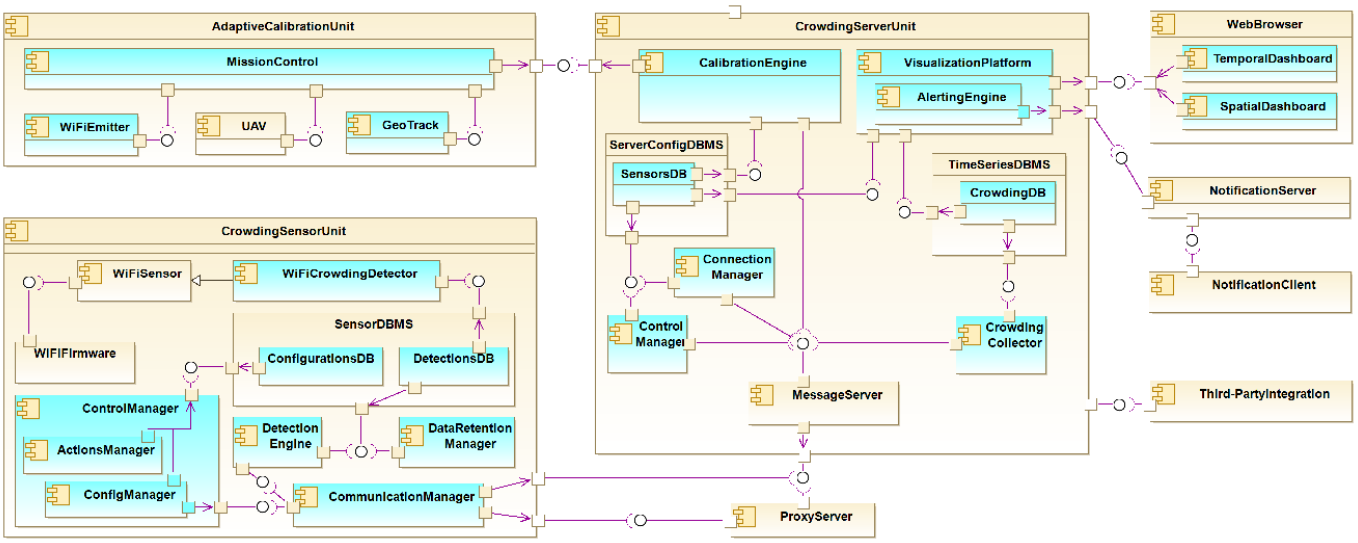


Fig. 2. Modelio Component Diagram

of MAC address randomization and repeated transmissions, a robust fingerprinting method is required to distinguish between unique devices and to avoid overcounting.

Device identification is performed using a novel fingerprinting method based on Information Elements (IEs) extracted from Wi-Fi probe requests. Previous approaches have typically used fewer IEs because they treated their content as undifferentiated byte sequences [17, 18, 23]. In contrast, this work introduces a more granular strategy, where fingerprint construction is not limited to the value of IEs as a whole but extends to their internal structure. Specifically, the method considers three distinct components from each IE: the ID Value (indicating the IE's presence), the ID Length, and the individual byte-level content. This allows for a more precise device representation, improving both distinctiveness and resilience to minor variability.

The selection of which IEs, and which components within them, should be used was performed automatically using the Information Elements Automatic Analyser (IEAA), a tool developed as part of this work. The IEAA was applied to a public dataset [25], comprising labelled probe requests for a set of different devices, as well as to data collected in a controlled local environment (described in this section). For each candidate IE, the tool evaluated three key criteria: its presence rate across devices, its temporal stability for individual devices, and its uniqueness at the byte level across devices. IE data that were frequently observed, showed consistent values over time for the same device, and exhibited high variability across different devices were considered suitable for fingerprinting.

The final selected set is presented in Table I, which also indicates which parts of each IE metadata (ID and Length) and content (specific bytes) are included. These extracted components are concatenated into a feature string, which is then hashed using the t1ha0 function from the t1ha library, known for its performance [26].

TABLE I
DATA IN SELECTED WI-FI INFORMATION ELEMENTS (IE) USED FOR FINGERPRINTING

Information Element (IE)	IE Data Considered			Byte #
	ID	Length	Content	Removed
Supported Rates	Yes	Yes	Yes	-
Extended Supported Rates	Yes	Yes	Yes	-
DS Parameter Set	Yes	No	No	-
HT Capabilities	Yes	Yes	Yes	5^{th}
VHT Capabilities	Yes	Yes	Yes	-
Extended Capabilities	Yes	No	Yes	-
RM Enabled Capabilities	Yes	Yes	Yes	-
Interworking	Yes	Yes	Yes	-
Vendor Specific	Yes	Yes	Yes	6^{th}, 8^{th}

This process yields a compact, deterministic, and data-driven rule-based fingerprint that enables passive device identification while preserving privacy, thereby enhancing the accuracy and reliability of crowd size estimation.

Each Sensor, shown in Fig. 3, is based on a Raspberry Pi and integrates an Alfa Network AWUS036ACH Wi-Fi adapter for probe request capture. To support multi-RAN connectivity, a Raspberry Pi IoT LoRa pHAT is included for LoRaWAN communication, operating alongside the onboard Wi-Fi interface. The Raspberry Pi is to be equipped with a Geekworm X1202 4-Cell 18650 UPS HAT, providing an off-grid power supply suitable for autonomous operation. All components are housed in a weather-resistant enclosure, designed for portable deployment during temporary urban events. The units run Kali Linux, selected for its comprehensive support of low-level networking tools, and employ a customised version of aircrack-ng specifically developed to extract IEs at the byte level for fingerprinting. The local software stack comprises

Python-based modules for control, configuration, and data handling, along with a local SQLite database for managing fingerprint data and sensor settings. An MQTT client ensures communication with the central Crowding Server Unit (Server), publishing anonymised and aggregated counts, and subscribing to control and configuration messages.



Fig. 3. Open and closed views of the crowding sensor unit.

Dataset Construction for Robust Fingerprinting

To implement effective device fingerprinting, training datasets with reliable ground truth are essential. These datasets may vary in nature, ranging from synthetic data [23] to real-world probe request traces [25]. While synthetic data allows for flexibility and scalability, datasets collected from actual environments offer greater realism. However, publicly available real-world datasets are limited in number and often small in size. To address this, the current work extends existing public datasets by collecting new data under controlled laboratory conditions.

To ensure data quality and avoid contamination from external signals, a custom-designed Faraday cage was constructed specifically for Wi-Fi probe request acquisition. This enclosure blocks electromagnetic radiation from external sources, enabling the collection of high-fidelity ground-truth data. The resulting datasets will be processed and made publicly available to support reproducibility and further research.

Fig. 4 shows the experimental setup used for collecting Wi-Fi probe requests within the Faraday cage. The structure is engineered to provide robust signal isolation, thereby ensuring the integrity of fingerprinting data collected in a controlled setting. It employs a multilayer shielding approach consisting, from the outside in, of an aluminium shell, an outer layer of RF shielding fabric, a wooden frame for structural support, a layer of conductive dissipative foam, and a final inner layer of RF shielding fabric.

This layered shielding design combines different electromagnetic interference (EMI) mitigation mechanisms:

- **Reflection**, via the aluminum shell and conductive RF fabrics, deflects incident high-frequency signals.
- **Absorption**, provided by the dissipative foam layer, reduces residual electromagnetic energy through dielectric and resistive losses within the material.
- **Dissipation**, distributed across all lossy layers, further attenuates any remaining signals by spreading and weakening signal energy as it traverses the shield.



Fig. 4. Faraday cage setup used to collect training data for fingerprinting.

This multi-mechanism approach reduces the intensity of incoming electromagnetic signals far more effectively than any single material could, ensuring a highly controlled environment for capturing uncontaminated wireless traces.

The dataset currently being collected using this Faraday cage, combined with existing public data, is essential to support the IEAA tool and enhance fingerprinting robustness. This has already been demonstrated by the success of the currently implemented Wi-Fi fingerprinting algorithm in the international CONFRONT challenge, where it achieved 295 out of 300 points across tasks of increasing difficulty, by accurately estimating the number of distinct devices in an unlabeled dataset.

However, any dataset is inherently prone to obsolescence as the use of wireless protocols and device characteristics evolves, making it increasingly difficult to maintain calibrated sensors capable of reliably distinguishing and counting unique devices. Addressing this limitation is critical to ensuring that fingerprinting methods remain effective across successive device generations and software updates, thereby supporting consistent long-term accuracy in crowd estimation. By proactively constructing a larger and more diverse dataset, this work lays the groundwork for applying machine learning techniques aimed at improving the adaptability of fingerprinting. In particular, semi-supervised learning approaches are well-suited to this setting, where large volumes of unlabeled data are readily obtainable, while labeled data remain limited. These methods offer a promising direction for enhancing fingerprint resilience without relying on extensive manual annotation. Specifically, the expansion of fingerprinting models using Siamese networks and the exploration of transfer learning across datasets are foreseen as key strategies for improving long-term sensor resilience.

Server Implementation and Deployment Validation

All data collected by the individual Sensors is transmitted to the central Crowding Server Unit (Server), a Virtual Private Server (VPS) running Ubuntu OS. The Server is responsible for aggregating, storing, and making the information available for real-time visualization, analysis, sensor control, and configuration. Communication between the Sensors and the Server is achieved via a lightweight message broker (Mosquitto MQTT), allowing transparent multi-RAN communications. The Server hosts a local time-series database (InfluxDB), which receives data from a Telegraf agent, and serves web-based interfaces implemented using Grafana dashboards, enabling access to both historical and real-time crowding metrics for each sensor location. Sensor configuration and calibration modules, developed in Python, rely on a PostgreSQL database.

One of the deployment scenarios tested so far consisted of two Wi-Fi crowding sensors (named Kiosk and Help-point) installed at strategic locations within Pena Park and Palace, one of the most visited tourist locations in Portugal, during a monitored public event of relevance to the park's management. For the data visualization of the crowding information collected by the two sensors, a customized Grafana dashboard was created, based on templates defined for different monitoring scenarios, and is presented in Fig. 5. This dashboard includes temporal graphs that show crowd variation over time, allowing users to observe fluctuations in crowding during specific periods, as well as interactive maps displaying the real-time location of each sensor along with their corresponding crowding levels.

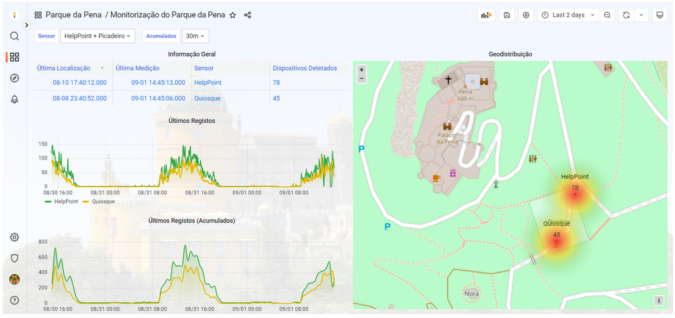


Fig. 5. Dashboard for crowding data visualization at Pena Park.

Ground-truth data were also obtained by manually counting the number of individuals in the vicinity of the monitored areas at 5-minute intervals. To evaluate the effectiveness of the fingerprinting and counting methodology, statistical analysis was performed. Pearson correlation coefficients between the sensor detections and the actual number of people were calculated for each sensor, revealing strong positive correlations: 0.9448 and 0.8502 for the Kiosk and Help-point sensors, respectively. These results demonstrate a high degree of accuracy in the system's ability to estimate relative crowding levels.

In addition, the Dynamic Time Warping (DTW) algorithm was applied to assess the temporal similarity between the sensor detections and the actual number of people over time. The DTW distances obtained were 1.084 for Kiosk and 1.924 for Help-point, confirming good alignment between the trends in the estimated and actual counts, even in the presence of minor timing offsets or variability in crowd flow.

The superior performance of the Kiosk sensor can be directly attributed to its favourable placement conditions and better RF propagation, highlighting the critical importance of sensor calibration tailored to the specific characteristics of each deployment site. This reinforces the need for adaptive calibration mechanisms that adjust sensor sensitivity and detection thresholds according to environmental factors, ensuring consistent and accurate crowd estimation across diverse urban scenarios.

Adaptive Calibration via UAV Support

Crowd sensor placement critically affects accuracy. The work presented here, developed under the MoniCrowd project, [8], aims at the integration of a UAV-assisted calibration strategy that includes signal scans across a grid over the deployment area. Based on collected data, a sensitivity map is computed per sensor, and optimal detection thresholds are derived using optimization algorithms.

To obtain the sensitivity map, the project is implementing a customized DJI Mavic 3T drone that flies over these points, using Real-Time Kinematics (RTK) to achieve centimetre-level positioning accuracy. Onboard, the drone carries a Google Pixel running GrapheneOS, which gives us better control over MAC address identification and the emission of Wi-Fi probe requests. Sensors, when placed in calibration mode, send timestamped SNR information of Probe Requests sent by the Google Pixel device. This information is stored in the ServerConfigDBMS (PostgreSQL database), to be used later by the *CalibrationEngine*.

Fig. 6 illustrates preliminary information manually collected for the calibration process at 2 pilot deployment sites in the Lisbon region: National Palace and Gardens of Queluz (top row) and Park and Palace of Monserrate (bottom row). For each site, the 3 views shown from left to right correspond to: (1) the monitored location, (2) the spatial matrix defining the drone's signal emission grid of calibration points, and (3) the resulting sensitivity map with the signal reception levels at one sensor.

These maps will then feed into an optimization algorithm, implemented in the *CalibrationEngine* module at the Server, that adjusts sensor sensitivity. The sensitivity of each sensor will be updated over the air (OTA), using downlink communication with the Sensor, ensuring better coverage and minimized detection overlap of crowd counting.

Deployment Costs Estimation

To assess the economic feasibility and scalability of the proposed crowd monitoring solution, a detailed cost estimation was conducted, as summarised in Table II. The analysis considers multiple deployment scenarios with varying numbers of sensors (1, 10, 20, and 30), reflecting both small-scale and large-scale applications. The total cost for each scenario includes hardware components, protective cases, communication fees, and server infrastructure.

Core sensor hardware (€179.23 per unit) includes a Raspberry Pi 5 (4GB) with active cooling (SC1148), an Alfa Net-



Fig. 6. Calibration setup at Queluz (top) and Monserrate (bottom): location view, drone emission grid, and resulting sensitivity map.

TABLE II
TOTAL COSTS ESTIMATION FOR DIFFERENT SETUPS

# Sensors	1	10	20	30
Sensor Power	€90.75	€907.50	€1 815.00	€2 720,50
Sensor Hardware	€179.23	€1 792.30	€3 584.60	€5 376.90
Sensor Cases	€90.00	€900.00	€1 800.00	€2 700.00
Communications	€0.98	€9.80	€19.60	€29.40
Server Hosting	€72.98	€72.98	€72.98	€72.98
Total = 1st Year	**€432.96**	**€3 681.60**	**€7 291.20**	**€10 900.80**
Total > 1st Year	**€72.98**	**€81.80**	**€91.60**	**€101.40**

work AWUS036ACH Wi-Fi adapter for probe request capture, a Raspberry Pi IoT LoRaWAN pHAT for additional connectivity, a UART GPS NEO-8M module for location metadata, and a RTC Backup Battery (SC1163). Each unit may also have a Geekworm X1202 UPS and four Samsung INR18650-35E batteries to support off-grid deployment, adding an extra cost of €90.75 per sensor. Weather-resistant enclosures, for outdoor event scenarios, cost €90 per unit.

Communication costs are calculated based on expected MQTT uplink and downlink data traffic over LoRaWAN, using Helium's public pricing model [27], yielding an average of €0.98 per sensor annually. Server hosting is based on a yearly subscription for a VPS used to operate the Crowding Server Unit, currently priced at €72.98.

The first-year total combines the one-time hardware and case investments with ongoing communication and server expenses. In subsequent years, only the operational costs (communication and server) apply, significantly reducing the maintenance burden. This cost analysis demonstrates that the solution remains economically viable even at larger scales, supporting its applicability in diverse urban event scenarios.

V. CONCLUSION AND FUTURE WORK

This paper presented an adaptive, privacy-compliant system for real-time crowd monitoring in urban settings, tailored for temporary events. The proposed architecture combines self-configuring Wi-Fi sensing units, multi-RAN connectivity, and a UAV-assisted calibration mechanism. A robust fingerprinting method based on detailed Information Elements analysis was developed and validated in controlled conditions by ranking first in the CONFRONT international challenge on Wi-Fi frame fingerprinting. The system was also tested in real-world conditions at a UNESCO heritage site (Pena Park), demonstrating strong correlation with manual crowd counts, even under MAC address randomization.

In addition to its technical design, the system offers practical scalability, supported by cost-effective components and infrastructure. The modular architecture and remote management features enable rapid deployment with minimal technical overhead, making it suitable for municipalities and civil protection services. A key contribution was also the design and construction of a controlled test environment using a Faraday cage, which enables the creation of a high-quality labeled dataset. This dataset not only supports the current rule-based fingerprinting approach but also lays the groundwork for the training and validation of future machine learning models.

Future work will focus on: (i) extending the fingerprinting engine using semi-supervised learning techniques such as Siamese networks to improve generalization to evolving device patterns; (ii) completing the automation of the UAV-based calibration process; and (iii) validating the system during large-scale public events.

Acknowledgments

This work was funded by the MoniCrowd project supported by measure "RE-C05-i08.M04 – "Support the launch of a program of R&D projects aimed at the development and implementation of advanced cybersecurity, artificial intelligence and data science systems in public administration, as well as a scientific capacity building program", of the Recovery and Resilience Plan – PRR, within the framework of the financing contract signed between the Recuperar Portugal Mission Structure (EMRP) and the Portuguese Foundation for Science and Technology.

References

[1] UNWTO, "Number of international tourist arrivals worldwide from 2005 to 2024, by region (in millions)," 2025. [Online]. Available: https://www.statista.com/statistics/186743/international-tourist-arrivals-worldwide-by-region-since-2010/

[2] O. Tokarchuk, J. C. Barr, and C. Cozzio, "How much is too much? Estimating tourism carrying capacity in urban context using sentiment analysis," *Tourism Management*, vol. 91, p. 104522, 2022.

[3] M. Biendicho, A. Derqui, O. Puig, and C. Estève, "Tourists Go Home! Examining Antitourism in Barcelona from an Emotions Perspective," *Tourism Culture & Communication*, vol. 22, no. 3, pp. 275–295, 2022.

[4] H. Seraphin, A. King, D. Pilato, and L. Sheeran, "Over-tourism and the fall of Venice as a destination," *Journal of Destination Marketing & Management*, vol. 9, pp. 374–376, 2018.

[5] F. Brito e Abreu, R. Neto Marinheiro, I. Boavida-Portugal, A. Lopes, T. Mestre dos Santos, D. Sampaio de Almeida, and R. Simões, "A digital transformation approach to scaffold tourism crowding management: pre-factum, on-factum, and post-factum," in *Proc. of the Joint Int. Conf. ECTI DAMT & NCON*. IEEE, 2024, pp. 586–591.

[6] RESETTING project, "RESETTING: Smart tourism tools for overtourism monitoring," 2025. [Online]. Available: https://sites.google.com/iscte-iul.pt/resetting-project

[7] T. Mestre dos Santos, R. Neto Marinheiro, and F. Brito e Abreu, "Wireless Crowd Detection for Smart Overtourism Mitigation," in *Smart Life and Smart Life Engineering: Current State and Future Vision*, ser. LNBIP, E. Kornyshova, R. Deneckère, and S. Brinkkemper, Eds. Springer, 2025, pp. 237–258.

[8] R. Neto Marinheiro, "Project MoniCrowd: Adaptive System for Crowding Monitoring Using User's Devices Fingerprinting," Funded by IACDC – Inteligência Artificial, Ciência dos Dados e Cibersegurança na Administração Pública, 2024, schedule [31 Mar 2025, 30 Jan 2026].

[9] L. Atzori and L. Pintor, "CONFRONT – Challenge ON wifi FRame fingerprinting for people cOunting aNd Tracking," Mar. 2024. [Online]. Available: https://sites.unica.it/net4u/confront-challenge-on-wifi-frame-fingerprinting-for-people-counting-and-trackingconfront/

[10] R. Dias da Silva, R. Neto Marinheiro, and F. Brito e Abreu, "Crowding Detection Combining Trace Elements from Heterogeneous Wireless Technologies," in *Proc. of the 22nd Int. Symposium on Wireless Personal Multimedia Communications (WPMC)*. IEEE, Nov. 2019.

[11] R. Agarwal, S. Kumar, and R. M. Hegde, "Algorithms for Crowd Surveillance Using Passive Acoustic Sensors Over a Multimodal Sensor Network," *IEEE Sensors Journal*, vol. 15, no. 3, pp. 1920–1930, 2015.

[12] A. López-Cifuentes, M. Escudero-Viñolo, J. Bescós, and P. Carballeira, "Semantic-driven multi-camera pedestrian detection," *Knowledge and Information Systems*, vol. 64, no. 5, pp. 1211–1237, 2022.

[13] S. Xu, S. Li, W. Huang, and R. Wen, "Detecting spatiotemporal traffic events using geosocial media data," *Computers, Environment and Urban Systems*, vol. 94, p. 101797, 2022.

[14] S. Park, M. Bourqui, and E. Frias-Martinez, "MobInsight: Understanding Urban Mobility with Crowd-Powered Neighborhood Characterizations," in *Proc. of the 16th Int. Conf. on Data Mining Workshops (ICDMW)*. IEEE, 2016, pp. 1312–1315.

[15] U. Singh, J.-F. Determe, F. Horlin, and P. D. Doncker, "Crowd Monitoring: State-of-the-Art and Future Directions," *IETE Technical Review*, vol. 38, no. 6, pp. 578–594, 2021.

[16] F. D. Meersman, G. Seynaeve, M. Debusschere, P. Lusyne, P. Dewitte, Y. Baeyens, A. Wirthmann, C. Demunter, F. Reis, and H. I. Reuter, "Assessing the Quality of Mobile Phone Data as a Source of Statistics," in *Proc. of the European Conf. on Quality in Official Statistics (Q2016)*, Madrid, Spain, Jun. 2016. [Online]. Available: https://api.semanticscholar.org/CorpusID:43992550

[17] M. Vega-Barbas, M. Álvarez Campana, D. Rivera, M. Sanz, and J. Berrocal, "AFOROS: A Low-Cost Wi-Fi-Based Monitoring System for Estimating Occupancy of Public Spaces," *Sensors*, vol. 21, no. 11, 2021.

[18] T. Bravenec, J. Torres-Sospedra, M. Gould, and T. Fryza, "What Your Wearable Devices Revealed About You and Possibilities of Non-Cooperative 802.11 Presence Detection During Your Last IPIN Visit," in *Proc. of the 12th Int. Conf. on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE, 2022, pp. 1–7.

[19] Y. Cai, M. Tsukada, H. Ochiai, and H. Esaki, "MAC address randomization tolerant crowd monitoring system using Wi-Fi packets," in *Proc. of the 16th Asian Internet Engineering Conf. (AINTEC)*. ACM, 2021, p. 27–33.

[20] T. He, J. Tan, and S.-H. G. Chan, "Self-Supervised Association of Wi-Fi Probe Requests Under MAC Address Randomization," *IEEE Transactions on Mobile Computing*, vol. 22, no. 12, pp. 7044–7056, 2023.

[21] M. Uras, R. Cossu, E. Ferrara, O. Bagdasar, A. Liotta, and L. Atzori, "WiFi Probes sniffing: an Artificial Intelligence based approach for MAC addresses de-randomization," in *Proc. of the 25th Int. Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD)*. IEEE, 2020, pp. 1–6.

[22] M. Uras, E. Ferrara, R. Cossu, A. Liotta, and L. Atzori, "MAC address de-randomization for WiFi device counting: Combining temporal- and content-based fingerprints," *Computer Networks*, vol. 218, p. 109393, 2022.

[23] R. Rusca, D. Gasco, C. Casetti, and P. Giaccone, "Privacy-preserving WiFi fingerprint-based people counting for crowd management," *Computer Communications*, vol. 225, pp. 339–349, 2024.

[24] T. M. Santos, "Smart tourism toolkit for crowd-monitoring solutions," Master's thesis, Iscte – Instituto Universitário de Lisboa, 2023. [Online]. Available: http://hdl.handle.net/10071/29505

[25] L. Pintor and L. Atzori, "A dataset of labeled device Wi-Fi probe requests for MAC address de-randomization," *Computer Networks*, vol. 205, p. 108783, 2022.

[26] L. Yuriev, "t1ha – Fast Positive Hash," 2024. [Online]. Available: https://github.com/erthink/t1ha

[27] Helium Foundation, "Data Transfer Calculator — Helium Documentation," 2024. [Online]. Available: https://docs.helium.com/tokens/data-credit/#data-transfer-calculator