

Room Occupancy Prediction: Exploring the Power of Machine Learning and Temporal Insights

Siqi Mao^{1,#}, Yaping Yuan^{2,#}, Yinpu Li^{1,*,#}, Ziren Wang³, Yuanxin Yao¹, Yixin Kang¹

¹Department of Statistics, Florida State University, Tallahassee, USA

²Department of Mathematics and Statistics, University of Massachusetts Amherst, Amherst, USA

³Department of Statistics, Rice University, Houston, USA

*Corresponding author: yinpulee@gmail.com

Received December 12, 2023; Revised January 15, 2024; Accepted January 22, 2024

These authors contributed equally to this work

Abstract Energy conservation in buildings is a paramount concern to combat greenhouse gas emissions and combat climate change. The efficient management of room occupancy, involving actions like lighting control and climate adjustment, is a pivotal strategy to curtail energy consumption. In contexts where surveillance technology isn't viable, non-intrusive sensors are employed to estimate room occupancy. In this study, we present a predictive framework for room occupancy that leverages a diverse set of machine learning models, with Random Forest consistently achieving the highest predictive accuracy. Notably, this dataset encompasses both temporal and spatial dimensions, revealing a wealth of information. Intriguingly, our framework demonstrates robust performance even in the absence of explicit temporal modeling. These findings underscore the remarkable predictive power of traditional machine learning models. The success can be attributed to the presence of feature redundancy, the simplicity of linear spatial and temporal patterns, and the advantages of high-frequency data sampling. While these results are compelling, it's essential to remain open to the possibility that explicitly modeling the temporal dimension could unlock deeper insights or further enhance predictive capabilities in specific scenarios. In summary, our research not only validates the effectiveness of our prediction framework for continuous and classification tasks but also underscores the potential for improvements through the inclusion of temporal aspects. The study highlights the promise of machine learning in shaping energy-efficient practices and room occupancy management.

Keywords: room occupancy prediction, classification, support vector machine, random forest, XGBoost, Internet-of-Things (IoT)

Cite This Article: Siqi Mao, Yaping Yuan, Yinpu Li, Ziren Wang, Yuanxin Yao, and Yixin Kang, "Room Occupancy Prediction: Exploring the Power of Machine Learning and Temporal Insights." American Journal of Applied Mathematics and Statistics, vol. 12, no. 1 (2024): 1-9. doi: 10.12691/ajams-12-1-1.

1. Introduction

Buildings account for as much as 40% of the world's total energy consumption [1] and contribute to 30% of greenhouse gas emissions [2]. Therefore, diminishing the energy consumption within the construction sector will significantly contribute to addressing global energy usage and environmental carbon emissions concerns [3]. In the pursuit of creating energy-efficient and comfortable indoor environments, the management of Heating, Ventilation, and Air Conditioning (HVAC) systems has emerged as a pivotal area of research and application. As the world faces increasingly complex challenges related to climate change, resource conservation, and sustainability, the optimization of HVAC systems has become an imperative task for both researchers and practitioners. One of the fundamental aspects of HVAC system optimization is the intelligent prediction and adaptation of room occupancy patterns. Occupancy in a room or building

directly influences the HVAC system's operation. Efficiently managing temperature, ventilation, and airflow in spaces with varying levels of occupancy can lead to substantial energy savings, reduce environmental impact, and enhance the overall quality of indoor environments.

The advancement in Internet-of-Things (IoT) and communication technologies has played a pivotal role in facilitating the integration of various wireless technologies into occupancy sensing. These wireless technologies include Radio Frequency Identification (RFID) [4], Wi-Fi [5], and Bluetooth Low Energy (BLE) [6], all with the aim of enhancing the range, speed, precision, level of detail, and energy efficiency of existing sensing methods. Nevertheless, despite their merits, terminal-based approaches encounter constraints due to the necessity of deploying specialized sensors and implementing third-party software components, such as BLE beacons, Wi-Fi access points, and mobile applications. These requirements lead to increased implementation costs, and disruption of occupants' daily routines, and raise concerns regarding privacy.

Furthermore, there are other methods rely on passive

sensing technologies, which encompass CO₂ sensors [7], Passive Infrared (PIR) sensors [8], ultrasonic detection sensors [9], sound detection sensors [10], camera systems [11], and intelligent power meters [12]. These non-terminal strategies indirectly accumulate occupancy data for specific zones within a building where these sensors are strategically positioned.

Machine learning techniques have been widely used in numerous sectors recently and achieved gratifying results [13,14,15,16,17]. In addition to the progress made in sensing technologies, the recent utilization of machine learning for occupancy detection has yielded significant enhancements in accuracy compared to conventional methods. These enhancements can be credited to machine learning algorithms' capacity to analyze extensive datasets gathered from the building and discern patterns and relationships with the building's occupancy. Kadouce et al. [18] employed a support vector machine (SVM) classifier to deduce whether occupants were present or absent. They accomplished this by training the SVM with sensory datasets derived from a variety of sources, including motion sensors, pressure detectors, lighting controls, door and switch contactors, as well as flow meters. Sangogboye et al. [19] focused on modeling occupant presence in two commercial buildings using machine learning techniques applied to motion data. Their findings demonstrated that SVM exhibited robust performance in accordance with their results. Razavi et al. [20], conducted an analysis of electricity consumption data from over 5000 residential homes. They assessed the performance of a diverse range of machine learning models, including Support Vector Machines (SVM), K-Nearest Neighbors (KNN), Random Forest (RF), Gradient Boosting (GB), and neural networks. Their objective was to predict both current and future occupancy status. In a separate investigation, Park et al. [21] introduced an occupancy detection model based on Long Short-Term Memory (LSTM) architecture. They utilized energy consumption data obtained from smart plugs to discern the presence of occupants in residential homes.

Utilizing supervised learning techniques such as Random Forest, Decision Trees, and Bagging, Koklu et al. [22] conducted occupancy detection and assessed the models' performance through classification accuracy metrics. Assessing the accuracy of classification involves the computation of true positives (correctly recognized examples belonging to the class), true negatives (correctly recognized examples not part of the class), false positives (incorrectly assigned examples to the class), and false negatives (examples not recognized as class examples).

The central objective of this investigation is to demonstrate the robust applicability of the comprehensive framework introduced in [23]. This framework aims to enlighten individuals with a strong background in data science, who may have limited exposure to the field of sensor research, and aid in their adaptation for practical use across various research domains. We have employed this prediction framework to address the challenge of room occupancy prediction. While we have updated the framework by incorporating various classification methods, we have adhered to the established methodology. Our current study provides empirical evidence of the framework's high performance in predictive tasks. Additionally, our research has unveiled a counter-intuitive

revelation: machine learning models do not explicitly consider a temporal structure if the temporal structure of the target variable has been fully reflected in responses variable. This unexpected superiority challenges conventional assumptions and underscores the potential practical utility of machine learning models in the context of room occupancy prediction.

This research generalized a prediction framework which is proposed in [23]. An empirical data analysis is performed to test the framework by utilizing the dataset referenced in [24]. The dataset comprises more than 10,000 data points, each associated with 16 distinct features, representing the readings of specific sensors. The experiment incorporated five distinct sensor types: temperature, illumination, sound, CO₂, and passive infrared (PIR). The primary objective of this investigation is to accurately estimate the occupancy levels within a laboratory setting.

The remaining sections of the paper are organized as follows: In Section 2, we provide a comprehensive exploration of the acquired data and the prediction methodology. This section offers an extensive elucidation of the various models under the prediction framework in our study, including Multinomial Logistic Regression, Linear Discriminant Analysis, Multi-class Support Vector Machine (MSVM), Random Forest (RF), XGBoost, LightGBM, and a Multi-layer Perceptron classifier. It also encompasses a detailed description of the evaluation metrics we employed and presents a thorough overview of the entire implementation process. Furthermore, Section 3 includes a comprehensive performance comparison across different models, followed by an in-depth analysis. Additionally, we incorporate a SHAP (SHapley Additive exPlanations) analysis within Section 4. Section 5 explores why time series models may not be necessary for this specific dataset from two perspectives. The first perspective addresses the redundancy of time-dependent information in the target variable, it already presents in responses variables of the dataset, while the second considers the sampling frequency.

Lastly, in Section 6, we summarize the key conclusions derived from this study and outline potential directions for future research.

2. Data and Methodologies

2.1. Data

The data are collected based on the experiment from the research [24]. The purpose of the experiment was to estimate occupancy by the signal information from non-intrusive sensors. Seven sensor nodes and one edge node were arranged in the setup, and the sensor nodes used wireless transceivers to send data to the edge every thirty seconds. In this experiment, five distinct kinds of non-intrusive sensors were used: digital passive infrared (PIR), light, sound, temperature, and CO₂. Calibration of the PIR, sound, and CO₂ sensors required considerable labor. Before the CO₂ sensor was used for the first time, zero-point calibration was done by hand. This involved keeping the sensor in a clean environment for more than 20 minutes and then lowering the calibration pin (HD pin) for

more than 7 seconds. In essence, the sound sensor is a microphone with an analog amplifier of variable gain attached. Two trimpots are available for the PIR sensor: one for adjusting sensitivity and the other for adjusting how long the output remains high after motion detection. These two were both set to their maximum settings. The temperature, light, and sound sensors were located in sensor nodes S1–S4, a CO₂ sensor was located in sensor node S5, and one PIR sensor each was placed on the ceiling edges of S6 and S7 at an angle that optimized the sensor's field of view for motion detection. Over the course of four carefully planned days, the number of occupants in each room varied from zero to three. The manual recording of the room's occupancy count served as the basis for accuracy. The data description is provided in Table 1.

Table 1. Data Description

Variable Name	Type	Description	Units
Date	Date		YYYY/MM/DD
Time	Time		HH:MM:SS
S1_Temp	Continuous		C
S2_Temp	Continuous		C
S3_Temp	Continuous		C
S4_Temp	Continuous		C
S1_Light	Integer		Lux
S2_Light	Integer		Lux
S3_Light	Integer		Lux
S4_Light	Integer		Lux
S1_Sound	Continuous	amplifier output read by ADC	Volts
S2_Sound	Continuous	amplifier output read by ADC	Volts
S3_Sound	Continuous	amplifier output read by ADC	Volts
S4_Sound	Continuous	amplifier output read by ADC	Volts
S5_CO2	Integer		PPM
S5_CO2_Slope	Continuous	Slope of CO ₂ values	
S6_PIR	Binary	Binary value conveying motion detection	
S7_PIR	Integer	Binary value conveying motion detection	
Room_Occupancy_Count	Integer	Ground Truth	

Source: The table is summarized from the data source website <https://archive.ics.uci.edu/dataset/864/room+occupancy+estimation>.

2.2. Exploratory Data Analysis

There are 10129 entries in total and all features are numerical variables. In this study, 70% is selected as the training set. The summary statistics of features in the training set are displayed in Table 2. The features "Date" and "Time" are excluded all the classification models. The rationales and discussions are provided in Section 5. The values of the target variable "Room_Occupancy_Count"

are integers and ranges from 0 to 3.

Table 2. Summary Statistics of Training data. Count = 8103

Statistics	Mean	STD	Min	Q1	Median	Q3	Max
S1_Temp	25.45	0.35	24.94	25.19	25.38	25.63	26.38
S2_Temp	25.55	0.59	24.75	25.19	25.38	25.63	29.00
S3_Temp	25.05	0.43	24.44	24.69	24.94	25.38	26.19
S4_Temp	25.75	0.36	24.94	25.44	25.75	26.00	26.50
S1_Light	24.92	50.52	0.00	0.00	0.00	11.00	165.00
S2_Light	25.39	66.47	0.00	0.00	0.00	13.00	258.00
S3_Light	33.97	58.00	0.00	0.00	0.00	49.00	280.00
S4_Light	13.09	19.43	0.00	0.00	0.00	21.00	74.00
S1_Sound	0.17	0.32	0.06	0.07	0.08	0.08	3.88
S2_Sound	0.12	0.26	0.04	0.05	0.05	0.06	3.44
S3_Sound	0.16	0.41	0.04	0.06	0.06	0.07	3.67
S4_Sound	0.10	0.12	0.05	0.06	0.08	0.10	3.40
S5_CO2	461.0	201.1	345.0	355.0	360.00	465.0	1270.0
S5_CO2_Slope	-0.02	1.17	-6.30	-0.05	0.00	0.00	8.98
S6_PIR	0.09	0.28	0.00	0.00	0.00	0.00	1.00
S7_PIR	0.08	0.27	0.00	0.00	0.00	0.00	1.00

2.3. Methodologies

The initial introduction of the prediction framework can be traced back to the study presented in [23], where it was primarily employed for addressing a water quality prediction challenge, constituting a regression problem. In this paper, we adapt and extend the prediction framework to encompass classification problems, and its efficacy is evaluated using room occupancy prediction data. The classification models which are used in this paper is illustrated as follows.

Benchmark. The benchmark model provides a basic performance threshold to compare against more complex models. It uniformly predicts the most frequent class in the training data for all samples. While simple, it establishes a minimum accuracy level to improve upon.

Multinomial Logistic Regression. Multinomial logistic regression extends the logistic regression model to multi-class classification tasks. It estimates the probability of each class using a softmax activation function. The cross-entropy loss is optimized to find the regression coefficients. Regularization helps prevent overfitting.

Linear Discriminant Analysis (LDA). LDA finds a linear combination of features that maximizes the separation between classes. It assumes normal distributions with equal covariance for each class. The within-class and between-class scatter matrices are computed. Eigendecomposition provides projection directions to transform the data before classification.

Multi-layer Perceptron (MLP) Classifier. The MLP classifier uses a feedforward artificial neural network model. Multiple layers of nodes with nonlinear activation functions allow modeling complex relationships. The network is trained via backpropagation to optimize a cross-entropy loss function. Regularization methods and dropout prevent overfitting.

Multiclass Support Vector Machine (MSVM). MSVM extends SVM for multiple classes using the one-vs-one or one-vs-all approaches. It maximizes the margin between classes by mapping inputs to a high-dimensional feature

space with a kernel function. Optimization finds the maximum margin hyperplanes between classes.

Random Forest (RF). RF trains an ensemble of decision trees on random subsets of features and data. Each tree makes an independent prediction and the results are aggregated through voting or averaging to make the overall prediction. Randomness reduces overfitting and improves generalization.

LightGBM. LightGBM employs gradient boosting to train an ensemble of decision trees. Unlike traditional gradient boosting frameworks, trees are grown leaf-wise rather than level-wise to minimize loss and enhance accuracy. Through sampling, less informative instances are filtered out to mitigate noise. Additionally, multiple leaves can share the same child node to bolster generalization. Regularization is applied to curb overfitting, while parallel and GPU-based learning mechanisms significantly augment efficiency.

XGBoost. XGBoost stands for eXtreme Gradient Boosting, and as the name suggests, it implements gradient boosted decision trees with a focus on speed and performance. A notable feature is the regularization of the objective function to simplify the model complexity. Employing quantile sketching, it approximates the tree learning process, thus accelerating training. The framework is designed for parallel and distributed computing, leveraging multiple CPUs and machines. Out-of-core computing is another feature, wherein disk space is utilized to enhance performance further. The algorithm is finely tuned for sparse data handling and allows extensive customization through hyperparameter tuning.

2.4. Evaluation Metrics

Since we approach this task as a classification problem, we employ Balanced Accuracy, F-1 Score, and Area Under ROC Curve (AUC) as our evaluation metrics. Given that it involves multiple classes, we additionally incorporate the sample size of each class as weights to consolidate the measurements into a single metric:

$$\text{Balanced Accuracy} = \frac{\text{Specificity} + \text{Recall}}{2},$$

$$\text{F1-Score} = \frac{2 * \text{Precision} * \text{Recall}}{\text{Precision} + \text{Recall}},$$

$$\text{AUC} = \frac{\sum_{i \in \{C=1\}} \text{rank}_i - \frac{M(M+1)}{2}}{M \times N},$$

$$\text{Weighted Metric} = \frac{\sum_{j=1}^C w_j \text{Metric}_j}{\sum_{j=1}^C w_j},$$

where M is positive sample size and N is negative sample size, and set C represents different classes. In the context of all three metrics, our goal is to maximize them as much as possible. Balanced Accuracy takes into account imbalanced datasets, ensuring an accurate representation of model performance across different classes. F-1 Score provides a balance between precision and recall, making it particularly useful when dealing with situations where false positives and false negatives carry

different costs. AUC, on the other hand, quantifies the model's ability to distinguish between positive and negative instances across various thresholds. These metrics collectively offer a comprehensive perspective on the effectiveness of classification models.

2.5. Implementation

This section outlines the proposed prediction methodology, detailing its development through hyperparameter optimization, prediction generation, and assessment using specific scoring metrics. The implementation process involves the following steps:

1. Selection of Scoring Function: To fine-tune and assess hyperparameters, we employ the Area Under the Receiver Operating Characteristic Curve (AUC-ROC) for all candidate techniques.
2. Defining Hyperparameter Space: Each method's hyperparameter space is defined by specifying the relevant parameters.
3. Cross-validation: We utilize cross-validation to perform hyperparameter optimization, aiming to identify the optimal hyperparameters for each method.
4. Performance Evaluation: The evaluation metrics are computed on the test dataset for the best-tuned models within each method.
5. Results Storage: The outcomes, including the performance of each approach and the significance of its features, are saved for facilitating model comparison.

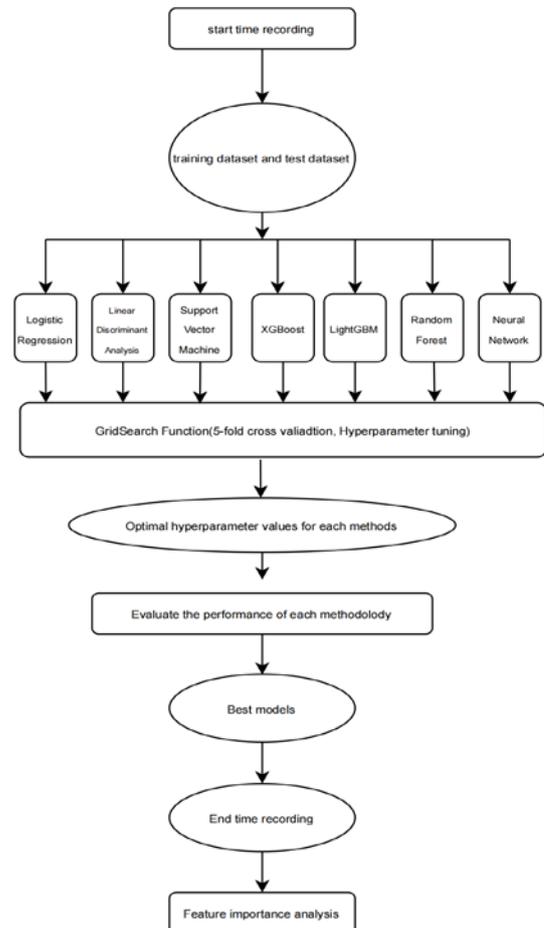


Figure 1. Flowchart of The Proposed Room Occupancy Prediction Framework

3. Results

The prediction results are summarized in this section. The 5-fold cross validation is employed in tuning the hyper-parameters. Due to the sensitivity of machine learning algorithms to the scale of the features, standardization is applied to all the features.

The model performance summary is displayed in Table 3. The model is sorted by the Weighted AUC score. Based on the performance of Benchmark model, we can see the necessity of employing advanced models. Furthermore, the random forest model achieves the best performance under the selected metrics.

The hyper-parameters tuning time is presented in Table 4. The number of fits in each model’s hyper-parameter tuning is indicated by the term “Total Fits” in the tables. The term “Tuning Time” refers to the duration of each model’s whole hyper-parameter tuning operation. “Fitting Time (Best Model)” refers to the amount of time spent training each model’s best hyper-parameters. The term “Average Tuning” refers to the typical tuning time for hyper-parameters. From the the tuning time table, the tree methods have satisfying results and efficiency.

Table 3. Room Occupancy Prediction Results (Numerical Features: Standardized)

Models	Weighted F_1 -Score	Weighted AUC	Balanced Accuracy
Benchmark	0.70988	N/A	0.25000
Logistic Regression	0.98177	0.99654	0.94155
LDA	0.98159	0.99714	0.94216
MSVM	0.99214	0.99951	0.97635
MLP	0.99704	0.99988	0.99196
LightGBM	0.99703	0.99994	0.99007
XGBoost	0.99703	0.99994	0.98981
RF	0.99852	0.99996	0.99490

Table 4. Running Time Summary in Seconds (Numerical Features: Standardized)

Models	Total Fits	Tuning Time	Average Tuning	Fitting Time (Best Model)
Logistic Regression	15	2.04	0.10	0.14
LDA	15	0.15	0.03	0.01
MSVM	20	3.27	1.36	0.16
MLP	40	4007.15	97.81	100.18
LightGBM	2000	295.56	0.23	0.15
XGBoost	5760	3199.18	4.52	0.56
RF	1200	120.17	0.53	0.10

4. SHapley Additive exPlanations (SHAP)

For single predicted point, The localized contributions to that specific forecast are quantified by the SHAP values [25]. We provide the definition for the SHAP value:

$$\begin{aligned} \Phi_i(F, x) &= \sum_{Z \subseteq X'} \left[\binom{M}{|Z|} \right]^{-1} [F(Z) - F(Z \setminus x_i)] \\ &= \sum_{Z \subseteq X'} \frac{(|Z|-1)!(M-|Z|)!}{M!} [F(Z) - F(Z \setminus x_i)]. \end{aligned}$$

In this equation, $\Phi_i(F, x)$ represents the SHAP value corresponding to feature x_i within the context of a model F constructed using a set of features X . Here, M stands for the total number of input features, X' denotes the set of all potential feature combinations that include feature x_i , and $|Z|$ represents the number of features within a specific feature combination denoted as Z . Furthermore, $F(Z)$ and $F(Z \setminus x_i)$ correspond to distinct predictive models trained on $|Z|$ and Z with feature x_i removed, respectively. Therefore, the SHAP value is determined by aggregating the marginal contributions $F(Z) - F(Z \setminus x_i)$ from all feasible feature combinations (Z) through a weighted average.

SHAP values are calculated using the Python implementation of SHAP, as described by Lundberg and Lee in [25]. In this study, we computed SHAP values for several candidate models.

Figure 2 is the illustration of SHAP values for the random forest model (best model). From the plot, the features “S1_Light” and “S2_Light” have notable impact in forecasting the room occupancy. Their impact are much higher than the other features.

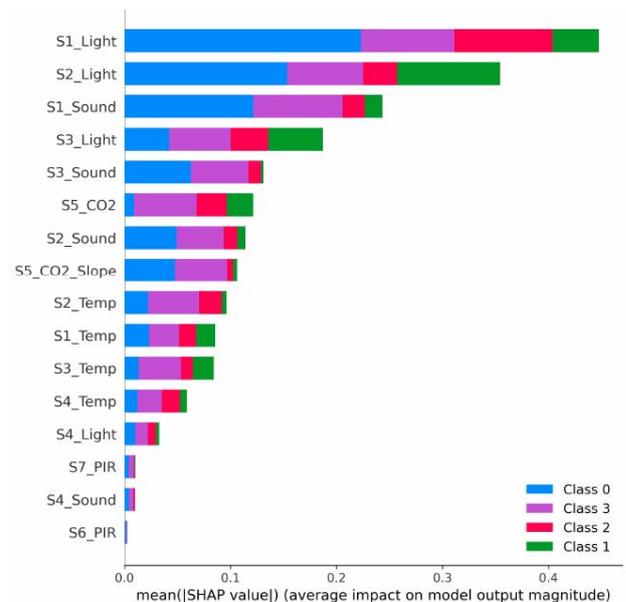


Figure 2. Average Impact (SHAP value) on model output for Random Forest

5. Discussion: Redundancy in Temporal Features

Throughout our study, the dataset has exhibited both temporal and spatial dimensions. Surprisingly, our prediction framework has demonstrated strong performance without explicitly incorporating the temporal dimension.

Our analysis reveals intriguing insights into the temporal aspects of our dataset, shedding light on the potential redundancy and high correlation (illustrated in Figure 3) between certain features. Notably, the autocorrelation results (detailed in Table 5) of our numeric columns provide evidence of the underlying temporal

dynamics. For instance, let's consider the 'S5_CO2' column, which pertains to carbon dioxide levels. It exhibits an extraordinarily high autocorrelation coefficient of approximately 0.999. This remarkable finding underscores that the current CO₂ level is intimately connected with its past values. Essentially, as more individuals enter a room, the 'S5_CO2' readings tend to follow a consistent pattern of escalation, while their departure leads to a noticeable decline in the readings. Similarly, the 'S5_CO2_Slope' representing the rate of change in CO₂ levels, also demonstrates a strong positive autocorrelation with a coefficient of about 0.998. This implies that the temporal aspect of our data, encapsulated by the CO₂ slope, is inherently linked to its past values. As more individuals arrive, the CO₂ slope tends to exhibit a higher positive value, whereas their departure results in a more pronounced negative slope. Furthermore, it's intriguing to note that various other sensors, such as temperature and light sensors ('S1_Temp,' 'S2_Temp,' 'S3_Light,' etc.), display substantial autocorrelation coefficients, indicating consistent temporal patterns in their readings over time. This pattern of high autocorrelation coefficients across multiple sensors suggests that our dataset may indeed have inherent redundancy or temporal dependencies.

In essence, our findings support the notion that the temporal and spatial cues within the dataset might be subtly embedded within other features. The machine learning model can adeptly learn and leverage these temporal dynamics to make accurate predictions, emphasizing the richness of information present in our data.

Furthermore, our dataset exhibits a notably high sampling frequency in the temporal dimension, with data

points recorded at 30-second intervals. This high-frequency data collection could potentially empower traditional models to implicitly incorporate temporal and spatial information, capitalizing on the fine granularity of the data points.

Table 5. Autocorrelation Coefficients of the non-time Features

Autocorrelation	Results:
S1_Temp:	0.9968
S2_Temp:	0.9982
S3_Temp:	0.9975
S4_Temp:	0.9939
S1_Light:	0.9975
S2_Light:	0.9955
S3_Light:	0.9970
S4_Light:	0.9996
S1_Sound:	0.4947
S2_Sound:	0.5516
S3_Sound:	0.5672
S4_Sound:	0.4689
S5_CO2:	0.9999
S5_CO2_Slope:	0.9978
S6_PIR:	0.6015
S7_PIR:	0.7439
Room_Occupancy_Count:	0.9962

The time series plots in Figure 4 illustrate the temporal patterns of various sensor data columns. These patterns are crucial in understanding the concept of redundant time information, where certain features exhibit high autocorrelation, indicating a strong linear relationship with their past values.

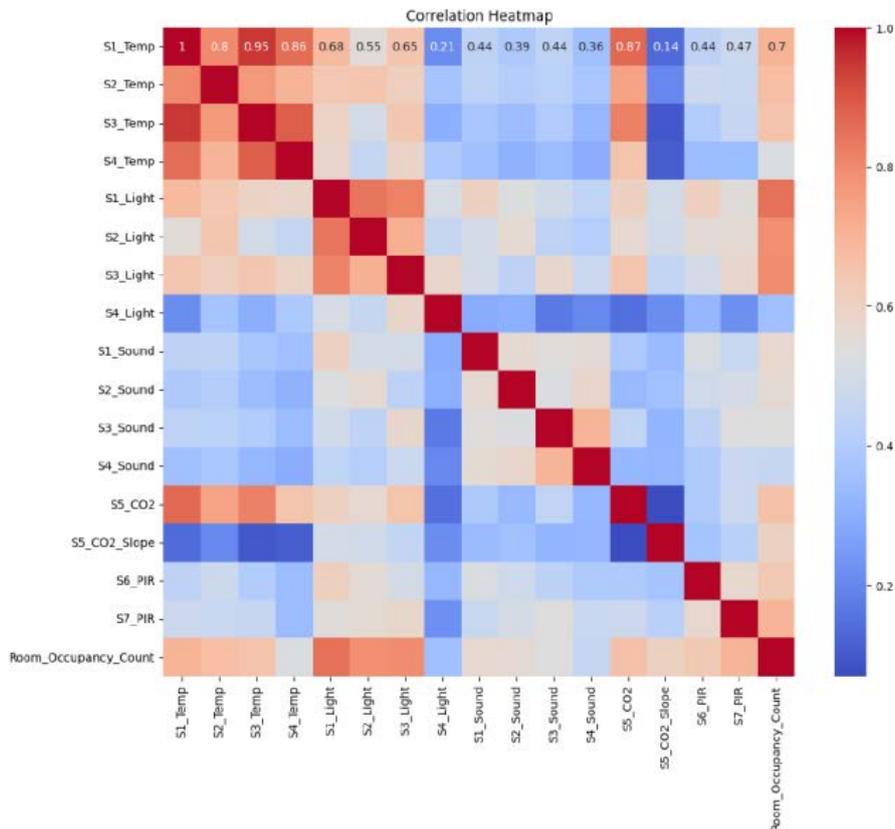


Figure 3. Flowchart of The Proposed Room Occupancy Prediction Framework

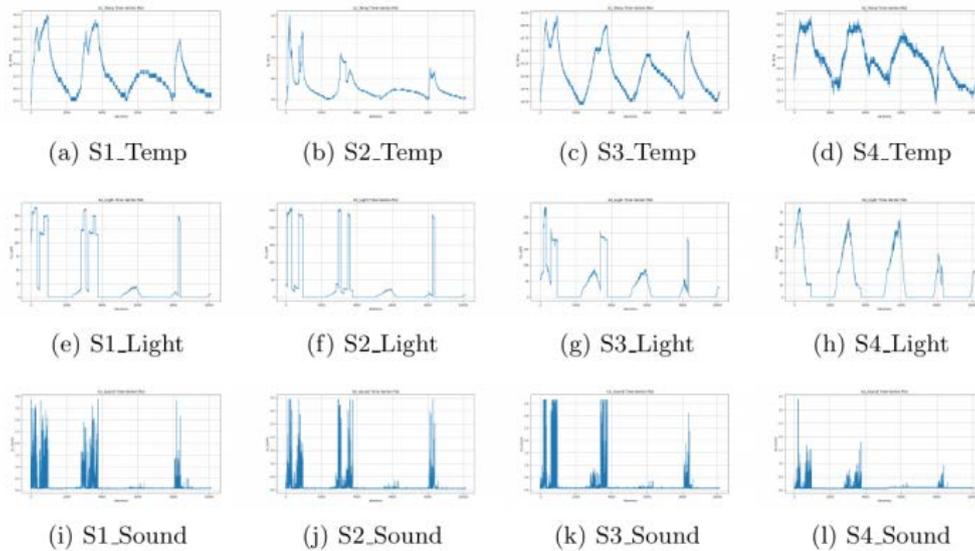


Figure 4. Time Series Plots for Selected Sensors

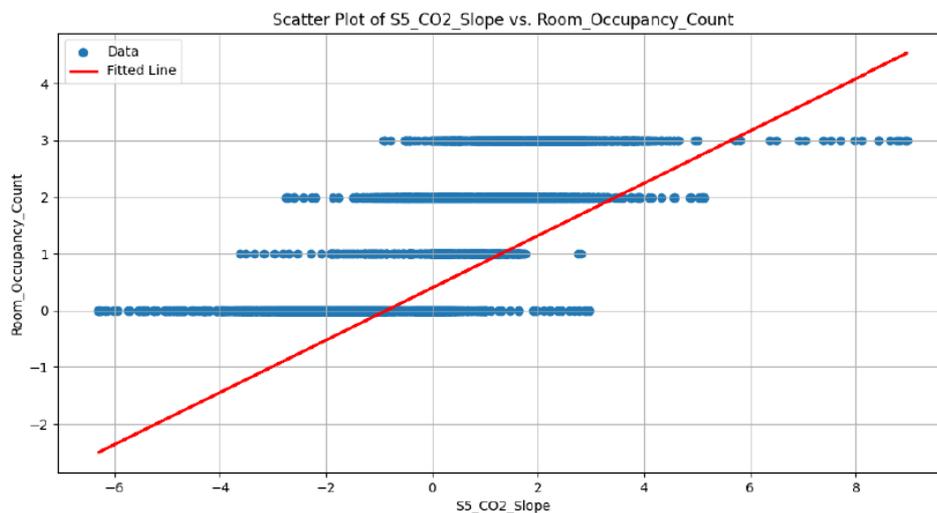


Figure 5. Time Series Plots for Selected Sensors

For instance, consider Figure 4a, which shows the time series plot for 'S1_Temp,' representing the temperature at Sensor 1. The consistently high autocorrelation observed in this plot suggests a significant temporal dependency, implying that the current temperature is closely related to its past values. Similar patterns can be observed in other sensors, such as 'S2_Temp,' 'S3_Temp,' and 'S4_Temp.'

These findings emphasize that the dataset contains valuable temporal information that traditional machine learning models can leverage, even when the specific significance of each feature is not explicitly defined. This temporal redundancy plays a vital role in understanding and predicting the behavior of our data.

Another noteworthy observation is that the spatial and temporal patterns within the dataset may follow relatively straightforward and linear relationships. In such cases, traditional machine learning models like logistic regression or decision trees might suffice in capturing these patterns without necessitating specialized spatial or temporal modeling. However, without comprehensive knowledge of the data collection, recording, and sensor aggregation procedures, we may not directly discern these patterns.

In Figure 5, we present a scatter plot depicting the relationship between the 'S5_CO2_Slope' and the 'Room_Occupancy_Count'. Each data point on the plot represents a specific measurement, where the x-axis corresponds to the 'S5_CO2_Slope' values, and the y-axis represents the 'Room_Occupancy_Count'.

Upon visual inspection, it becomes evident that there is a discernible linear increasing trend in the data points. This trend indicates that as the 'S5_CO2_Slope' value increases, there is a corresponding rise in the 'Room_Occupancy_Count'. The linear regression line fitted to the data further reinforces this observation.

The linear regression analysis provides numerical support for the observed trend. The intercept value of approximately 0.4008 indicates the expected 'Room_Occupancy_Count' when the 'S5_CO2_Slope' is zero. The positive slope of approximately 0.4611 quantifies the rate at which the 'Room_Occupancy_Count' increases concerning changes in the 'S5_CO2_Slope'. This suggests that, on average, for each unit increase in the 'S5_CO2_Slope' we can anticipate an increase of approximately 0.4611 in the 'Room_Occupancy_Count'.

This observation of a linear association between the 'S5_CO2_Slope' and the 'Room_Occupancy_Count' aligns with the idea that the dataset may exhibit linear spatial and temporal patterns. These linear patterns may be indicative of straightforward cause-and-effect relationships within the data. In cases like these, traditional machine learning models, such as linear regression, can effectively capture these patterns without the need for specialized spatial or temporal modeling techniques.

It's important to note that while we can observe these linear patterns, a comprehensive understanding of the data collection, recording, and sensor aggregation procedures is essential for precise interpretation. Without such knowledge, we may not be able to fully elucidate the underlying mechanisms driving these patterns.

In summary, the robust performance of traditional machine learning models, even in the absence of explicit consideration for the temporal dimension, can be attributed to the presence of feature redundancy, the simplicity of linear spatial and temporal patterns, and the advantages offered by a high-frequency data sampling approach. While these findings are intriguing, it is essential to remain open to the possibility that explicitly modeling the temporal dimension could unveil deeper insights or enhance predictive capabilities in certain scenarios.

6. Conclusion and Future Work

In our investigation of room occupancy prediction, we generalized the framework presented in [23], utilizing various machine learning approaches. Differing from its application in [23], our focus was on room occupancy prediction as a classification problem. Our study demonstrates that this framework is adaptable to classification tasks and yields excellent predictive performance. Notably, among all the methodologies we examined, the random forest model stood out with the best performance.

An interesting observation is that, when compared to the original paper of the data source [24], our predictive performance surpassed the original paper's results, even without accounting for time dependencies. Our analysis unveiled valuable insights into the temporal aspects of our dataset, highlighting the potential redundancies and significant correlations among certain features. Particularly, the autocorrelation results from our numeric columns provided evidence of underlying temporal dynamics. Furthermore, our dataset boasted a notably high temporal sampling frequency, with data points recorded at 30-second intervals. This high-frequency data collection potentially empowers traditional models to implicitly integrate temporal and spatial information, making the most of the fine granularity of the data points.

Our findings challenge the assumption made in the original paper that incorporating time-dependent and spatial features would invariably enhance predictive model accuracy.

Conversely, we observed limitations in the data collection process outlined in the original paper [24]. In that paper, sensor nodes were deployed at corners, the middle, and the ceiling of the room to predict the number of people in the room based on sensor readings. However, the position of individuals within the room was not considered. A straightforward implication is that by taking

into account the positions of people, more accurate sensor data could be obtained, potentially leading to improved prediction performance. For future research, we plan to redesign the data collection process, considering the positions of people. In this context, we may be able to predict not only the presence but also the specific number of individuals in the room, transforming the task from a simple classification problem to a more detailed occupancy estimation.

ACKNOWLEDGMENTS

The source code for this project is available on GitHub at the following URL: <https://github.com/YinpuLi/room-occupancy-prediction>. The repository contains the implementation of the algorithms and methods comparison discussed in this paper.

References

- [1] Cao, Xiaodong, Xilei Dai, and Junjie Liu. "Building energy-consumption status worldwide and the state-of-the-art technologies for zero-energy buildings during the past decade." *Energy and buildings* 128 (2016): 198-213.
- [2] Huovila, Pekka, et al. "Buildings and climate change: Summary for decision-makers." (2009).
- [3] Jacobson, Mark Z. "Review of solutions to global warming, air pollution, and energy security." *Energy & Environmental Science* 2.2 (2009): 148-173.
- [4] Li, Nan, and Burcin Becerik-Gerber. "Performance-based evaluation of RFID-based indoor location sensing solutions for the built environment." *Advanced Engineering Informatics* 25.3 (2011): 535-546.
- [5] Wang, Wei, et al. "Occupancy prediction through Markov based feedback recurrent neural network (M-FRNN) algorithm with WiFi probe technology." *Building and Environment* 138 (2018): 160-170.
- [6] Tekler, Zeynep Duygu, Raymond Low, and Lucienne Blessing. "An alternative approach to monitor occupancy using bluetooth low energy technology in an office environment." *Journal of Physics: Conference Series*. Vol. 1343. No. 1. IOP Publishing, 2019.
- [7] Nassif, Nabil. "A robust CO2-based demand-controlled ventilation control strategy for multi-zone HVAC systems." *Energy and buildings* 45 (2012): 72-81.
- [8] Raykov, Yordan P., et al. "Predicting room occupancy with a single passive infrared (PIR) sensor through behavior extraction." *Proceedings of the 2016 ACM international joint conference on pervasive and ubiquitous computing*. 2016.
- [9] Shih, Oliver, and Anthony Rowe. "Occupancy estimation using ultrasonic chirps." *Proceedings of the ACM/IEEE Sixth International Conference on Cyber-Physical Systems*. 2015.
- [10] Uziel, Sebastian, et al. "Networked embedded acoustic processing system for smart building applications." *2013 Conference on Design and Architectures for Signal and Image Processing*. IEEE, 2013.
- [11] Liu, Dixin, et al. "Measuring indoor occupancy in intelligent buildings using the fusion of vision sensors." *Measurement Science and Technology* 24.7 (2013): 074023.
- [12] Tekler, Zeynep Duygu, et al. "Plug-Mate: An IoT-based occupancy-driven plug load management system in smart buildings." *Building and Environment* 223 (2022): 109472.
- [13] Jordan, Michael I., and Tom M. Mitchell. "Machine learning: Trends, perspectives, and prospects." *Science* 349.6245 (2015): 255-260.
- [14] Li, Yinpu, Antonio R. Linero, and Jared Murray. "Adaptive conditional distribution estimation with Bayesian decision tree ensembles." *Journal of the American Statistical Association* 118.543 (2023): 2129-2142.
- [15] Linero, Antonio R., et al. "Bayesian survival tree ensembles with submodel shrinkage." *Bayesian Analysis* 17.3 (2022): 997-1020.

- [16] Chen, Diana Q., Si-Qi Mao, and Xu-Feng Niu. "Tests and classification methods in adaptive designs with applications." *Journal of Applied Statistics* 50.6 (2023): 1334-1357.
- [17] Mao, Siqi. *Time Series and Machine Learning Models for Financial Markets Forecast*. The Florida State University, 2022.
- [18] Kadouche, Rachid, Belkacem Chikhaoui, and Bessam Abdulrazak. "User's behavior study for smart houses occupant prediction." *annals of telecommunications-Annales des télécommunications* 65 (2010): 539-543.
- [19] Sangogboye, Fisayo Caleb, Kenan Imamovic, and Mikkel Baun Kjærgaard. "Improving occupancy presence prediction via multi-label classification." *2016 IEEE International Conference on Pervasive Computing and Communication Workshops (PerCom Workshops)*. IEEE, 2016.
- [20] Razavi, Rouzbeh, et al. "Occupancy detection of residential buildings using smart meter data: A large-scale study." *Energy and Buildings* 183 (2019): 195-208.
- [21] Park, Seunghyeon, et al. "Lstm-based office occupancy detection using smart plug data." *2021 International Conference on Information and Communication Technology Convergence (ICTC)*. IEEE, 2021.
- [22] Koklu, M., and K. Tutuncu. "Tree based classification methods for occupancy detection." *IOP Conference Series: Materials Science and Engineering*. Vol. 675. No. 1. IOP Publishing, 2019.
- [23] Li, Yinpu, et al. "Beyond Tides and Time: Machine Learning's Triumph in Water Quality Forecasting." *American Journal of Applied Mathematics and Statistics* 11.3 (2023): 89-97.
- [24] Singh, Adarsh Pal, et al. "Machine learning-based occupancy estimation using multivariate sensor nodes." *2018 IEEE Globecom Workshops (GC Wkshps)*. IEEE, 2018.
- [25] Lundberg, Scott M., and Su-In Lee. "A unified approach to interpreting model predictions." *Advances in neural information processing systems* 30 (2017).



© The Author(s) 2024. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).