Automated Classroom Attendance in Educational Environments Using a BLE-Based Context-Aware Middleware

Emanuel de Franceschi Vieira[1],
Tiago Antônio Rizzetti[1], **Renato Preigschadt De Azevedo**[1]

[1]Colégio Técnico Industrial de Santa Maria
Universidade Federal de Santa Maria (UFSM)
Av. Roraima nº 1000. Cidade Universitária. Camobi. Santa Maria - RS

emanuel.franceschi@acad.ufsm.br,
tiago.rizzetti@ufsm.br, renato@redes.ufsm.br

Abstract. *This paper presents a context-aware middleware for educational environments that uses Bluetooth Low Energy (BLE) signals to infer room-level presence and support pervasive applications. The architecture integrates ESP32 devices with ESPresense firmware, MQTT communication, Home Assistant for contextual interpretation, and a Python-based middleware responsible for filtering, persisting, and exposing contextual events through a REST API. To validate the proposal, an automated classroom attendance application was implemented on top of this infrastructure. The evaluation was conducted through three controlled simulation runs of 60 minutes each, involving 30 synthetic students and pseudo-random mobility patterns. Results showed stable operation, no WebSocket failures, and coherent attendance behavior across arrival, stay, and departure phases. Although 21,600 MQTT messages were published per run, only 14.38% generated persisted events, indicating effective noise reduction by the middleware compared to a raw-event persistence baseline. During the stay phase, the average attendance reached 27.43 out of 30 students (91.43%). The paper also discusses privacy, deployment cost, hardware portability, and limitations of BLE-only presence sensing in real classrooms.*

1. Introduction

One of the central goals of ubiquitous and pervasive computing is to integrate computational capabilities into the physical environment in a continuous, distributed, and minimally intrusive manner [Weiser 1991]. From this perspective, sensors, connected devices, event-driven services, and adaptation mechanisms make it possible to build spaces capable of perceiving contextual changes and reacting to them in a timely way [Schilit et al. 1994]. Rather than relying exclusively on explicit user interaction, such environments incorporate processes of observation, interpretation, and response based on information collected from the context of use itself [Abu-Rasheed et al. 2023].

According to [Kaur et al. 2022], among the different application scenarios of pervasive computing, educational environments constitute a relevant domain. Educational institutions bring together multiple users, shared spaces, intense circulation flows, and dynamic routines, which makes infrastructures capable of sensing presence, occupancy, and movement across environments become particularly valuable. When such information is treated as computational context, it becomes possible to develop adaptive services to support institutional management, environmental monitoring, and the provision of context-aware applications.

In this scenario, the automated identification of presence in indoor environments constitutes a practical and recurring problem. In educational institutions, especially those with large flows of people and multiple physical spaces, monitoring occupancy and presence at the room level can support administrative, operational, and pedagogical processes. However, commercial tracking and automation solutions often present limitations related to cost, dependence on proprietary infrastructure, difficulty of customization, and low adherence to the specific needs of public educational institutions. In addition, a large part of the literature focuses primarily on achieving higher indoor localization accuracy, while dedicating less attention to the construction of reusable architectures to support context-based pervasive applications.

At the same time, the evolution of open and low-cost technologies, such as ESP32 microcontrollers, MQTT communication, Bluetooth Low Energy (BLE) sensing, and event-driven automation platforms, expands the possibilities for deploying ubiquitous infrastructures in real environments. In particular, room-level presence identification by proximity, although it does not provide precise metric positioning, may be sufficient for a broad class of context-aware applications in educational environments. In such cases, more important than knowing the exact coordinates of an individual is determining, with consistency, in which environment that person is located and how this information can be transformed into useful contextual events.

This paper presents a context-aware middleware for educational environments based on BLE signal detection by ESP32 devices running ESPresense firmware, asynchronous communication via MQTT, intermediate interpretation through the Home Assistant platform, and contextual processing in a dedicated layer developed in Python. The proposal uses indoor localization not as an end in itself, but as a ubiquitous sensing mechanism to infer presence and context changes in real time. From this infrastructure, location events are organized, filtered, persisted, and made available through a REST API, allowing their consumption by applications decoupled from the data collection layer.

The core of the proposal lies in transforming presence readings into higher-level abstractions suitable for the development of context-sensitive services. Instead of directly exposing RSSI details, MQTT topics, or raw sensor states, the middleware provides information such as room presence, transitions between spaces, and occupancy history. This approach reduces the coupling between the physical sensing infrastructure and the final applications, favoring modularity, reuse, and scalability. In addition, the solution incorporates authentication and role-based access control (RBAC) mechanisms, which are important when dealing with data associated with the presence of individuals in educational environments.

As a form of validation, this paper uses the proposed infrastructure to implement and evaluate an automated classroom attendance application. In this application, the last valid location of each beacon associated with a student is used to infer presence in a specific class, allowing location events to be transformed into a concrete functionality of educational interest. Thus, the experimental focus of this paper is not on estimating highly accurate coordinates, but on verifying whether the proposed middleware is capable of sustaining a context-based pervasive application with coherent behavior and stable operation.

From a scientific perspective, the main contribution of this paper lies in the development and validation of an extensible middleware to support context-aware pervasive applications in educational environments, taking automated attendance as a case study. This middleware transforms location readings into higher-level contextual events, allowing applications decoupled from the collection infrastructure to consume information such as room presence, transitions between spaces, and occupancy history. In this way, the proposal demonstrates how open and accessible technologies can be integrated to build a context-aware computing foundation with low cost, replication potential, and adherence to real institutional scenarios.

Thus, the contributions of this paper can be summarized in four main points: *(I)* the development of a modular middleware for context perception in educational environments, integrating BLE sensing, MQTT communication, contextual interpretation through Home Assistant, and a dedicated processing layer; *(II)* the definition of an event abstraction mechanism capable of interpreting location readings, filtering redundant updates, persisting semantic transitions, and providing contextual information through a REST API; *(III)* the implementation of an automated classroom attendance application built on top of this middleware; and *(IV)* the experimental validation of the solution through simulated classroom mobility scenarios.

The remainder of this paper is organized as follows. Section 2 discusses related work. Section 3 describes the proposed architecture. Section 4 presents the adopted methodology. Section 5 discusses the experimental results. Finally, Section 6 presents the conclusions and future work.

2. Related Work

The literature related to this work can generally be organized into two main research directions. The first focuses on BLE-based indoor localization systems, with emphasis on improving accuracy, handling RSSI, and evaluating positioning techniques. The second includes proposals for smart classrooms and smart campuses, in which sensors and connected devices are used to automate specific functions, such as energy saving, environmental monitoring, and attendance recording. Although both lines are relevant to the present study, few works combine distributed sensing, contextual interpretation, and the standardized provision of events within the same architecture aimed at pervasive applications.

In the domain of intelligent educational environments, [Huang et al. 2019] propose a context-aware smart classroom architecture for smart campuses, structured around a technological integration model and supported by the use of Raspberry Pi as an interface element between devices and applications. The authors argue that the architecture is suitable for building and operating smart classrooms at campus scale and demonstrate its feasibility through an application focused on energy saving. Despite its conceptual proximity to the present work, the proposal is strongly centered on the smart classroom as a functional unit and on integration mechanisms for that scenario.

Also in the educational domain, [Puckdeevongs et al. 2020] present an automated classroom attendance system based on BLE indoor positioning. The proposed architecture is composed of two main blocks: an indoor positioning mechanism within the classroom and an attendance recording system, using RSSI fingerprinting and neural networks to identify student positions in a real environment. The work is relevant because it demonstrates that low-cost BLE technologies can be applied in real academic scenarios, including under interference conditions typical of classrooms. However, its focus remains associated with a specific application, namely attendance control.

From a perspective closer to IoT systems for indoor localization, [Teran et al. 2017] describe a modular BLE-based architecture composed of two main subsystems: acquisition and central server. The proposal includes modules for measurement, aggregation, and data transmission, storage, web visualization, and cloud services, employing an RSS-based approach to identify reference zones in indoor environments. This work is close to the present study because it emphasizes architectural aspects and treats indoor localization as part of a broader IoT ecosystem. Even so, the authors' proposal remains focused on identifying location in indoor environments and on validating an IoT localization solution, without advancing more explicitly toward a context interpretation layer.

In addition to these closely related works, recent reviews on BLE-based indoor localization indicate that most of the literature is still oriented toward achieving higher accuracy, especially through machine learning techniques, filtering, sensor fusion, and newer features of the BLE standard. This panorama reinforces that, although indoor localization is a widely explored topic, there is still room for contributions that prioritize less the achievement of precise coordinates and more the construction of reusable infrastructures for context-aware computing in real environments [Morgan 2024].

In this context, the main distinguishing feature of the present work lies in treating indoor localization not as an end in itself, but as a ubiquitous sensing mechanism for context perception in educational environments. Rather than proposing only a positioning system or an isolated application, this work presents a modular architecture that integrates BLE sensing, asynchronous communication via MQTT, contextual interpretation mediated by Home Assistant, and a dedicated middleware responsible for filtering, persisting, and exposing events through a REST API. In this way, the proposal broadens the scope of solutions found in the literature by offering a reusable architecture for different pervasive applications, such as occupancy monitoring, contextual alerts, presence management, and adaptive automation of physical resources in the environment.

In summary, the reviewed works contribute to the advancement of smart classrooms, automated attendance systems, and BLE-based indoor localization solutions. However, most of them remain focused either on specific applications, such as attendance control, or on positioning and IoT integration mechanisms primarily aimed at location identification. In contrast, the present work focuses on a middleware capable of abstracting location readings and making them available as reusable contextual events for applications decoupled from the collection infrastructure. In this sense, automated attendance is adopted not as the main contribution, but as a case study to validate this proposal.

3. System Architecture

The proposed architecture was conceived as a context perception infrastructure for educational environments, in which indoor localization is used as a ubiquitous sensing mechanism to infer room-level presence. The data flow begins at ESP32 locators running ESPresense firmware, which detect BLE signals emitted by beacons and publish these readings via MQTT to a Mosquitto broker.

In general terms, the system is organized into five main layers: *(A)* a data collection layer, responsible for detecting BLE signals in the environment; *(B)* a communication layer, responsible for the asynchronous transport of the collected data; *(C)* a contextual interpretation layer, in which raw events are converted into presence states; *(D)* a middleware layer, responsible for organizing, persisting, controlling access to, and exposing the information; and *(E)* an application layer, composed of services that consume contextual events and implement adaptive behaviors. Figure 1 illustrates the system architecture.

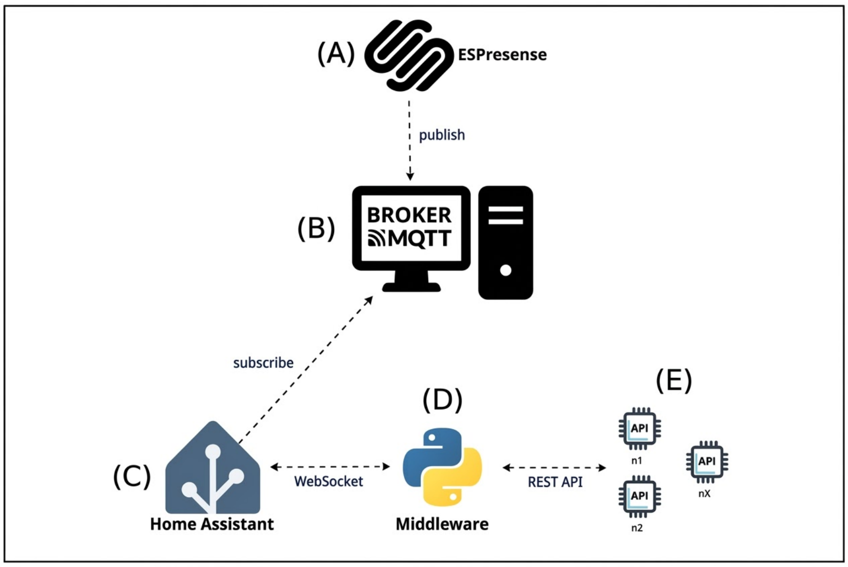

Figure 1. System Architecture

In the contextual interpretation layer, Home Assistant consumes MQTT topics and uses the `mqtt_room` component to consolidate the state of each beacon, producing `state_changed` events whenever a relevant change of environment occurs. This design choice reduces the coupling between the physical capture infrastructure and the final applications, while also simplifying the incorporation of new sensors into the monitored environment.

On top of this foundation, a middleware in Python with FastAPI was developed to maintain a persistent WebSocket connection with Home Assistant, filter redundant events, persist semantic transitions in a MySQL database, and maintain the current state of each beacon. The processing flow is event-driven: *(i)* only `state_changed` events related to beacon entities are accepted; *(ii)* each valid event updates the beacon snapshot timestamp (`last_seen_at`); *(iii)* persistence in `LocationEvents` occurs only when room state actually changes; *(iv)* transition semantics are encoded as `enter`/`leave` events. Since the collection infrastructure may produce successive updates caused by RSSI oscillations, estimated distance variations, and repeated state publications, the middleware does not persist every received update. Instead, it records only effective room changes, reducing redundant writes while preserving contextual transitions required by upper-layer applications. The middleware exposes the data through a REST API protected by JWT and role-based access control, allowing clients to query location history, current occupancy, and contextual events without directly accessing the low-level infrastructure components.

This modular organization made it possible to develop applications decoupled from the sensing layer. In this paper, the application used for validation was an automated attendance service, which queries the last valid location of the beacons associated with students and classifies each student as present or absent in a specific room. In this way, the proposed middleware no longer acts merely as support for tracking, but instead sustains context-oriented pervasive services in educational environments.

4. Methodology

The evaluation presented in this paper aimed to validate the feasibility of the proposed middleware as the basis for an automated attendance application. Rather than measuring metric positioning accuracy, the experiment focused on verifying whether the middleware is capable of sustaining a context-sensitive application by processing presence changes in real time and reflecting the expected behavior of a class throughout a lesson. Thus, the experimental campaign was designed to evaluate the proposal in terms of functional consistency, operational stability, and adherence to the configured mobility profile. To this end, a controlled simulation campaign was conducted, with pseudo-random movements of BLE beacons associated with synthetic students previously registered in the database.

The validation application was implemented on top of the same middleware described in Section 3. The automated attendance service uses the last known location of each beacon associated with the students in a class and considers a student present only if the corresponding beacon has a valid and recent location in the analyzed room. The adopted validity window was 180 seconds, which was sufficient to tolerate small variations in detection without turning old readings into current presence. The attendance function was made available both through an API and through a command-line interface, but the experimental campaign used the CLI for periodic execution and automated recording of results.

The experiments were orchestrated through simulation and data collection scripts. The analyzed campaign consisted of three independent repetitions, each lasting 60 minutes and involving 30 beacons associated with a synthetic class. This number was adopted because it represents a scenario compatible with the size of a real class in an educational environment, allowing the automated attendance application to be evaluated under a plausible simultaneous occupancy level. In each repetition, the simulator published MQTT messages every 5 seconds, recalculated mobility every 30 seconds, and executed the automated attendance procedure every 60 seconds, totaling 60 samples per run. Table 1 summarizes the general configuration of the experimental campaign.

The simulation used active BLE beacon identifiers in iBeacon-compatible format (e.g., `iBeacon:sim-XXXX`). As this campaign was synthetic, there was no physical classroom area, no physical ESP32 placement, and no RSSI-to-distance calibration step. Distance estimation was intentionally delegated to the ESPresense/Home Assistant stack, while this paper focused on room-level state transitions and middleware behavior.

Table 1. Experimental campaign configuration

Parameter	Value
Repetitions	3
Duration per repetition	60 min
Simulated beacons	30
MQTT publication interval	5 s
Mobility interval	30 s
Automated attendance interval	60 s
Signal validity window	180 s

A mobility model was implemented to reproduce classroom behavior in three phases: arrival, stay, and departure. The simulation adopted durations of 10, 35, and 15 minutes for these phases, respectively, totaling the 60 minutes of the scenario. Since the campaign considered only one monitored room, the possible states were reduced to presence in the room or absence. Transition probabilities were heuristically defined in the simulator in order to produce a plausible dynamic of gradual entry, class stabilization throughout the lesson, and gradual departure at the end of the period. Thus, during the arrival phase, beacons outside the room had a 55% probability of entering, while beacons already inside had a 95% probability of remaining there. During the stay phase, the probability of a beacon remaining in the room increased to 97.5%, while absent beacons had a 25% chance of returning. During the departure phase, this probability dropped to 84%, modeling the gradual emptying of the room at the end of the class.

The analysis considered two groups of indicators. The first group evaluated the internal behavior of the middleware and its integration with the infrastructure, including the total number of events persisted in the `LocationEvents` table, the availability of the WebSocket connection with Home Assistant, and the relationship between published MQTT messages and events effectively stored. It should be noted that the ESPresense/MQTT-based collection layer had already been evaluated by the authors in a previous performance study, in which latency, reliability, and scalability metrics were analyzed in a real environment [Vieira and Rizzetti 2025]. The second group evaluated the behavior of the application built on top of the middleware, considering the number of present students per collection, the mean attendance across repetitions, and the temporal distribution of results during the arrival, stay, and departure phases.

To provide a baseline for comparison, a *raw persistence baseline* is defined in which every MQTT update is stored as a historical record. In the evaluated campaign, this baseline corresponds to 21,600 stored records per run. The proposed middleware is compared against this baseline by quantifying how many records are actually persisted after semantic filtering.

The experimental workflow relied on dedicated scripts for data preparation, simulation, roll-call execution, and campaign orchestration. Since this was a controlled simulation without external observation in a real classroom, the validation was not conducted as a definitive measure of real-world accuracy, but rather as a verification of functional feasibility and operational stability.

5. Results and Discussion

Each repetition executed 30 beacons over 60 minutes, with MQTT publication every 5 seconds. This resulted in 21,600 published messages per run, corresponding to an average load of 6 messages per second. Despite this volume, the middleware persisted an average of 3,105 events per repetition (standard deviation of 53.78), which corresponds to approximately 0.863 events per second. In other words, only 14.38% of the published messages generated persisted events, while 85.62% were absorbed without generating new historical records. This result is relevant because it highlights the middleware's ability to act as an abstraction layer between raw data collection and application-level data consumption.

This result confirms, in operational terms, the strategy described in the system architecture: although the infrastructure receives frequent updates associated with RSSI oscillations, estimated distance variations, and state republishes, the middleware continuously maintains an up-to-date location snapshot and persists only effective room changes, such as entries and exits. From an architectural perspective, this behavior is relevant because it reduces the volume of historical data without compromising the contextual information required by the applications.

Middleware operation also remained stable throughout the simulations. None of the three repetitions recorded WebSocket connection errors with Home Assistant during the executed runs. This result indicates that, in the evaluated scenario, the architecture was able to sustain event ingestion and concurrent execution of the automated attendance application without interruptions in the real-time event channel. Although the considered load is still moderate, the absence of communication failures during three 60-minute runs is a favorable indication of robustness for the intended class of applications. Table 2 summarizes the main results observed across the three campaign repetitions.

Table 2. Consolidated results of the three repetitions

Indicator	Rep. 1	Rep. 2	Rep. 3	Average
Persisted events	3133	3043	3139	3105.00
Overall mean attendance per collection	24.03	24.12	24.08	24.08
Mean attendance after the first collection	24.15	24.24	24.20	24.20
Mean absences per collection	5.97	5.88	5.92	5.92
WebSocket errors	0	0	0	0

Regarding the automated attendance application, the results coherently followed the behavior defined by the `classroom` mobility model. Considering the entire window after the first collection of each repetition, the mean attendance was 24.20 students out of 30 simulated beacons, with very low variation across runs (standard deviation of 0.05), which corresponds to 80.66% of the class. This aggregate value, however, combines moments of arrival, stay, and departure. When the results are observed by phase, the interpretation becomes more informative.

During the arrival phase, the attendance procedure recorded an average of 25.77 students present per collection, corresponding to 85.89% of the class, reflecting the gradual entry of students into the monitored environment. During the stay phase, which represents the most important period for the attendance functionality, the average increased to 27.43 students present per collection, or 91.43% of the class, with minimal variation across repetitions (standard deviation across runs of 0.03) and values ranging from 23 to 30 present students in each collection. During the departure phase, the average dropped to 15.13 students present per collection, or 50.44%, reproducing the gradual emptying of the room at the end of the lesson. Taken together, these results show that the application not only remained operational under load, but also adequately reflected the expected temporal dynamics of the simulated educational scenario. Table 3 summarizes the automated attendance results across the different phases of the simulated mobility.

Table 3. Automated attendance results by phase

Phase	Mean present students per collection	Mean percentage
Arrival	25.77	85.89%
Stay	27.43	91.43%
Departure	15.13	50.44%

Figure 2 complements this analysis by illustrating the temporal evolution of attendance in a representative repetition of the campaign, highlighting the transitions between the arrival, stay, and departure phases.

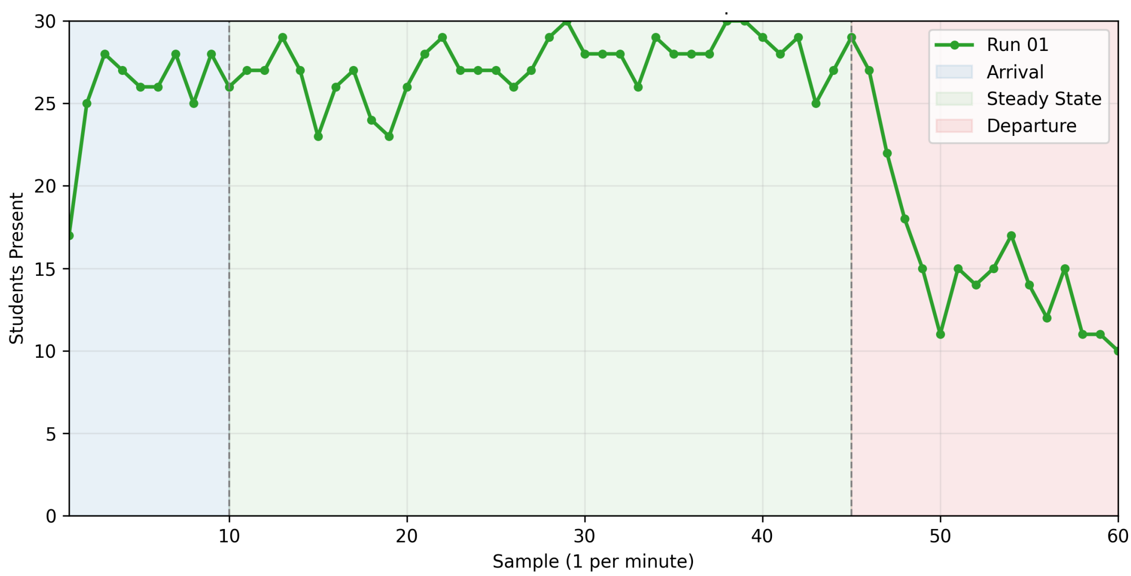

Figure 2. Attendance over time in a representative run.

Another relevant result concerns the reasons for absence. Throughout the campaign, the recorded absences were compatible with the adopted mobility model, being concentrated in the arrival period and, especially, in the departure period. Since the scenario considered only one monitored environment, students outside the room were treated as absent, which explains the gradual reduction in attendance observed at the end of the simulated lesson. In addition, `stale_signal` and `other_room` remained null in the analyzed runs, indicating that absences were dominated by modeled out-of-room behavior (`unknown_room`) rather than timeout artifacts.

The baseline comparison also reinforces the role of the middleware: persisting only 3,105 events on average, instead of 21,600 raw updates, represents an 85.62% reduction in historical write volume while preserving room transition semantics. This behavior directly addresses message redundancy and database pressure, and is aligned with edge-oriented deployment in which processing remains in the local institutional network.

Despite these results, BLE-only attendance has known practical limitations in real use. A student may leave the beacon in the room and physically leave, which may produce false positives. Therefore, in high-assurance scenarios the proposed architecture should be seen as one contextual signal among others, and can be extended with multimodal sensing (e.g., camera or access-control events) for cross-validation.

Privacy and governance are also central in educational deployments. Although this study used synthetic identities, real operation should adopt data minimization (room-level events instead of continuous coordinates), strict role-based access, short retention windows for sensitive traces, and clear institutional policies with transparency to students.

Regarding scalability and portability, the architecture is hardware-flexible because the middleware consumes standardized events (`state_changed`) instead of vendor-specific firmware internals. Thus, equivalent BLE gateways or other event producers can be integrated as long as they publish compatible room-state abstractions. For larger classes and distinct mobility patterns, the same orchestration scripts already support campaigns with higher numbers of synthetic students/beacons, enabling progressive stress tests before real deployment.

In summary, the results show that the proposed middleware was capable of sustaining an automated attendance application based on contextual information, with reproducible behavior and no observed communication failures during the campaign. At the same time, the evaluation also highlights the limits of the study. The campaign was conducted with synthetic mobility, a single monitored room, and 30 beacons, so the results should not be interpreted as a final measure of performance in a real deployment. Still, for the purposes of this paper, they are sufficient to demonstrate that the proposed middleware not only organizes and provides contextual events, but also effectively supports a concrete pervasive application of educational interest.

6. Conclusion

This paper presented a context-aware middleware for educational environments, based on BLE signal detection, MQTT communication, contextual interpretation through Home Assistant, and a dedicated layer responsible for filtering, persisting, and exposing location events. On top of this infrastructure, an automated attendance application was developed to infer room-level presence from the last valid location of the beacons associated with students.

The validation was conducted through a controlled simulation campaign composed of three 60-minute repetitions with 30 beacons and pseudo-random mobility. The results showed middleware operational stability, absence of WebSocket connection failures during the experiments, and coherent behavior of the automated attendance application throughout the arrival, stay, and departure phases. In addition, the difference between the volume of published MQTT messages and the number of effectively persisted events showed that the middleware contributes to reducing operational noise and transforming raw readings into semantically relevant events for the applications.

In summary, the results indicate that the proposal is capable of sustaining a concrete pervasive application of educational interest, using indoor localization as a mechanism for context perception rather than merely as a positioning problem. More specifically, the study showed that the developed middleware is capable of abstracting location readings, organizing them as contextual events, and making them available in a reusable way for applications decoupled from the collection infrastructure.

From a practical perspective, the solution is viable for educational institutions because it is based on open components and low-cost hardware (BLE tags, ESP32 gateways, MQTT broker, and open-source middleware stack), with incremental deployment potential by room or building. However, real adoption should include privacy controls and multimodal validation in contexts where BLE-only presence is insufficient.

As future work, future work includes experiments in real classrooms with ground-truth observation and explicit accuracy/error metrics, as well as the extension to scenarios with multiple monitored rooms operating simultaneously and the comparison of the automated attendance application with conventional attendance control procedures.

Acknowledgments

This work was carried out with the support of the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES)

References

Abu-Rasheed, H., Weber, C., and Fathi, M. (2023). Context based learning: a survey of contextual indicators for personalized and adaptive learning recommendations – a pedagogical and technical perspective. *Frontiers in Education*, Volume 8 - 2023.

Huang, L.-S., Su, J.-Y., and Pao, T.-L. (2019). A context aware smart classroom architecture for smart campuses. *Applied Sciences*, 9(9).

Kaur, A., Bhatia, M., and Stea, G. (2022). A survey of smart classroom literature. *Education Sciences*, 12(2).

Morgan, A. A. (2024). On the accuracy of ble indoor localization systems: An assessment survey. *Computers and Electrical Engineering*, 118:109455.

Puckdeevongs, A., Tripathi, N. K., Witayangkurn, A., and Saengudomlert, P. (2020). Classroom attendance systems based on bluetooth low energy indoor positioning technology for smart campus. *Information*, 11(6).

Schilit, B., Adams, N., and Want, R. (1994). Context-aware computing applications.

Teran, M., Aranda, J., Carrillo, H., Mendez, D., and Parra, C. (2017). Iot-based system for indoor location using bluetooth low energy.

Vieira, E. and Rizzetti, T. (2025). Performance evaluation of an espresense/mqtt architecture for real-time locating system. In *Anais da XXII Escola Regional de Redes de Computadores*, pages 68–73, Porto Alegre, RS, Brasil. SBC.

Weiser, M. (1991). The computer for the 21 st century. *Scientific american*, 265(3):94–105.