Received 11 June 2025, accepted 17 June 2025, date of publication 10 July 2025, date of current version 17 July 2025.

Digital Object Identifier 10.1109/ACCESS.2025.3588031

RESEARCH ARTICLE

Cross-Room CO_2-Based Presence Detection for Occupancy Profiling

JELLE VANHAEVERBEKE, EMIEL DEPROST, STEVEN VERSTOCKT, AND SOFIE VAN HOECKE
IDLab, Ghent University-imec, 9052 Ghent, Belgium

Corresponding author: Jelle Vanhaeverbeke (jelle.vanhaeverbeke@ugent.be)

This work was supported in part by the Flanders Innovation & Entrepreneurship (VLAIO) O&O Project iCosy2, and in part by the VLAIO O&O Project Watchdog. The work of Emiel Deprost was supported by the Doctoral Fellowship of the Research Foundation-Flanders (FWO).

ABSTRACT Control systems for building services, such as heating and cooling, often rely on fixed timing schemes. While such approaches are convenient, they make strong assumptions about room usage, often leading to inadequate comfort and energy efficiency. This study addresses this limitation by presenting two contributions aimed at automating the configuration of building control systems. The first contribution involves the development of a presence detection model based on CO_2 data, which is easy to measure and non-privacy intrusive. Unlike existing literature, which typically focuses on single-room applications, this work introduces a dataset and machine learning methodology demonstrating the generalizability of a presence detection model across various real-world rooms, even among different building types. Sliding window normalization of the sensor data is the key to achieve this unsupervised cross-room adaptability. As second contribution, we propose an occupancy profiling technique that relies on the predicted presence information. This approach facilitates the automated configuration of building control systems by using historical presence probabilities to anticipate future occupancy. In contrast to fixed timing schemes, these occupancy profiles dynamically adapt over time, accommodating changes in occupant behavior. As such, this work improves the configuration of building control systems, leading to a more comfortable and energy-efficient environment.

INDEX TERMS Building control, CO_2, occupancy profiling, presence detection.

I. INTRODUCTION

For building control systems, the integration of presence information plays an essential role in achieving both comfort and energy efficiency, especially for the regulation of heating and cooling. However, actual presence information per room is often not available, leading to the use of estimated occupancy schedules or necessitating manual interventions, such as adjusting thermostatic valves. While manual adjustments are straightforward, they are easily overlooked, and fixed programmed schedules often fail to account for the unique occupancy patterns of individual rooms [1].

These shortcomings can be mitigated through more advanced approaches, such as those proposed by Baldi et al. [2], who demonstrated that incorporating dynamic occupancy behavior into HVAC (Heating, Ventilation, and Air Conditioning) control strategies allows systems to adapt to actual needs. This not only minimizes energy waste but also maintains optimal environmental conditions. However, such strategies depend on the availability of accurate room-level occupancy information, which is the focus of this work.

Various techniques exist for detecting human presence in indoor environments, including passive infrared (PIR) sensors, camera-based systems, and radar-based solutions. Each option has its unique characteristics and applicability in different contexts, as detailed in Section II-A.

Within this work, we opt for a different approach, investigating a methodology centered around CO_2-based presence detection. This approach proves viable as humans

The associate editor coordinating the review of this manuscript and approving it for publication was Mostafa M. Fouda.

naturally exhale CO_2 during breathing, which causes a gradual increase in the CO_2 levels of the room. This can be easily measured with cost-effective and commercially available CO_2 sensors, which are often already installed in modern buildings or can be easily retrofitted in older properties. In contrast to motion-dependent sensors like PIR, the CO_2-based approach effectively detects presence even when individuals are sleeping or at rest, as it does not rely on motion. Unlike camera-based systems, CO_2-sensing is not affected by line of sight. Moreover, this proves a non-intrusive approach that does not raise privacy concerns as it only measures the CO_2 levels in the air. Despite these advantages, it is important to note that CO_2-based presence detection has its own set of limitations, including a slower response time and susceptibility to external factors, as discussed in Section VI.

Several prior studies have employed measured CO_2 levels, sometimes in combination with other sensor data, to identify the presence of persons in a room [3], [4], [5], [6], [7], [8], [9], [10]. These works present physical or machine learning models and assess their efficacy in single or multiple (simulated) rooms, demonstrating promising outcomes. However, the ability of these models to generalize to new, unseen spaces is either unexplored or requires fine-tuning and adaptation with labeled data from the new environment to achieve optimal performance. This is significant given the challenges in obtaining accurate presence ground truth data. Consequently, the need to retrain and reconfigure a model for each distinct room poses scalability issues when applied to entire buildings.

This paper advances existing research by extending the applicability of a binary presence detection model from one room to unobserved rooms in a fully unsupervised fashion. Notably, we are the first to use sliding window normalization for CO_2-based presence detection, demonstrating its efficacy in helping the generalization across diverse rooms. Importantly, our evaluation is conducted using real-world data collected from multiple rooms, spanning both office and residential buildings.

In addition, we introduce a novel methodology for room occupancy profiling using the designed presence detection model. This approach aims to characterize the average usage patterns of rooms over time, which can then be used, for instance, to automatically configure the heating schedule. This approach offers the advantage of a dynamically adaptive schedule that aligns with users' evolving behavior, which is particularly beneficial during periods of remote work and flexible hours. Alternatively, the occupancy profile can serve to propose optimized timing schedules to end users, helping them with establishing an effective schedule that aligns with their preferences and routines.

The relevant data and code of this research can be accessed on GitHub at https://github.com/predict-idlab/cross-room-CO2-presence-occupancy for reproducibility and to support future research in this domain.

The main contributions of this work are:

- a CO_2-based presence detection methodology that generalizes to unseen rooms in a fully unsupervised manner, evaluated on real-world data;
- an approach for room occupancy profiling, leveraging CO_2-based presence detection to characterize room usage patterns over time;
- an open-source repository that includes both the dataset and the implementation of our research, ensuring transparency and reproducibility.

The remainder of this paper is structured as follows. First, related work on presence detection and occupancy profiling is summarized and discussed in Section II. Next, Section III gives information on the used data, and how it is collected and preprocessed. Afterwards, the key aspects of the proposed methodology for CO_2-based presence detection and occupancy profiling are introduced in Section IV. This is followed by a validation in Section V of which the results are discussed in three parts. First, the proposed model and features are leveraged for single-room presence detection. Second, the model is validated on unseen rooms and the cross-room presence detection performance is presented. Lastly, two example occupancy profiling approaches are discussed and evaluated. After the results, Section VI takes a deeper look at the limitations of CO_2-based presence detection. Finally, Section VII concludes the work and presents future research possibilities.

II. RELATED WORK

A. PRESENCE DETECTION SENSORS

Various solutions exist for detecting human presence in indoor environments, with one common, widely adopted, method being the passive infrared sensor [11], [12]. PIR sensors are compact, cost-effective electronic devices that identify presence by detecting variations in infrared radiation in their immediate surroundings. Being passive, they do not emit infrared light but solely monitor these variations in radiated heat. PIR sensors are frequently employed as motion-activated light switches. Despite PIR sensors are affordable and easy-to-use, a limitation of these is their reliance on changes in radiated heat, requiring active movement of a heat source nearby [13]. Consequently, PIR sensors may not reliably detect individuals who are stationary, making them less suitable for environments such as offices or living rooms where people often sit still for extended periods.

An alternative category of presence detection solutions involves camera-based systems, encompassing thermal vision [14], [15] or conventional visible-light spectrum cameras [16], [17]. In these systems, a computer vision model processes the camera stream to identify humans in the captured images. This detection yields both presence detection as well as person count output. To safeguard occupant privacy, image processing should be performed on-device, and images should never be transmitted externally. While camera-based solutions offer high accuracy, they are comparatively expensive due to the need for powerful hardware for capturing and processing image data, as opposed to the simpler PIR sensor. Furthermore, acceptance suffers if images are not purely processed on the edge.

Additional techniques for presence detection include radar-based systems, where radio frequency signals are emitted, received, and analyzed [18], [19]. These systems offer several advantages, such as their ability to detect movement through obstacles and in various environmental conditions, but require sophisticated signal processing techniques and are costly to implement.

Acoustic signals can also be utilized, capturing and processing sounds of human origin, such as speech and footsteps [20], [21]. These are generally considered less intrusive than camera's, but still pose privacy issues if the data is not handled carefully. Additionally, acoustic systems might face challenges in noisy environments, leading to potential false detections.

Monitoring energy consumption is another approach, as human presence often correlates with increased energy usage in a room [22], [23]. This method provides indirect and non-intrusive sensing, yet might require infrastructure modifications to measure the energy usage accurately.

Lastly, wireless networks can be leveraged for presence detection, either by relying on the presence of connected smart devices like smartphones or laptops [24], [25] or by directly analyzing the interference of the wireless signal without relying on device connections [26], [27]. Monitoring connected smart devices offers the advantage of utilizing existing infrastructure and can achieve high accuracy. However, it raises privacy concerns as it involves tracking of personal devices. Analyzing the interference of the wireless network closely resembles the radar-based approach, but offers the advantage of using existing hardware infrastructure.

B. CO_2-BASED PRESENCE DETECTION

Numerous studies have explored the detection of presence using CO_2 sensor values, either independently or in combination with other environmental sensor data. A brief overview of select works is discussed below. For a comprehensive overview, readers are directed to existing review studies [28], [29], [30], [31], [32], [33].

Some models leverage physical equations to capture the relationship between the environment and occupants. Typically, the literature within this category employs the mass balance equation which considers the various factors influencing CO_2 levels [34]. Calì et al. propose a dynamic algorithm that optimizes the parameters based on actual CO_2 levels and ground truth presence data [35]. Occupancy is subsequently detected by rewriting the mass balance equation and using the actual CO_2 level as input, yielding promising results across their dataset of five rooms. Nienaber et al. enhance this by refining air exchange rate estimations and expanding the dataset's quality and size [36].

Hybrid approaches extend physical models by incorporating data-driven parameter estimation, enhancing robustness against measurement and model uncertainties. Ebadat et al. propose a nonlinear hybrid model, utilizing Maximum Likelihood Estimation (MLE) to approximate the parameters based on CO_2 and ventilation data [37]. Their work also focuses on adapting the model to other rooms without additional ground truth data, showcasing promising results albeit in a simulated environment. Wolf et al. employ stochastic differential equations to model the mass balance, also using MLE for parameter estimation [38].

Despite their accuracy and interpretability, physical models depend on the correctness of underlying equations and assumptions, which may not always hold in real-world environments. To address these limitations, research has explored learning-based approaches, which excel at capturing complex relationships in a data-driven manner.

For instance, Arief-Ang et al. [3] developed the DA-HOC++ methodology, leveraging CO_2 data. Their model is initially trained on a single-person office and subsequently adapted to new environments using smaller datasets from a cinema, two study zones, and two classrooms. This focus on adaptability aligns with our research objective. While Arief-Ang et al.'s semi-supervised approach demonstrates promising results, it nevertheless requires extra adaptation data for each new room. In contrast, our approach aims to eliminate this requirement, proposing a methodology that generalizes effectively to new environments without the need for additional data.

Looking at more recent studies, Sayed et al. [4] introduced a real-time, edge-based deep learning system for presence detection. Their solution uses a custom sensing board collecting CO_2, motion, sound, and other environmental data, processed with a 2D-CNN architecture to achieve high detection performance. However, they only evaluate their model on the same single-occupant office used for training, leaving its generalizability to other spaces untested.

Similarly, Liang et al. [5] presented an occupancy prediction approach based solely on temperature, humidity, and CO_2 data collected from two similar university classrooms over a three-month period. Among five evaluated techniques, Long Short-Term Memory (LSTM) networks performed best. Nonetheless, their evaluation was limited to the rooms used for training, without investigation into performance on unseen environments.

Karasoulas et al. [6] applied a Hidden Markov Model (HMM) for presence detection, relying exclusively on CO_2 measurements to ensure practicality. Their experiments used data from a publicly available single-office dataset as well as a self-collected, one-month dataset obtained from a 20-desk bank office. Evaluations were performed separately for each office, providing no insight into how well their approach transfers between different rooms.

Kim et al. [7] conducted occupancy detection experiments in a living lab equipped with multiple sensors, including CO_2, temperature, volatile organic compounds (VOCs), as well as system information such as ventilation status and differential pressure. Their models, trained and validated on 55 days of manually labeled data using random forest and neural network algorithms, benefited from the integration of ventilation data. However, their analysis was limited to a single room, offering no evidence regarding the model's ability to generalize to new environments.

In contrast to much of the prior work, our research is aimed at developing a learning-based solution that effectively generalizes to unseen rooms through an unsupervised approach, thereby eliminating the need for additional adaptation data. Our model is trained and evaluated on diverse real-world data covering a wide range of room types, sizes, and occupant behaviors, providing realistic insights into presence detection performance.

C. OCCUPANCY PROFILING

Occupancy profiling can help to enhance the precision of building HVAC and energy consumption simulations or optimize building control system schedules.

Diraco et al. proposed an occupancy profiling methodology utilizing 3D depth cameras to measure occupant count, trajectory, and density [39]. Subsequently, they model the occupancy profiles using inhomogeneous Markov chains which allow the transition probabilities between states to vary with time. However, the deployment of 3D depth cameras is costly, and may face challenges by occlusions and lighting conditions.

In contrast, Kang et al. leveraged mobile positioning data sourced from social media platforms, employing the k-means clustering technique to derive typical weekly occupancy profiles for non-residential buildings [40]. While this approach uses readily available data, it suffers from limitations such as privacy concerns and biases in the social media data, which may not accurately represent the actual occupancy patterns of the environment.

Yang et al.'s work shares similarities with ours, particularly in the use of environmental sensors [41]. They designed a custom sensor box measuring eight room parameters, including CO_2, sound, temperature, and PIR signals. These sensors were mounted at the entrances of single-person offices, accompanied by a camera for ground truth collection. Utilizing this data, pruned decision tree models are trained to predict occupancy status. The resulting occupancy patterns were then analyzed to generate personalized occupancy profiles for both weekdays and weekends. Furthermore, they assessed model generalization by testing it on different offices, demonstrating promising results. However, relying solely on single-person offices for data collection may limit the generalizability of the model to larger spaces with varying occupancy dynamics.

In our study, we focus on a single sensor, specifically CO_2, to offer a cost-effective solution while also preserving privacy. Additionally, we assess the generalizability across various building types within uncontrolled, multi-person environments.

D. CONCLUSION

Numerous studies have presented learning-based CO_2-based presence detection methods, with some addressing the challenge of transferring models to unseen rooms. However, to the best of our knowledge, none of these studies have introduced an approach that can generalize to unseen, real-world rooms in an unsupervised manner. Additionally, our work goes beyond existing research by extending the presence detection approach with an occupancy profiling methodology, aiming to make the configuration of building control systems more accurate and dynamic.

III. DATASET

This research uses two datasets comprising environmental and occupancy data from two different building types. The first dataset was collected in an office setting, while the second dataset was acquired in a residential context. All data is captured in real-world environments, wherein occupants could perform their work and daily activities without any prescribed instructions or constraints. The subsequent sections describe the methodology employed in the collection and processing of these datasets.

A. OFFICE DATA

The office dataset consists of three rooms situated on the tenth floor of a high-rise office building. Two of these rooms, referred to as office L1 and L2 throughout this paper, are large and adjacent offices with a capacity of up to 15 persons each. The third office, designated as office S3, is smaller and designed for individual use.

For environmental monitoring, commercially available sensors from Netatmo[1] are employed to collect data in each office. Given the limited size of the offices, one sensor is installed per office in a central position at desk height. These sensors are configured to record temperature, humidity, CO_2 and loudness at 10-minute intervals. In the context of this study, only the CO_2 measurement will be used.

The dataset annotations were generated using a Steinel HPD2 sensor,[2] which is a camera-based detection system. This sensor employs computer vision to count the number of persons within its field of view. Person detection occurs on-device, ensuring that images are not transmitted to safeguard the privacy of occupants. In offices L1 and L2, two Steinel sensors each are deployed to eliminate blind spots, while the smaller office S3 only requires one sensor. For compatibility with the environmental data, the occupancy data is resampled to a 10-minute interval by taking the mode. As we want to perform presence detection and do not need the exact amount of people in the room, the people count is transformed into a binary presence signal.

[1] https://www.netatmo.com/smart-indoor-air-quality-monitor

[2] https://www.steinel.de/en/sensors-professional/products/series-sensors-professional/hpd/hpd2-033200.html

The data collection within the three offices occurred concurrently, yielding an equal quantity of samples, as illustrated in Table 1.

TABLE 1. Number of samples and presence distribution per room in the dataset.

Room	Sample count	No presence / presence distribution
Office L1	30234 (± 210 days)	82.2% / 17.8%
Office L2	30234 (± 210 days)	81.9% / 18.1%
Office S3	30234 (± 210 days)	92.8% / 7.2%
Home 1	25917 (± 180 days)	82.2% / 17.8%

B. RESIDENTIAL DATA

The residential dataset comprises data collected from a single room within a house, specifically, the living room featuring a large window leading to an outdoor terrace. The environmental parameters are captured using a proprietary sensor capable of measuring temperature, humidity, CO_2, total volatile organic compounds (TVOC) and brightness. For our analysis, only the CO_2 data from this dataset is retained. The sensor is installed against the wall, away from the sliding window.

In addition to the environmental data, presence annotations are obtained using a radar sensor. This sensor is also fixed to the wall and has a small blind spot in the sitting area.

To ensure compatibility with the office dataset, all data is resampled at 10-minute intervals using the mean of the CO_2 measurements and the mode of the presence. However, it is worth mentioning that the logging for this particular room started at a different time and had a slightly shorter runtime, as shown in Table 1.

C. DATA IMBALANCE

In the final column of Table 1, the distribution of the *presence* and *no presence* classes for each room in the dataset is presented. Notably, the *no presence* class occurs on average seven times more than the *presence* class. This distribution aligns with our expectations.

For the two large offices, a standard 40-hour work week would suggest a presence rate of approximately 24%. However, in reality, this estimation is expected to be lower due to factors such as flexible working hours and holiday periods. Therefore, the observed presence distribution of about 18% appears reasonable. In contrast, the small office is utilized less frequently, typically by a single occupant with a schedule filled with numerous meetings, resulting in less time spent within the office room. Regarding the living room, assuming an average occupancy of approximately 4 hours per day, the expected presence rate comes to around 17%, closely corresponding with the measured presence rate.

Given the imbalanced nature of the data, balanced accuracy (BA) is chosen as the metric for validating the methodology in this work. Since this is the average between sensitivity and specificity, it offers a more insightful evaluation compared to metrics such as accuracy.

D. DATA AVAILABILITY

The office dataset is accessible for academic research purposes through the GitHub repository provided at https://github.com/predict-idlab/cross-room-CO2-presence-occupancy. However, the residential data, collected through collaboration with a private company, is restricted from publication.

IV. METHODOLOGY

The methodology proposed in this study comprises of two components. Firstly, a presence detection pipeline based on CO_2 is introduced. This approach is further expanded to incorporate additional temporal information and improve its applicability across unseen rooms. Secondly, the developed presence detection approach is used to obtain occupancy profiles for a given space.

A. CO_2-BASED PRESENCE DETECTION

1) MACHINE LEARNING PIPELINE

Given the limited size of the dataset, our approach focuses on the usage of traditional machine learning techniques. Moreover, it has been shown that for time series classification tasks, deep learning models do not necessarily perform better than traditional approaches when good features are employed [42]. Therefore, the first part of the pipeline extracts meaningful features from the time series sensor data. These extracted features aim to capture certain characteristics and temporal aspects of the data, with the objective to inform the model to make a good prediction.

In particular, a strided window feature extraction methodology is employed, wherein features for each window of interest are derived. This procedure is performed using the tsflex Python library, enabling convenient definition of features, window sizes, and strides [43].

Given that the whole dataset is resampled at 10-minute intervals, the stride of this feature extraction is set to 10 minutes as well. Figure 1 illustrates the extraction of three features for every 10-minute prediction window. Firstly, the CO_2 value of the current window is used. Subsequently, the mean and slope are computed for a 1-hour window based on historical data. These three base features provide the model with information regarding the current CO_2 level, as well as the longer-term CO_2 level and trend. To maintain simplicity, no additional features are calculated, ensuring a minimal feature set to limit overfitting and enhance model interpretability. Lastly, presence annotations per 10 minutes are retrieved and serve as the target variable during both the training and evaluation processes.

CatBoost is used as classification model, known for its robust performance with minimal parameter tuning [44].

2) TEMPORAL SHIFT FEATURES

To further augment the model's understanding of temporal aspects related to CO_2 levels and trends, the mean and slope of preceding and/or subsequent hours are incorporated as

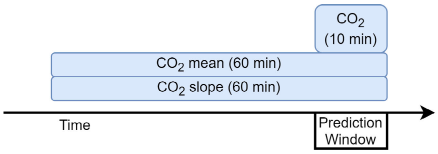

FIGURE 1. **Visualization showing the base features to predict the presence of a 10-minute prediction window. The current CO_2 value is included, and both the CO_2 mean and slope of a 1-hour window are calculated.**

well. We refer to these features as temporal shift features, which can be used to, for example, include data from 2 hours before the prediction window.

When only historical temporal shift features are integrated, the model can perform real-time presence detection, which we define as prospective analysis in this paper. However, if future shift features are included, a delay in inference is required equivalent to the duration of future data used for the prediction. Consequently, we term this methodology as retrospective. Figure 2 visually illustrates this concept. The prediction window signifies the 10-minute interval for which presence is predicted. Historical temporal shift features are incorporated to the left of it, while temporal shift features based on future data after the prediction window are calculated to the right.

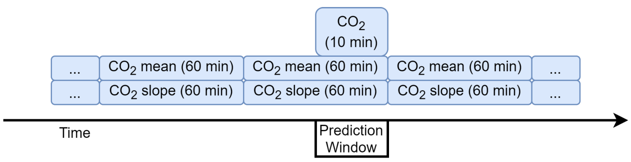

FIGURE 2. **Visual representation showing both prospective (using historical data) and retrospective (using "future" data) temporal shift features. The prediction is delayed in order to calculate and incorporate the retrospective shift features.**

We have chosen to only shift by complete hours, avoiding any overlap between the windows, since more frequent features are not needed to capture the temporal trends effectively. The number of historical and/or future temporal shift features can be seen as a hyperparameter. The decision to use historical and/or future data depends entirely on the use case, considering factors such as the need for real-time prediction and the acceptable level of delay.

3) SLIDING WINDOW NORMALIZATION

To achieve cross-room generalization of the CO_2-based presence detection model, we introduce a different normalization methodology known as sliding window normalization. With this, the common practice of computing normalization parameters once on the training data and reusing them during evaluation or inference is not followed. Instead, we dynamically calculate the normalization parameters, i.e., mean and standard deviation, for each CO_2 measurement based on a sliding window of historical data. This window adjusts with every new sample, as depicted in Figure 3. While sliding window normalization is not a new technique and has been successfully employed in various studies [45], [46], our work is the first to implement it for CO_2-based presence detection.

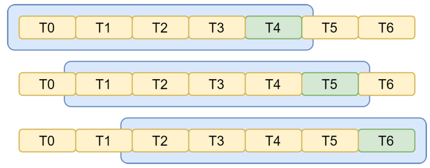

FIGURE 3. **Visualization showing the concept of sliding window normalization. The green box denotes the sample undergoing normalization, which is the window for prediction. Meanwhile, the blue box represents the historical data window employed to compute the normalization parameters. The normalization window slides in tandem with the sample being predicted.**

This approach offers several benefits for presence detection. By basing normalization on a window of historical CO_2 data, it automatically adapts to different environments, user behavior changes, seasonal differences and sensor drift. In our implementation, the window size is set to 30 days, striking a balance between adaptability (short enough) and stability (long enough).

Moreover, our method ensures that the normalization parameters are based on the room where the model is applied on. This helps to reduce some differences between rooms by aligning the characteristics of the data which should improve generalization. While this effect can be partially achieved by transforming based on the normalization parameters of the target room, it lacks the temporal self-adaptation and is impractical for real-world deployment, requiring the capturing of a historical dataset for parameter calculation.

In contrast, our approach is easily deployed. After installing the CO_2 sensor in a room, the normalization gradually adapts to the current environment over a few weeks, enhancing model performance. Although there is an initial cold-start period, this is resolved after a few weeks of operation.

B. OCCUPANCY PROFILING

Our cross-room presence detection methodology can be practically applied in real-world scenarios, without the need for retraining the model for each new room. One valuable application of this model is occupancy profiling, where the focus is beyond real-time presence detection by analyzing average room usage patterns over an extended period. This information is helpful for configuring building control systems, such as heating and cooling, in a smart, automated, and self-adapting manner.

For this occupancy profiling, the retrospective model is leveraged since real-time prediction is not required and a higher performance is achievable. To generate the occupancy profile, the presence probability is calculated at each 10-minute interval by grouping corresponding intervals from multiple days and averaging the presence, as illustrated in Figure 4. For instance, all presence at 8:10 AM is averaged over the past months to acquire the presence probability of that time window. The occupancy profile is obtained by doing this for every other interval as well, and is then resampled to a more practical 30-minute window by calculating the mean.

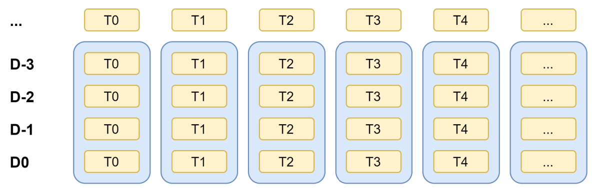

FIGURE 4. Visualization demonstrating the computation process for occupancy profiling. The yellow boxes depict distinct 10-minute intervals throughout a day, while the blue boxes indicate the window used to group and average occupancy, determining the presence probability for each interval. In this example, the window size ranges from the current day *D0* to three days prior *D-3*. This base grouping method can be customized to distinguish between workdays, weekends, and other variations, as detailed in Section V-C1 and V-C2.

The window size can be freely chosen, allowing to balance between stability and adaptability. A larger window size yields a more stable occupancy profile, while a smaller window size increases adaptability. In this paper, a window size of four months (a quarter) is selected, providing a stable profile which yet responds fairly quickly to changes in occupancy patterns.

The grouping methodology is also adjustable based on the specific needs. At a minimum, grouping is done per 10-minute interval, as shown in Figure 4, but additional factors like day of the week or working/weekend days can be considered. For instance, grouping based on working and weekend days provides insight into average occupancy during a working day, without distinguishing between the different weekdays. For more fine-grained insights, grouping can be performed separately for each weekday, allowing investigation into the individual occupancy profiles of weekdays. Moreover, the grouping can extend across multiple rooms and floors to create a more general occupancy profile for an entire floor or building.

To effectively use these occupancy profiles for configuring schedules in building control systems, several approaches can be considered. A simple yet effective method is to apply a threshold, such as 10%, on the occupancy probabilities to determine appropriate activation and deactivation times of building systems. Alternatively, more sophisticated strategies can be employed, such as occupancy-driven Model Predictive Control (MPC), which dynamically adjusts control schedules based on occupancy data to better align with actual needs [2].

V. RESULTS

A. SINGLE-ROOM PRESENCE DETECTION

This section investigates the outcomes of the proposed model and techniques employed for single-room presence detection. The features are calculated in accordance with the methodology detailed in Section IV-A1, and subsequently used to train and assess a CatBoost model, employing data from the same room. Furthermore, the study explores the impact of incorporating temporal shift features.

For each room, the first two-thirds (2/3) of the data are allocated for training, while the last one-third (1/3) is reserved for evaluation. Note that the data is chronologically split rather than randomly sampled, aiming to reduce bias of the evaluation process. During this phase, sliding window normalization is not introduced yet. However, the features are z-normalized, with normalization parameters determined on the training set solely and subsequently applied to both the training and evaluation datasets. Every CatBoost model is trained for 100 iterations since higher numbers do not result in better learning performance, as shown in Figure 11.

This approach is adopted to assess the performance of our CO_2-based presence detection methodology in a single-room configuration. It is important to note that this analysis does not offer insights into the cross-room generalization capabilities of the model, which is later covered in Section V-B.

1) BASELINE

Table 2 highlights the single-room performance of some key configurations. The baseline results only use three features: the current 10-minute CO_2 level, and the 1-hour CO_2 mean and slope.

Among the baseline results, office L1 and home 1 show the highest performance, achieving respective evaluation balanced accuracy (BA) scores of 67.6% and 65.3%. Office L2, which is closely related to office L1, lags behind with a BA score of 59.6%. The one-person office, office S3, demonstrates the lowest performance with a BA score of 52.3%, indicating that the model failed to learn a relevant correlation to detect presence based on CO_2 data.

The low performance of office S3 can be attributed to the frequent opening of windows, allowing fresh air into the room. This disturbs the anticipated rise in CO_2 levels that typically occurs in a closed environment. Given that the presence detection model relies on such a rise in CO_2, the data becomes more challenging for the model to interpret and discern presence accurately.

2) TEMPORAL SHIFT FEATURES

To assess the impact of historical temporal shift features on model performance, we progressively introduce an increasing number of shift features, as depicted in Figure 5. Overall, an increase in BA score is noticeable when more temporal shift features are provided to the model, showing the benefits of the additional temporal information. However, there is no optimal quantity of shift features across all rooms. Offices L1 and L2 show improved performance with the inclusion of more temporal information, whereas office S3 and Home 1 reach a plateau early when shift features are incorporated. Consequently, we empirically evaluated that including up to 8 hours of temporal shift features is a good trade-off between performance enhancement and feature vector length.

When allowing for prediction delays, the integration of future shift features yields additional improvement in

TABLE 2. Training and validation balanced accuracy (BA) scores across all rooms employing different feature configurations. The baseline model solely comprises the current 10-minute CO_2, along with the 1-hour CO_2 mean and slope. The subsequent results incorporate temporal shift features (TSF) in prospective, retrospective, and feature-reduced configurations.

	Office L1		Office L2		Office S3		Home 1	
Configuration	Train BA	Val BA	Train BA	Val BA	Train BA	Val BA	Train BA	Val BA
Baseline, no TSF	79.3%	67.6%	82.0%	59.6%	57.3%	52.3%	83.9%	65.3%
TSF up to 8 hours ago	91.1%	77.4%	91.5%	**69.9%**	72.4%	56.3%	90.0%	66.1%
TSF up to 8 hours ago+later	97.7%	79.0%	99.0%	69.4%	92.7%	60.0%	96.8%	69.4%
TSF 1/2/4/8 hours ago+later	94.6%	**80.7%**	95.9%	69.5%	78.8%	**60.6%**	93.2%	**70.4%**

FIGURE 5. Validation BA scores across all rooms, highlighting the influence of incorporating an increasing number of historical temporal shift features. At $x = 0$, no historical shift features are incorporated, aligning with the baseline model outlined in Table 2. As x values increase, historical temporal shift features are cumulatively included, extending up to 12 hours ago.

FIGURE 6. Validation BA scores across all rooms, showcasing the impact of increasing the number of future temporal shift features. At $x = 0$, no future shift features are incorporated, aligning with the model containing historical shift features up to 8 hours ago. As x values increase, future temporal shift features are cumulatively integrated, extending up to 12 hours later.

performance, as illustrated in Figure 6. The BA scores for all rooms increase when future temporal information is incorporated. The most substantial improvement is observed with the inclusion of the first hour which is due to the delayed response time of a CO_2 signal. To elaborate, while a person's presence in a room is immediate, CO_2 levels increase gradually over time which is why the first future hour is so valuable. Similar to the earlier results, no single configuration is optimal across all rooms. Consequently, we opt for the 8-hour variant once again, as it consistently yields a notable improvement across all rooms on average.

3) FEATURE REDUCTION

The previous section introduced a set of temporal shift features aimed at enhancing the model's performance. While this objective is achieved, the many new features also increase the potential for overfitting. Consequently, several configurations with a reduced set of features are assessed, as outlined in Table 3.

Analysis of the BA scores reveals that a substantial number of temporal shift features can be omitted without compromising performance compared to the initial feature set. This improvement in performance can be attributed to the reduced overfitting, which is shown by the smaller gap between the training and validation BA scores.

Limiting the usage of temporal shift features to 1, 2, 4 and 8 hours ago and later strikes the best balance between feature reduction and overall performance, motivating the use of this configuration throughout the remainder of this study. Moreover, by using this reduced set of features, both data preprocessing and model training times are minimized, with the entire process now completing in approximately 3.3 seconds on an Intel Core i7-8650 CPU.

Table 2 presents the BA scores for the final feature selection, alongside the most noteworthy prospective and retrospective configuration, as well as the baseline. The results demonstrate that incorporating past and future temporal shift features yield notable improvements, ranging from 4.1% to 11.4% over the baseline. Moreover, refining the feature set to include only the most important feature shifts result in an additional improvement of up to 1.7%.

B. CROSS-ROOM PRESENCE DETECTION

In the previous section, we demonstrated the detection capability of presence through CO_2 sensors. However, training a model for each individual room is highly impractical for an

TABLE 3. Training and validation balanced accuracy (BA) scores across all rooms with a reduced set of temporal shift features. The first row represents the complete set of shift features up to 8 hours ago and later, as specified in Section V-A2. Subsequent rows eliminate combinations of shift features.

	Office L1		Office L2		Office S3		Home 1	
Configuration	Train BA	Val BA	Train BA	Val BA	Train BA	Val BA	Train BA	Val BA
TSF up to 8 hours ago+later	97.7%	79.0%	99.0%	69.4%	92.7%	60.0%	96.8%	69.4%
TSF 1/2/3/4/5/7/8 hours ago+later	97.8%	80.2%	98.8%	**69.8%**	91.0%	**61.5%**	96.8%	69.7%
TSF 1/2/4/5/7/8 hours ago+later	97.6%	78.9%	98.9%	69.0%	90.1%	60.8%	96.8%	69.0%
TSF 1/2/4/7/8 hours ago+later	97.1%	80.6%	98.8%	69.5%	89.7%	61.0%	96.9%	69.5%
TSF 1/2/4/8 hours ago+later	94.6%	**80.7%**	95.9%	69.5%	78.8%	60.6%	93.2%	**70.4%**
TSF 1/4/8 hours ago+later	94.6%	79.4%	95.6%	67.6%	76.8%	59.8%	93.1%	70.1%

entire building, as this would require the collection of labeled data for every room.

To address this constraint, it is essential to develop a model capable of cross-room generalization. Such a model should effectively predict presence in rooms beyond the training set without requiring specific labeled data for each room.

To investigate the cross-room generalization of the CO_2-based presence detection methodology, we evaluate a model trained on one room across other rooms in the dataset. The following sections present the results of this cross-room evaluation, using two distinct data normalization approaches.

1) TRADITIONAL NORMALIZATION APPROACH

A first and straightforward strategy involves directly assessing the single-room model's performance on an unseen room. In this approach, the normalization parameters from the training data are applied to the data of the unseen room.

The results obtained from this traditional method, as presented in Table 4, are in line with our expectations. In the majority of cases, the model demonstrates poor performance when applied on unseen rooms, with most of the BA scores around or below 60%. Such scores reflect the model's proficiency in identifying absence instances while facing challenges in accurately predicting presence.

2) SLIDING WINDOW NORMALIZATION APPROACH

To reduce variations in characteristics, sensor operation, and user behavior across different rooms, sliding window normalization is introduced. Table 4 displays the cross-room results achieved by applying this sliding window normalization approach to various rooms.

Overall, a substantial improvement in BA scores is evident compared to the traditional normalization discussed in the previous section. The best result with the traditional normalization approach is a BA score of 81.4% when assessing the model from office L1 on office L2. In contrast, the sliding window approach here yields a higher performance, achieving a BA score of 84.6%. Similarly, evaluating the model of office L2 on office L1 results in a BA score of 83.4%, demonstrating robust generalization between these related offices irrespective of the training room.

Another interesting observation is the models' ability to generalize across different types of rooms. The traditional approach yields a BA score of 65.4% when evaluating the model from office L1 on home 1. However, with

TABLE 4. Comparison of cross-room validation BA scores employing both the traditional and sliding window normalization approaches. With the traditional method, normalization parameters are calculated for the training room and subsequently applied to the validation room. The sliding window normalization approach dynamically adjusts the normalization parameters based on a window of historic data. Each row specifies the training room and presents the validation results for other rooms. Additionally, the performance of the single-room model is provided as an indication of achievable results. However, they should not be directly compared with due to differences in the size of the validation set.

Training on	Validation BA on			
	Office L1	Office L2	Office S3	Home 1
Single-room with traditional normalization approach				
Same room	80.7%	69.5%	60.6%	70.4%
Cross-room with traditional normalization approach				
Office L1	X	81.4%	55.7%	65.4%
Office L2	63.0%	X	51.1%	75.7%
Office S3	62.0%	69.8%	X	70.6%
Home 1	50.0%	53.5%	50.0%	X
Cross-room with sliding window normalization approach				
Office L1	X	84.6%	66.8%	80.6%
Office L2	83.4%	X	64.1%	78.0%
Office S3	60.1%	62.2%	X	65.7%
Home 1	71.4%	70.5%	64.5%	X

the introduction of sliding window normalization, this performance increases to 80.6%. The reverse evaluation, using the model of home 1 on the offices, shows a substantial improvement as well from 50%, 53.5%, and 50% to 71.4%, 70.5%, and 64.5% respectively, with office S3 being the least improved due to CO_2 data influenced by the window.

Remarkably, most of the evaluation performances using sliding window normalization approximate or even exceed the indicative single-room performance, except when office S3 serves as the source room. Although these results are not exactly comparable, this demonstrates the significance of sliding window normalization not only in generalizing between rooms but also in adapting to changing conditions within a room over time.

C. OCCUPANCY PROFILING

Using the occupancy profiling methodology from Section IV-B, the CO_2-based presence detection approach is employed to characterise room usage patterns over a defined period of time. In the following sections, we will explore two example configurations. The first configuration involves grouping data into working and weekend days, while the second configuration creates separate profiles for each day of the week.

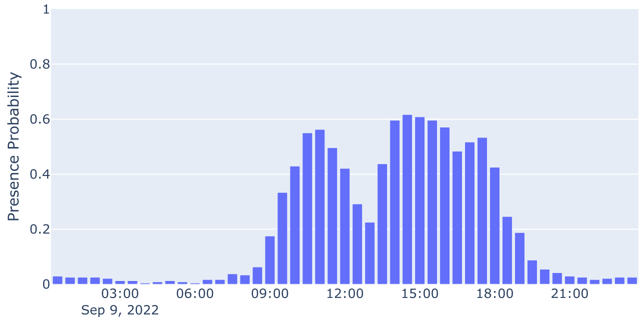

FIGURE 7. **Occupancy profile of the last Friday in the dataset for office L2 when grouping in working and weekend days. This profile is generated based on the predicted presence data using the cross-room model of office L1. The occupancy profile illustrates presence before and after noon, with a noticeable decrease during the lunch break.**

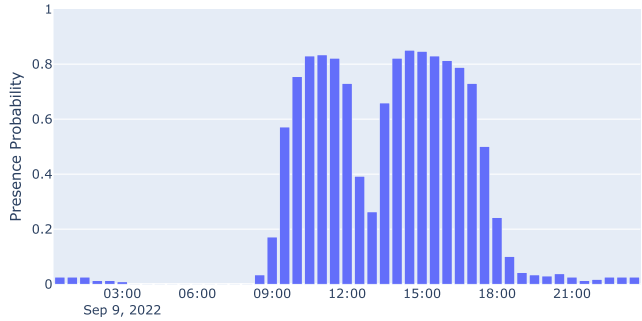

FIGURE 8. **Occupancy profile of the last Friday in the dataset for office L2 when grouping in working and weekend days. This profile is constructed using the ground truth presence data.**

These examples use the retrospective model of office L1 to predict the presence of office L2, representing a realistic usage scenario. The occupancy profiles are computed with a window size of four months. This implies that a quarter of historical presence data is considered when constructing the average room usage profile.

1) GROUPING PER WORKING OR WEEKEND DAY

In this first approach, the data is grouped into workdays (Monday to Friday) and weekend days (Saturday and Sunday). The occupancy profile for a specific day is thus based on the occupancy of all work or weekend days over the preceding four months. This grouping strategy is beneficial when configuring control systems that require different settings for weekdays and weekends. Furthermore, it remains stable across the various days within each category, making it predictable and reliable for occupants.

An example of the occupancy profile obtained through this method is presented in Figure 7. The profile shows the presence probability for the last Friday in the dataset of office L2 using predicted presence data over the past four months. As anticipated, the presence probability starts to gradually increase around 8:30 AM, followed by a decline approaching noon. Subsequently, a prolonged period of presence is visible in the afternoon, suggesting that employees often remain in the office until 6:00 PM or later. This insight is important for configuring a heating system, as a conventional 9-to-5 schedule would leave employees working late in cold conditions.

While Figure 7 shows the occupancy profile based on predicted presence data, reflecting real-world application, Figure 8 displays the profile based on the ground truth presence data of office L2. Although both profiles exhibit similar trends, the predicted occupancy profile generally shows lower presence probabilities, indicating that the model did not predict all instances of presence accurately. Nevertheless, the advantage of occupancy profiling lies in its temporal averaging, ensuring that overall trends remain visible and useful for configuring building control systems.

When comparing the ground truth and predicted occupancy profiles numerically across all data for office L2, the mean absolute error (MAE) of the presence probabilities is 0.043, with a standard deviation of 0.079. While the MAE is low, the relatively large standard deviation suggests inconsistencies in the model's predictions. However, as mentioned earlier, the trends remain visible given the prolonged averaging period.

2) GROUPING PER WEEKDAY

For a more fine-grained approach, the data can also be grouped per weekday, resulting in distinct occupancy profiles for each day of the week (Monday through Sunday). This approach is valuable in scenarios where presence patterns vary significantly across weekdays. For instance, if an office remains unused on certain weekdays due to remote work, this method is capable of identifying days with lower presence probabilities and building control systems can be configured accordingly.

In the case of office L2 in our dataset, while there are no designated remote working days, clear variations exist in the 4-month average occupancy profiles across the different weekdays, as illustrated in Figure 9. Mondays and Tuesdays exhibit the highest presence probabilities, with a noticeable decline towards the end of the workweek, indicating a greater variability and uncertainty in occupancy. However, there are no workdays with near-zero presence probability, showing the need for building control systems to be prepared if employees do arrive. In contrast, the occupancy profile of the weekend shows that the office is then unused.

When employing grouping per weekday, there is a slight increase in both the mean absolute error and the standard deviation, reaching 0.052 and 0.103 respectively. This increase can be attributed to the more fine-grained profiling methodology.

VI. LIMITATIONS OF CO_2-BASED PRESENCE DETECTION

This work demonstrated the feasibility of CO_2-based presence detection in both office and residential settings, with the ability to apply the model to unseen rooms. However, notable performance differences can be seen in the evaluation results. Office L1 shows the highest performance, followed

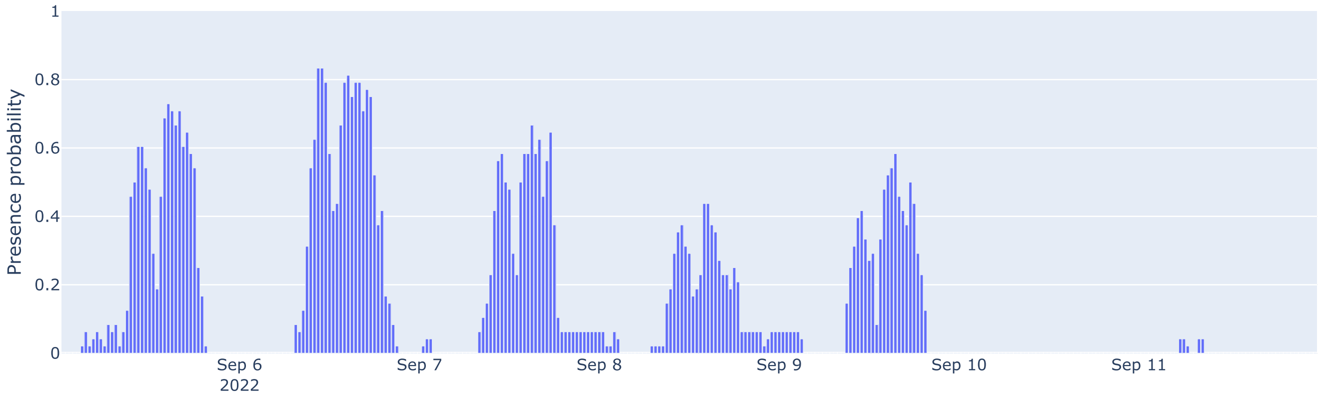

FIGURE 9. Occupancy profile showing last complete week (Monday to Sunday) in the dataset for office L2 when grouping per weekday. This profile is generated based on the predicted presence data using the cross-room model of office L1. Each workday exhibits a distinct occupancy profile, clearly showing variations in presence between the beginning and end of the workweek. The profile of the weekend indicates virtually no presence in the office.

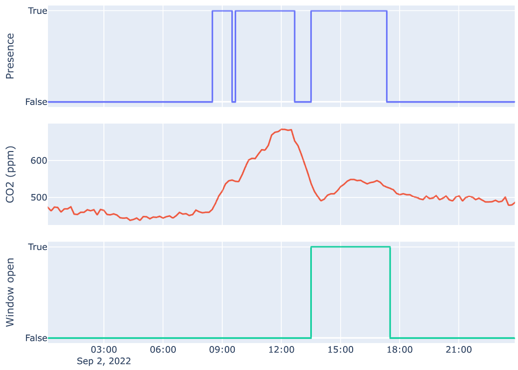

FIGURE 10. Example situation in office L1 demonstrating the influence of an open window on the CO_2 concentration. Notably, the presence in the afternoon with an open window leads to significantly lower CO_2 levels. The status of the window is monitored through a contact sensor.

by office L2 and home 1, while office S3 consistently shows the lowest performance. This poor performance for office S3 is attributed to the disturbed CO_2 data in spaces where windows are frequently open, which significantly influences the expected CO_2 trends within a room. An example of this phenomenon is shown in Figure 10. Before noon, presence in the room with the windows closed results in clearly rising CO_2 levels, which is ideal for presence detection. However, in the afternoon, despite the continued presence, the open windows disturb the natural rise of the CO_2 levels, making presence detection more challenging.

In addition to open windows, ventilation directly influences the CO_2 levels in a room for the same reasons. Moreover, open doors can impact CO_2 levels by allowing it to disperse into adjacent rooms, making it more difficult to detect the rising trend. Similarly, performing CO_2-based presence detection is more challenging in large rooms, where CO_2 can diffuse more extensively, and the presence of a single individual has a minimal effect on the overall CO_2 concentration in the room.

In an ideal case for CO_2-based presence detection, the room would be a sealed environment with no external influences. However, in reality, this is neither feasible nor desirable, as individuals require a constant supply of fresh air for comfort and well-being. The impact of windows and ventilation on CO_2 levels is often not a binary condition. Some degree of influence is acceptable, provided it does not significantly disrupt the CO_2 signal. For instance, if a window is open and the CO_2 level rises gradually, the model will take a longer time to detect presence, but it should eventually succeed. However, when a window is open and the CO_2 level begins to lower, the model's outcome depends on the speed and extent of the CO_2 decrease. If the decline is too rapid and/or the CO_2 level becomes too low, distinguishing between an open window and people leaving the room becomes challenging for the model.

Next to the impact of external factors on the CO_2 level, another drawback is the delayed response time associated with CO_2 sensors. These sensors have a latency in detecting occupancy as they rely on the accumulation of CO_2 in the air, which occurs after individuals have been present for a certain duration. The impact of this response delay is noticeable in Figure 6, where the introduction of temporal shift features of one hour later significantly improved performance.

A. ADDRESSING THE LIMITATIONS

To address the limitations of relying solely on CO_2 data, future work may explore incorporating additional sources of information. For scenarios where instant presence detection is required, the slower response time of CO_2 signals may not be sufficient. In such cases, the presence detection approach could be extended with alternative sensor technologies, such as PIR or radar sensors, which offer real-time detection capabilities. However, this also increases hardware requirements and associated costs.

In instances where a delay in detection is acceptable, or when real-time presence information is not required, the

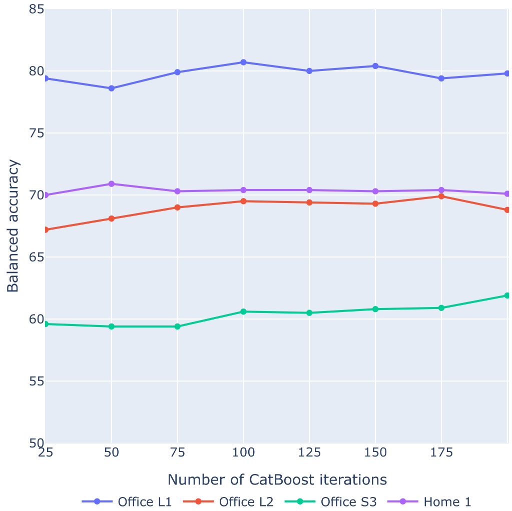

FIGURE 11. Validation BA scores across all rooms using the final single-room model described in Section V-A3. The results demonstrate that increasing the number of CatBoost iterations does not lead to substantial improvements in learning performance. For all other experiments, 100 iterations are used as it provides a good balance between learning capacity and model complexity.

methodology could also be further refined. For example, to account for the impact of open windows on CO_2 levels, window state information can be captured using contact sensors. Such data can be integrated as a feature into the proposed machine learning approach, enabling the model to make more informed decisions. Additionally, window size is another significant factor. This information could be extracted from a building information model (BIM) and incorporated into the detection model by specifying the area (e.g., in square meters) that is open.

To better handle the influence of ventilation on CO_2 dynamics, integrating ventilation system data, such as airflow rates, would be beneficial. Similar to CO_2 data, ventilation data is a time series from which relevant window-based features, such as levels and trends, can be extracted. Additionally, temporal shift features could be employed to capture long-term trends. These features can easily be incorporated into the proposed machine learning framework.

Many modern ventilation systems already measure CO_2 concentrations to regulate ventilation rates. These systems present significant potential for implementing automated presence detection, as they have access to both CO_2 and ventilation rate data. A well-functioning ventilation system can also reduce the need for window ventilation, thereby minimizing its impact on presence detection performance and improving reliability. Furthermore, since such systems are often part of an HVAC solution, the presence information could be shared with, for example, the heating system to optimize timing schedules, aligning with the objectives of the proposed work.

Although these potential improvements offer promising directions for future research, they fall outside the scope of this study. The focus of this work remains on providing a methodology that is broadly applicable, requiring only CO_2 data.

VII. CONCLUSION AND FUTURE WORK

Fixed configuration schemes for building control systems, such as heating, are still a common practice and lead to energy inefficiency and/or suboptimal comfort. This work proposes an approach to automatically configure building control systems based on a room's occupancy profile generated using CO_2-based presence information.

To achieve this objective, the primary focus was on developing a CO_2-based presence detection methodology with good generalization across diverse rooms. We showed that traditional models achieve good predictive capabilities using simple features, namely the current CO_2, 1-hour CO_2 mean and 1-hour CO_2 slope, augmented with temporal shift features to incorporate past and/or future information into the model's prediction. Additionally, the importance of a robust normalization strategy was identified. Employing sliding window normalization significantly improved the model's ability to generalize to unseen rooms and makes practical deployment more convenient, as the model autonomously adapts to room-specific conditions over time.

The generalization of the model depends on the similarity between rooms. For instance, applying the model developed for office L1 to office L2 yields a high BA score of 84.6%. Performance on different room types is slightly lower but remains acceptable, as illustrated by applying the model of office L1 to home 1, resulting in an BA score of 80.6%. However, CO_2-based presence detection has its limitations, such as in cases like office S3, where open windows affect the performance.

The CO_2-based presence predictions were shown to be valuable for the generation of occupancy profiles, indicating the presence probability during the day. Various calculation approaches are available, allowing flexibility in determining whether the resulting profiles should be more coarse-grained or fine-grained. Based on the predicted occupancy profiles for office L2 using the model of office L1, an average error of about 5% with a standard deviation of about 10% is observed. These profiles align closely with the trends of actual room usage, making them a reliable basis for automatically configuring building control systems based on historical presence probability. As the occupancy profiles evolve over time with changing occupant behavior, there is no need for manual reconfiguration of building control systems.

Future research could explore the performance of neural networks operating on raw time series data to validate if these are more effective than the proposed methodology based on feature extraction. Furthermore, to overcome the limitations associated with CO_2 for presence detection, the integration of additional sensors and building-related information can be considered. Moreover, the integration of calendar metadata, such as holidays and vacation periods, can provide valuable support to the occupancy profiling methodology.

We believe that the insights from this study on generalizable CO_2-based presence detection and occupancy profiling will contribute to the shift from fixed timing schemes to more dynamic and smart configuration of building control systems.

ACKNOWLEDGMENT

The authors would like to thank Thibault Blyau and Jeroen Van Der Donckt for their contributions to the research and feedback on the manuscript.

REFERENCES

[1] T. Peffer, M. Pritoni, A. Meier, C. Aragon, and D. Perry, "How people use thermostats in homes: A review," *Building Environ.*, vol. 46, no. 12, pp. 2529–2541, Dec. 2011.

[2] S. Baldi, C. D. Korkas, M. Lv, and E. B. Kosmatopoulos, "Automating occupant-building interaction via smart zoning of thermostatic loads: A switched self-tuning approach," *Appl. Energy*, vol. 231, pp. 1246–1258, Dec. 2018.

[3] I. B. Arief-Ang, M. Hamilton, and F. D. Salim, "A scalable room occupancy prediction with transferable time series decomposition of CO_2 sensor data," *ACM Trans. Sensor Netw.*, vol. 14, nos. 3–4, pp. 1–28, Nov. 2018.

[4] A. N. Sayed, F. Bensaali, Y. Himeur, and M. Houchati, "Edge-based real-time occupancy detection system through a non-intrusive sensing system," *Energies*, vol. 16, no. 5, p. 2388, Mar. 2023.

[5] X. Liang, J. Shim, O. Anderton, and D. Song, "Low-cost data-driven estimation of indoor occupancy based on carbon dioxide (CO_2) concentration: A multi-scenario case study," *J. Building Eng.*, vol. 82, Apr. 2024, Art. no. 108180.

[6] C. Karasoulas, C. Keroglou, E. Katsiri, and G. C. Sirakoulis, "Hidden Markov models for presence detection based on CO_2 fluctuations," *Frontiers Robot. AI*, vol. 10, pp. 1–18, Oct. 2023. [Online]. Available: https://www.frontiersin.org/journals/robotics-and-ai/articles/10.3389/frobt.2023.1280745/full

[7] J. Kim, J. Bang, A. Choi, H. J. Moon, and M. Sung, "Estimation of occupancy using IoT sensors and a carbon dioxide-based machine learning model with ventilation system and differential pressure data," *Sensors*, vol. 23, no. 2, p. 585, Jan. 2023.

[8] Q. Huang, M. Syndicus, J. Frisch, and C. van Treeck, "Spatial features of CO_2 for occupancy detection in a naturally ventilated school building," *Indoor Environ.*, vol. 1, no. 3, Oct. 2024, Art. no. 100018.

[9] J. Dutta and S. Roy, "OccupancySense: Context-based indoor occupancy detection & prediction using CatBoost model," *Appl. Soft Comput.*, vol. 119, Apr. 2022, Art. no. 108536.

[10] P. Kumari, S. R. N. Reddy, and R. Yadav, "Indoor occupancy detection and counting system based on boosting algorithm using different sensor data," *Building Res. Inf.*, vol. 52, nos. 1–2, pp. 87–106, Feb. 2024.

[11] J. Caniou, *Passive Infrared Detection: Theory and Applications*. Cham, Switzerland: Springer, Mar. 2013.

[12] A. Shokrollahi, J. A. Persson, R. Malekian, A. Sarkheyli-Hägele, and F. Karlsson, "Passive infrared sensor-based occupancy monitoring in smart buildings: A review of methodologies and machine learning approaches," *Sensors*, vol. 24, no. 5, p. 1533, Feb. 2024.

[13] T. Teixeira, G. Dublon, and A. Savvides, "A survey of human-sensing: Methods for detecting presence, count, location, track, and identity," *ACM Comput. Surveys*, vol. 5, p. 59, Jan. 2010.

[14] A. A. Trofimova, A. Masciadri, F. Veronese, and F. Salice, "Indoor human detection based on thermal array sensor data and adaptive background estimation," *J. Comput. Commun.*, vol. 5, no. 4, pp. 16–28, 2017.

[15] S. Singh and B. Aksanli, "Non-intrusive presence detection and position tracking for multiple people using low-resolution thermal sensors," *J. Sensor Actuator Netw.*, vol. 8, no. 3, p. 40, Jul. 2019.

[16] D. Lefloch, F. A. Cheikh, J. Y. Hardeberg, P. Gouton, and R. Picot-Clémente, "Real-time people counting system using a single video camera," *Real-Time Image Process.*, vol. 6811, Feb. 2008, Art. no. 681109.

[17] J. Zou, Q. Zhao, W. Yang, and F. Wang, "Occupancy detection in the office by analyzing surveillance videos and its application to building energy conservation," *Energy Buildings*, vol. 152, pp. 385–398, Oct. 2017.

[18] J. W. Choi, D. H. Yim, and S. H. Cho, "People counting based on an IR-UWB radar sensor," *IEEE Sensors J.*, vol. 17, no. 17, pp. 5717–5727, Sep. 2017.

[19] J.-H. Choi, J.-E. Kim, and K.-T. Kim, "Deep learning approach for radar-based people counting," *IEEE Internet Things J.*, vol. 9, no. 10, pp. 7715–7730, May 2022.

[20] Q. Huang, "Occupancy-driven energy-efficient buildings using audio processing with background sound cancellation," *Buildings*, vol. 8, no. 6, p. 78, Jun. 2018.

[21] J. M. Sim, Y. Lee, and O. Kwon, "Acoustic sensor based recognition of human activity in everyday life for smart home services," *Int. J. Distrib. Sensor Netw.*, vol. 11, no. 9, Sep. 2015, Art. no. 679123.

[22] W. Kleiminger, C. Beckel, T. Staake, and S. Santini, "Occupancy detection from electricity consumption data," in *Proc. 5th ACM Workshop Embedded Syst. Energy-Efficient Buildings*, New York, NY, USA, Nov. 2013, pp. 1–8.

[23] A. Akbar, M. Nati, F. Carrez, and K. Moessner, "Contextual occupancy detection for smart office by pattern recognition of electricity consumption data," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Jun. 2015, pp. 561–566.

[24] B. Balaji, J. Xu, A. Nwokafor, R. Gupta, and Y. Agarwal, "Sentinel: Occupancy based HVAC actuation using existing WiFi infrastructure within commercial buildings," in *Proc. 11th ACM Conf. Embedded Networked Sensor Syst.*, New York, NY, USA, Nov. 2013, pp. 1–14.

[25] X. Lu, H. Wen, H. Zou, H. Jiang, L. Xie, and N. Trigoni, "Robust occupancy inference with commodity WiFi," in *Proc. IEEE 12th Int. Conf. Wireless Mobile Comput., Netw. Commun. (WiMob)*, Oct. 2016, pp. 1–8.

[26] S. Depatla, A. Muralidharan, and Y. Mostofi, "Occupancy estimation using only WiFi power measurements," *IEEE J. Sel. Areas Commun.*, vol. 33, no. 7, pp. 1381–1393, Jul. 2015.

[27] Y. Zhang, X. Wang, J. Wen, and X. Zhu, "WiFi-based non-contact human presence detection technology," *Sci. Rep.*, vol. 14, no. 1, p. 3605, Feb. 2024.

[28] L. Rueda, K. Agbossou, A. Cardenas, N. Henao, and S. Kelouwani, "A comprehensive review of approaches to building occupancy detection," *Building Environ.*, vol. 180, Aug. 2020, Art. no. 106966.

[29] D. Trivedi and V. Badarla, "Occupancy detection systems for indoor environments: A survey of approaches and methods," *Indoor Built Environ.*, vol. 29, no. 8, pp. 1053–1069, Oct. 2020.

[30] T. Li, X. Liu, G. Li, X. Wang, J. Ma, C. Xu, and Q. Mao, "A systematic review and comprehensive analysis of building occupancy prediction," *Renew. Sustain. Energy Rev.*, vol. 193, Apr. 2024, Art. no. 114284.

[31] L. Zhao, Y. Li, R. Liang, and P. Wang, "A state of art review on methodologies of occupancy estimating in buildings from 2011 to 2021," *Electronics*, vol. 11, no. 19, p. 3173, Oct. 2022.

[32] A. N. Sayed, Y. Himeur, and F. Bensaali, "Deep and transfer learning for building occupancy detection: A review and comparative analysis," *Eng. Appl. Artif. Intell.*, vol. 115, Oct. 2022, Art. no. 105254.

[33] J. Caballero-Peña, G. Osma-Pinto, J. M. Rey, S. Nagarsheth, N. Henao, and K. Agbossou, "Analysis of the building occupancy estimation and prediction process: A systematic review," *Energy Buildings*, vol. 313, Jun. 2024, Art. no. 114230.

[34] S. Wang, J. Burnett, and H. Chong, "Experimental validation of CO_2-based occupancy detection for demand-controlled ventilation," *Indoor Built Environ.*, vol. 8, no. 6, pp. 377–391, Nov. 1999.

[35] D. Calì, P. Matthes, K. Huchtemann, R. Streblow, and D. Müller, "CO_2 based occupancy detection algorithm: Experimental analysis and validation for office and residential buildings," *Building Environ.*, vol. 86, pp. 39–49, Apr. 2015.

[36] F. Nienaber, M. T. Wesseling, D. Müller, and D. Calì, "Validation and optimization of air quality sensor based occupancy detection algorithms," in *Proc. Roomvent Ventilation*, Jan. 2018, pp. 109–114.

[37] A. Ebadat, G. Bottegal, M. Molinari, D. Varagnolo, B. Wahlberg, H. Hjalmarsson, and K. H. Johansson, "Multi-room occupancy estimation through adaptive gray-box models," in *Proc. 54th IEEE Conf. Decis. Control (CDC)*, Dec. 2015, pp. 3705–3711.

[38] S. Wolf, D. Calì, J. Krogstie, and H. Madsen, "Carbon dioxide-based occupancy estimation using stochastic differential equations," *Appl. Energy*, vol. 236, pp. 32–41, Feb. 2019.

[39] G. Diraco, A. Leone, and P. Siciliano, "People occupancy detection and profiling with 3D depth sensors for building energy management," *Energy Buildings*, vol. 92, pp. 246–266, Apr. 2015.

[40] X. Kang, D. Yan, J. An, Y. Jin, and H. Sun, "Typical weekly occupancy profiles in non-residential buildings based on mobile positioning data," *Energy Buildings*, vol. 250, Nov. 2021, Art. no. 111264.
[41] Z. Yang and B. Becerik-Gerber, "Modeling personalized occupancy profiles for representing long term patterns by using ambient context," *Building Environ.*, vol. 78, pp. 23–35, Aug. 2014.
[42] J. Van Der Donckt, J. Van Der Donckt, M. Rademaker, G. Vandewiele, and S. Van Hoecke, "Do not sleep on linear models: Simple and interpretable techniques outperform deep learning for sleep scoring," *Biomed. Signal Process. Control*, vol. 81, Jul. 2023, Art. no. 104429.
[43] J. Van Der Donckt, J. Van Der Donckt, E. Deprost, and S. Van Hoecke, "Tsflex: Flexible time series processing & feature extraction," *SoftwareX*, vol. 17, Jan. 2022, Art. no. 100971. [Online]. Available: https://github.com/predict-idlab/tsflex
[44] L. Prokhorenkova, G. Gusev, A. Vorobev, A. V. Dorogush, and A. Gulin, "CatBoost: Unbiased boosting with categorical features," in *Proc. 32nd Int. Conf. Neural Inf. Process. Syst.*, vol. 31, Dec. 2018, pp. 6639–6649.
[45] T. Tanaka, I. Nambu, Y. Maruyama, and Y. Wada, "Sliding-window normalization to improve the performance of machine-learning models for real-time motion prediction using electromyography," *Sensors*, vol. 22, no. 13, p. 5005, Jul. 2022.
[46] S. Simtharakao and D. Sutivong, "Exploring normalization techniques in neural networks for Bitcoin candlestick price prediction," in *Proc. Int. Conf. Artif. Intell. Inf. Commun. (ICAIIC)*, Feb. 2023, pp. 483–488.

JELLE VANHAEVERBEKE received the M.Sc. degree in electronics and ICT engineering, in 2019. He is currently pursuing the Ph.D. degree with IDLab, IMEC, Ghent University, under the supervision of Prof. Sofie Van Hoecke and Prof. Steven Verstockt. His current research focuses on the development and application of spatiotemporal event detection solutions across various domains, including industrial processes, building control systems, and sports analytics.

EMIEL DEPROST received the M.Sc. degree in electronics and ICT engineering, in 2019. He is currently pursuing the Ph.D. degree with IDLab-MEDIA. In 2020, he obtained funding from the Research Foundation-Flanders (FWO). His doctoral research focuses on the use of machine learning to improve the measurement and preservation of plant biodiversity in grasslands. He has a high interests in computer vision, time-series data, embedded systems, and loves to make interactive visualizations.

STEVEN VERSTOCKT received the master's degree in informatics from Ghent University, Belgium, in 2003, and the Ph.D. degree in multi-modal video fire analysis from the Multimedia Laboratory (iMinds), Electronics and Information Systems Department, Ghent University. Since 2012, he has been a Postdoctoral Researcher with iMinds, Electronics and Information Systems Department, Ghent University. In September 2015, he started a tenure track professorship in multimedia. His current research focuses on multimodal data processing, video analysis, computer vision, and GEO-ICT. Over the last years, his research group has been getting involved in several sports, media, and heritage-related projects/proposals to share their expertise on spatio-temporal data collection, filtering, classification, enrichment, mapping, and visualization.

SOFIE VAN HOECKE received the degree from the Engineering Department, Ghent University, in 2003, and the Ph.D. degree in computer science engineering from the Department of Information Technology, Ghent University, with a focus on efficient service management in healthcare. After being a Postdoctoral Research Engineer with the Department of Information Technology, she started as a Lecturer in ICT and an ICT Research Coordinator with the University College West Flanders. Currently, she is an Associate Professor with IDLab-Data Science Laboratory, Ghent University. Her specialties are: multi-sensor and service-oriented architectures, novel services, condition monitoring, emotion recognition, machine learning, semantic dashboards, and the fusion of machine learning and semantic technologies, applied in both predictive maintenance and predictive healthcare.