Energy & Buildings 327 (2025) 115086

Contents lists available at ScienceDirect

Energy & Buildings

journal homepage: www.elsevier.com/locate/enbuild

Optimizing indoor environmental prediction in smart buildings: A comparative analysis of deep learning models

Roupen Minassian *, Adriana-Simona Mihăiţă, Arezoo Shirazi

University of Technology Sydney, 61 Broadway Str, Sydney, NSW 2007, Australia

ARTICLE INFO

Keywords:
Timeseries forecasting
Deep learning
Smart buildings
Indoor environmental quality
Temperature prediction
Spatio-temporal data

ABSTRACT

This paper presents a comprehensive investigation into the application of deep learning models for predicting indoor environmental quality in smart buildings. Using data collected from a network of microclimate sensors deployed across a university campus in Sydney, we evaluated the performance of Convolutional Neural Network (CNN), Long Short-Term Memory (LSTM), and hybrid CNN-LSTM models. Our study encompassed various aspects of model development, including data preparation, architecture design, hyperparameter optimization, and model interpretability. Contrary to common assumptions in time series forecasting, our results demonstrate that CNN models consistently outperformed LSTM and hybrid models in predicting indoor temperature. We found that multivariate input configurations enhanced prediction accuracy across all model types, highlighting the importance of capturing complex interactions between environmental parameters. Through SHapley Additive exPlanations (SHAP) analysis, we identified temperature, humidity, and Heating, Ventilation, and Air Conditioning (HVAC) status as the most influential features for predictions. Our experiments also revealed optimal configurations for historical input length and prediction horizon, providing practical guidelines for model implementation. This research contributes valuable insights for the development of more efficient and accurate smart building management systems, potentially leading to improved energy efficiency and occupant comfort in built environments.

1. Introduction

Buildings account for over one-third of global energy consumption, making energy efficiency critical for reducing carbon emissions [21]. Enhancing HVAC systems through demand-driven control strategies, which adjust operations based on real-time occupancy and environmental data, can significantly improve efficiency [25]. These strategies require comprehensive data on current and future conditions, typically gathered from advanced sensor networks.

The proliferation of ubiquitous sensing devices and the growing prevalence of crowdsourcing have paved the way for the collection of vast amounts of spatio-temporal data [7]; [45]. These data sources often exhibit intricate patterns, including cyclical variations due to seasonal effects and spontaneous changes arising from external factors. Accurately forecasting such patterns holds immense practical value across numerous domains, from urban planning to building and infrastructure energy management and efficiency enhancement. Modern HVAC systems leverage Internet of Things (IoT) devices and machine learning to predict temperature, humidity, and other microclimate variables, enabling more responsive management. Various machine learning techniques, including Artificial Neural Network (ANN), k-Nearest Neighbor (KNN), and Support Vector Machine (SVM), have shown promise in optimizing HVAC performance [36]; [4]. For example, neural network and particle swarm optimization algorithm showed that it is possible to reduce HVAC systems' energy consumption by 7.8 percent without compromising indoor environmental conditions [2]. In commercial buildings, machine learning-based occupancy predictions have resulted between 7 and 52 percent energy savings as compared to the conventionally-scheduled cooling systems [34].

Other studies used ANN for predicting hourly indoor air temperature and relative humidity in modern building in humid regions to good accuracy [29]. Considering breadth of micro-climate inputs in modeling algorithms, a comprehensive environmental sensing testbed at Carnegie Mellon University's Intelligent Workplace integrates advanced IT systems and sensing technologies, including sensors for Carbon Dioxide (CO2), Carbon Monoxide (CO), Total Volatile Organic Compounds

* Corresponding author.
E-mail address: roupen.minassian@uts.edu.au (R. Minassian).

https://doi.org/10.1016/j.enbuild.2024.115086
Received 27 August 2024; Received in revised form 17 November 2024; Accepted 18 November 2024
Available online 22 November 2024

(TVOC), Particulate Matter (PM), and acoustics, to achieve energy efficiency and provide a productive environment [14]. Overall, there have been a wide array of experimental approaches that can contribute to model performance.

Although machine learning methods have proven widely applicable in controlling buildings' indoor environment, advanced deep learning models have not yet been investigated properly in this field [14]. Deep learning has demonstrated remarkable capabilities in extracting complex patterns from diverse data sources, including spatio-temporal datasets [45]. Deep learning models, such as CNN, LSTM, and their hybrid variants, have shown promising results in modeling and predicting a wide range of time series phenomena [26]; [18]. However, the performance of these models can be significantly influenced by the choice of input features, model architecture, and the temporal characteristics of the target variables.

Existing research has explored the application of deep learning to time series forecasting, highlighting both the benefits and limitations of various model architectures and input configurations [8]; [20]. For instance, Benidis et al. [8] emphasizes the importance of appropriate normalization techniques to ensure the comparability of input features, while Faloutsos et al. [16] underscores the challenges of leveraging large and diverse data sources for building effective forecasting systems.

Recent research has highlighted the complex interplay of factors affecting indoor environmental quality, energy efficiency, and occupant comfort in buildings. Ma et al. [28] conducted a comprehensive review of variables and models for thermal comfort and indoor air quality, emphasizing the importance of considering a wide range of parameters beyond just temperature and humidity. Their study identified key factors such as air velocity, mean radiant temperature, and personal factors (e.g., clothing insulation and metabolic rate) as crucial for accurate thermal comfort predictions. Additionally, they highlighted the significance of indoor air pollutants, including particulate matter, Volatile Organic Compound (VOC), and CO2, in determining overall indoor environmental quality. Ganesh et al. [19] further investigated the factors affecting human comfort in indoor environments, stressing the need for a broad approach that considers thermal, visual, and acoustic comfort alongside indoor air quality. Their critical review emphasized the importance of integrating these diverse factors into building management systems to optimize occupant well-being and productivity. Yang et al. [51] demonstrated the potential of combining machine learning techniques with model predictive control to optimize both building energy efficiency and comfort. Their adaptive approach showed significant improvements in energy savings while maintaining or enhancing occupant comfort levels. In the context of energy efficiency and comfort optimization, Brandi et al. [10] explored the application of deep reinforcement learning for indoor temperature control and heating energy consumption. Their research showcased the potential of advanced AI techniques in balancing the often conflicting goals of energy conservation and occupant comfort. Furthermore, Dimitroulopoulou et al. [12] provided a comprehensive appraisal of indoor air quality guidelines from around the world, considering their implications for energy saving, health, productivity, and comfort. Their work underscores the importance of addressing indoor environmental quality comprehensively, taking into account regional variations and evolving standards. These studies collectively highlight the multifaceted nature of indoor environmental quality and the need for sophisticated modeling approaches that can capture the complex interactions between various environmental parameters, occupant behavior, and building systems. By incorporating insights from this body of research, our study aims to develop more accurate and comprehensive predictive models for indoor temperature and environmental quality, ultimately contributing to the advancement of smart building management systems.

This paper contributes to the field of smart building management by systematically exploring the application of deep learning techniques to predict indoor environmental quality. Our work distinguishes itself through a comprehensive comparative analysis of multiple deep learning models (CNN, LSTM, and hybrid CNN-LSTM) for temperature prediction, addressing a gap in existing literature where such extensive comparisons are often lacking. We evaluate various input configurations and conduct a thorough parametric analysis, offering insights into optimal model implementations for real-world scenarios [40,38]. Unlike many previous studies that relied on limited datasets, our analysis leverages diverse data from a university campus, encompassing different room sizes, usage patterns, and environmental conditions, ensuring robust and generalizable findings [55]. Furthermore, we emphasize model interpretability through SHAP analysis, providing valuable insights into feature importance [27]. By addressing these aspects collectively, our research provides a more holistic understanding of deep learning applications in indoor environmental quality prediction, potentially leading to improved energy efficiency and occupant comfort in built environments [37]; [23].

Our focus on accurate temperature prediction serves as a crucial stepping stone towards the development of fully integrated smart building management systems. By establishing a robust foundation for environmental forecasting, we enable the creation of more sophisticated control algorithms that can simultaneously optimize energy efficiency and occupant comfort [23]. Accurate temperature predictions allow building systems to anticipate thermal needs, proactively adjust HVAC operations, and maintain optimal comfort levels while minimizing energy consumption [1]. This predictive capability is essential for the next generation of intelligent building systems that can adapt in real-time to changing environmental conditions and occupant requirements [11]. Moreover, the insights gained from our comparative analysis of deep learning models and input configurations provide valuable guidelines for researchers and practitioners working towards integral building management solutions that integrate multiple environmental parameters beyond temperature [46]. By demonstrating the effectiveness and interpretability of deep learning approaches in this critical aspect of building management, our work contributes to the ongoing transition towards more sustainable, efficient, and comfortable built environments.

2. Methodology

2.1. Data collection

In this study, a comprehensive data collection setup was established using Hibou portable microclimate sensors within a city-based university campus in Sydney, Australia, from September 2022 to July 2023. The deployment included ten indoor sensors in classrooms, computer labs and lecture theaters to monitor air quality parameters. The floor layout and the spaces equipped with sensors are displayed in Fig. 1. The building has a central air duct system that uses ductwork to circulate cooled or heated air from a central air conditioning unit or furnace throughout the building, utilizing an Air Handling Unit (AHU) system for temperature control. The system operates on a fixed schedule, running from 6 am to 6 pm on weekdays and remaining off during weekends. The set point temperature is maintained between 22 and 23 degrees Celsius during operational hours. These spaces were occupied during regular university operational hours, and there were no significant interfering activities like cooking or smoking.

The sensors measured and recorded localized temperature, humidity, lighting, ambient pressure and particulate matters. They recorded data at two-minute intervals, continuously transmitting it to a cloud-based database via a secure WiFi connection, facilitated by a separate DeviceNet protocol. Table 1 presents the specifications of the Hibou indoor and outdoor sensors used in this study. Additionally, occupancy was measured using people counting sensors (XOVIS 3D PC) installed at room entrances, employing computer vision technology. The manufacturer reports a minimum 98% accuracy. These sensors recorded data at half-hourly intervals, which was then resampled to hourly frequency to align with other environmental measurement granularity.

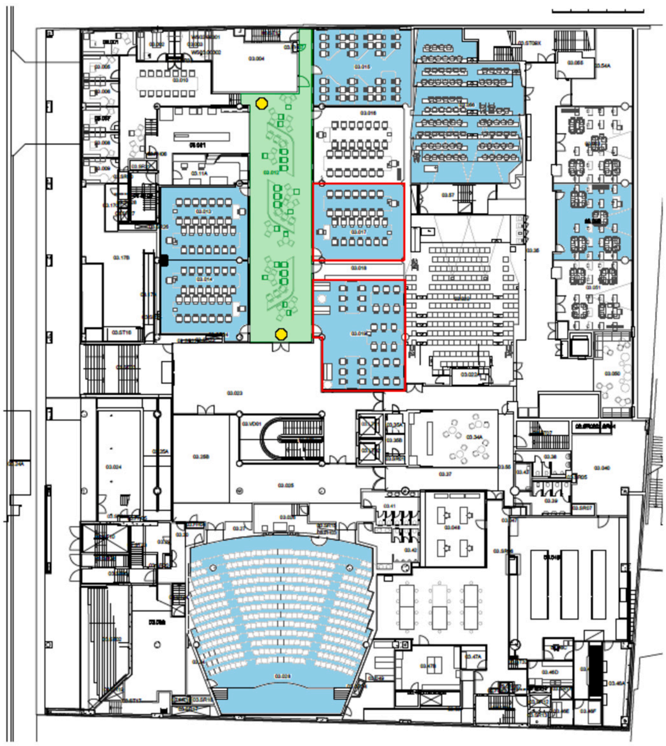

Fig. 1. Building floor layout. Rooms equipped with sensors are marked in blue. Two specific rooms are highlighted with red borders: room 03.017 is positioned in the northern section, with room 03.019 situated immediately to its south. These rooms flank a shared common space (marked in green) which contains two environmental sensors (indicated by yellow markings). The north most sensor of the communal space is positioned at the back of the common space, while the southern sensor is located near the entrance point of this area. Sensors throughout the building are placed to ensure comprehensive environmental data collection. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

Table 1
The specifications for Hibou indoor sensors (https://www.hibouair.com/).

Parameter	Resolution/Output	
Particulate Matter (PM2.5)	resolution 0.3 µg/m3.	Max Error ±10%
Humidity range	0...100% R.H.	Accuracy: ±3% R.H.
Temperature range	0...50 °C.	Accuracy: ±1 °C
Pressure range	300...1100 hPa.	Accuracy ±0.6 hPa
Ambient light	resolution 100mLux	as U.V. Index (WHO standard)
Volatile Organic Compounds	Air Quality Index	

2.2. Data preparation

Data integrity was maintained through live data uploads to the Hibou web interface, with sensor registration on DeviceNet ensuring secure and consistent data exchange. For data retrieval, researchers accessed the sensor data through an Application Programming Interface (API). To address instances of missing data, a KNN Imputer technique was employed, leveraging correlated sensor behaviors to estimate missing values effectively. The initial phase of our methodology focuses on the preparation and splitting of the dataset. This ensures that the data is properly scaled, sequenced, and divided into training, validation, and test sets.

2.2.1. Enhancing temperature measurement precision

While each Hibou sensor has a manufacturer-stated accuracy of ±1 °C, we leverage high-frequency sampling at 2-minute intervals, and hourly data aggregation to substantially improve measurement precision. This method allows us to capture fine-grained temperature fluctuations and minimize the impact of individual measurement errors (see Section 3.1).

2.2.2. Normalization

Both the features and the target variable are normalized to a range between 0 and 1 using MinMax scaling. This normalization is crucial for improving model convergence during training.

2.2.3. Sequence preparation

Sequences of data are created by sliding a window of a specified history length over the normalized dataset, capturing temporal dependencies. Each sequence's target is the value at a specified prediction length following the history window. In our experiments, the historical length and prediction length both range from 1 to 12 time steps with hourly granularity.

The historical length includes all time points in that range, while the prediction length is a point forecast of that specific time point and does not include all previous time points.

Input-output pairs are generated by iterating over the dataset, where each input sequence consists of time steps of features, and the corresponding output is the target value. The dataset is divided into training, validation, and test sets based on predefined ratios (60% training, 20% validation, 20% testing). This split ensures that the model can be trained and validated effectively while reserving an unbiased test set for final evaluation.

2.3. Parameter selection

Based on the comprehensive literature review and considering the complex interplay of factors affecting indoor environmental quality, energy efficiency, and occupant comfort [28,19,51,10,12], we have selected the following parameters for our machine learning model:

- **Temperature**: This is the target variable for our prediction model. Temperature is the main driver of thermal comfort, with optimal indoor comfort generally being achieved between 20-25 °C [5]. As highlighted by Yang et al. [51], accurate temperature prediction is crucial for both energy efficiency and occupant comfort.
- **Humidity**: Humidity significantly affects perceived temperature by influencing the body's ability to cool itself through evaporation. Ma et al. [28] identified humidity as a key factor in thermal comfort models. The ideal range for comfort is typically between 30-60% [17,44].
- **Air Pressure**: While not directly impacting comfort, changes in atmospheric pressure can affect indoor air circulation, which may influence temperature stability and HVAC performance [5]. Ganesh et al. [19] noted the importance of considering such indirect factors in holistic indoor environment quality assessments.
- **Ambient Light**: As discussed by Ganesh et al. [19], visual comfort is an integral part of overall indoor environmental quality. Ambient light, especially natural sunlight, contributes to heat gain within indoor spaces and can significantly impact both energy consumption and occupant comfort [43].
- **Volatile Organic Compounds (VOC)**: VOC are crucial indicators of indoor air quality, as emphasized by Ma et al. [28] and Dimitroulopoulou et al. [12]. Poor air quality can degrade comfort and indirectly affect temperature by reflecting the efficiency of ventilation systems [47].
- **Particulate Matter (PM1, PM2.5, PM10)**: Dimitroulopoulou et al. [12] highlighted the importance of considering particulate matter in indoor air quality assessments. Airborne particulates can impact HVAC system efficiency and lead to uneven temperature distribution [48].

- **CO2 Levels**: Ma et al. [28] identified CO2 as a key factor in indoor air quality models. High CO2 levels often indicate poor ventilation, which can trap heat inside a building and decrease comfort [39].
- **Booking Status**: To account for scheduled use of spaces, which can help predict occupancy patterns and their impact on indoor environment.
- **People Count**: Direct measurement of occupancy, which Yang et al. [51] demonstrated as significant for both energy efficiency and comfort optimization.
- **HVAC Status**: A binary input indicating active system operation. Brandi et al. [10] showed the importance of considering HVAC operation in models for optimizing indoor temperature control and energy consumption.

By incorporating this diverse set of variables, our models aim to capture the complex interplay of factors affecting indoor temperature, energy efficiency, and occupant comfort, as emphasized in recent literature [28,19,51,10,12]. This comprehensive approach allows us to develop more accurate and robust predictive models, potentially leading to improved smart building management systems that can balance energy efficiency with occupant comfort and well-being.

2.4. HVAC status representation in the model

In our approach to modeling building energy consumption, we have introduced a simplified representation of the HVAC system's operational state, referred to as 'HVAC status'. This section elucidates the rationale behind this representation and its implications for the model's performance.

The 'HVAC status' in our model is defined as a binary input variable that indicates whether the building's HVAC system is actively operating or not during each hour. Specifically:

- A value of 1 indicates that the HVAC system is active and maintaining the indoor environment at the setpoint temperature.
- A value of 0 indicates that the HVAC system is inactive, allowing the indoor temperature to fluctuate more naturally.

The adoption of a binary representation for HVAC system status in our model is underpinned by several key considerations, despite the inherent complexity of HVAC systems. This approach offers a balance of simplicity and effectiveness, providing our model with critical information about active indoor environment control periods versus times of natural temperature fluctuation, without introducing unnecessary input parameter complexity. The binary input significantly enhances the model's capacity to differentiate between active temperature control periods and more passive conditions, thereby improving predictions of energy consumption patterns. Furthermore, this method maintains computational efficiency by avoiding the inclusion of multiple detailed HVAC parameters, striking a balance between model complexity and processing requirements. Importantly, for the building in our study, detailed HVAC operational data beyond its active/inactive state was neither readily available nor essential for achieving our modeling objectives. It is worth noting that the temperature predictions in our study refer to room temperature measured by our independent sensor network, which is separate from the sensors implemented in the HVAC system itself. This distinction is crucial as it allows our model to capture the actual experienced room temperature, which may differ from the HVAC system's set points or internal measurements due to factors such as sensor placement, local temperature variations within the room, and the dynamic response of the space to HVAC operations.

It is noteworthy that our model does not explicitly incorporate setpoint temperature as a distinct feature. This decision was informed by the specific characteristics of the building under investigation. The structure in question operates with a fixed setpoint temperature rather than a dynamic one. In this context, the HVAC status effectively serves as an indicator of the periods during which the system is actively maintaining this constant setpoint temperature. By utilizing the HVAC status as a proxy, we indirectly account for the influence of the setpoint temperature on both the system's operation and energy consumption patterns. This approach allows us to capture the essential thermal management dynamics without the need for additional setpoint-specific variables.

2.5. Model architecture

We investigated prediction models using both univariate and multivariate inputs to compare their capabilities in capturing complex spatio-temporal dependencies. Univariate refers to a prediction model developed using only the target variable as input - in this case temperature - whereas multivariate incorporates all other environmental measurements collected through sensors, as well as data from the building information platforms discussed in the paper. Additionally, we introduced Shapley value analysis to provide novel insights into model interpretability, elucidating the contribution of individual features to prediction accuracy. Our investigation encompasses three distinct deep learning model architectures: CNN, LSTM, and hybrid CNN-LSTM models. Each architecture is tailored to leverage specific strengths in handling spatio-temporal data.

2.5.1. CNN model

The CNN model employs multiple convolutional layers to capture local temporal patterns in the data [24]. Each convolutional layer is followed by max-pooling layers to reduce dimensionality and computational complexity. The final layers are fully connected (dense) layers that map the extracted features to the target variable.

2.5.2. LSTM model

The LSTM model is designed to capture long-term dependencies in the time series data [22]. It consists of multiple LSTM layers, each configured to either return sequences or output only the final state based on the specified layer depth. The model concludes with dense layers that transform the LSTM outputs into the final prediction.

2.5.3. CNN-LSTM hybrid model

This model processes the data through time-distributed convolutional and max-pooling layers, treating each timestep as a separate instance to extract spatial features. The extracted features are then fed into LSTM layers to capture temporal dependencies. Finally, dense layers produce the output prediction (Fig. 2).

2.6. Model training and evaluation

2.6.1. Training configuration

The models are trained and tested using various configurations. A Bayesian optimization approach is utilized to efficiently navigate the hyperparameter search space. The search space includes:

- LSTM units: $\{1, 2, 4, 8, 16, 32, 64, 128\}$
- Learning rate: $[1 \times 10^{-5}, 0.1]$
- Batch size: $\{16, 32, 64\}$
- Filters: $\{32, 64, 128\}$
- Kernel size: $\{3, 5, 7\}$
- CNN kernel size: $\{1, 2, 3\}$
- CNN layers: $\{1, 2, 3, 4, 5\}$
- LSTM layers: $\{1, 2, 3, 4, 5\}$
- Dense units: $\{64, 128, 256\}$
- Optimizer: {adam, sgd, adadelta, lion, ftrl, nadam, adamax, adagrad, rmsprop}
- Activation: {elu, selu, gelu, leaky_relu, relu, tanh}
- Dilation rate: $\{1, 2, 4\}$

Mean Squared Error (MSE) is used as the loss function, and different optimizers (e.g., Adaptive Moment Estimation (Adam), Stochastic

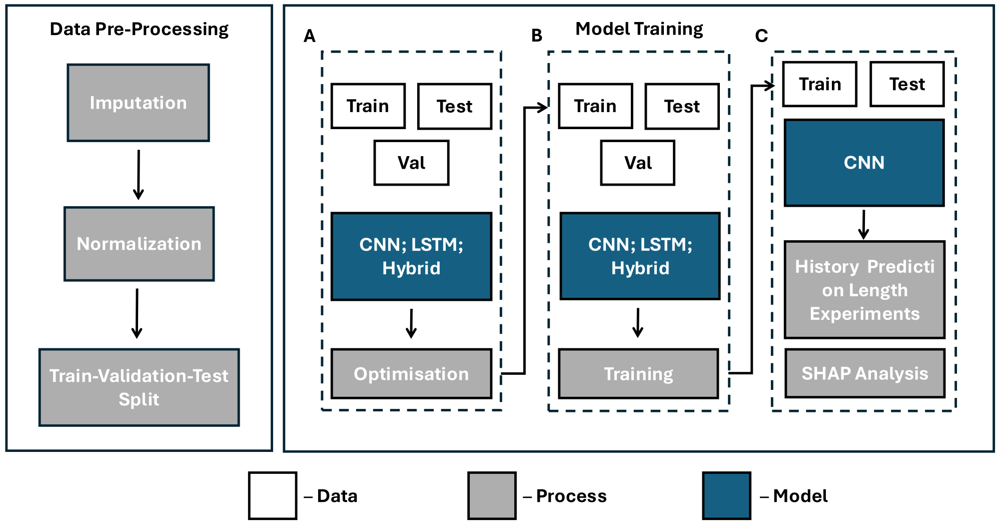

Fig. 2. Methodology workflow for predictive modeling. This figure outlines our study's methodology, which consists of three main stages. The initial stage, data pre-processing, involves imputation of missing data, normalization of features, and splitting the dataset into train, validation, and test sets. Following this, the model training stage implements and trains CNN, LSTM, and hybrid models, optimizing their parameters using the training set and validating performance on the validation set (A, B). The final stage encompasses our experimental approach and analysis (C). Here, we explore various historical input lengths and prediction horizons, test the optimized models on the held-out test set, and apply SHAP analysis for model interpretability. The workflow illustrates an iterative process between the model training and experimental stages, indicating continuous refinement of the models based on experimental results.

Gradient Descent (SGD)) are employed based on the configuration to minimize this loss.

The model is compiled with the selected optimizer and loss function to begin the training process. The model is trained for 50 epochs using the training data, with validation data provided to monitor performance and prevent over-fitting. The training process includes callbacks for logging metrics and saving the best model based on validation loss. Model training metrics and experiments were logged and tracked using Weights & Biases for comprehensive experiment monitoring [9].

2.6.2. Naive predictions

For comparison, naive predictions were created using mean aggregates based on hour-of-day and day-of-week. The data used for these naive predictions was drawn from the validation set - a period prior to the test set, ensuring an unbiased baseline for performance evaluation.

2.7. Model evaluation and analysis

Post-training, the models are evaluated on the test set using metrics such as Root Mean Squared Error (RMSE), Mean Absolute Percentage Error (MAPE), and Symmetric Mean Absolute Percentage Error (SMAPE). These metrics provide a comprehensive assessment of the model's predictive accuracy and robustness.

The RMSE is calculated as follows:

$$\text{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \tag{1}$$

where y_i is the actual value, $\hat{y}_i$ is the predicted value, and n is the number of observations.

The MAPE is calculated as:

$$\text{MAPE} = \frac{100\%}{n}\sum_{i=1}^{n}\left|\frac{y_i - \hat{y}_i}{y_i}\right| \tag{2}$$

where y_i is the actual value, $\hat{y}_i$ is the predicted value, and n is the number of observations.

The SMAPE is calculated as:

$$\text{sMAPE} = \frac{100\%}{n}\sum_{i=1}^{n}\frac{|y_i - \hat{y}_i|}{(|y_i| + |\hat{y}_i|)/2} \tag{3}$$

Table 2
Temperature measurement statistics by sensor.

Sensor ID	Average SEM (°C)	Minimum Temperature (°C)	Maximum Temperature (°C)
01D997	0.020	18.68	27.91
01E60A	0.019	17.64	27.51
052E9D	0.026	18.56	28.24
05500B	0.016	18.41	27.36
0589C3	0.012	18.01	26.78
43344A	0.019	17.81	28.69
0565F6	0.012	18.18	25.00
4336AB	0.022	17.58	28.55
432541	0.015	16.58	31.90
4350B8	0.021	16.13	28.46

where y_i is the actual value, $\hat{y}_i$ is the predicted value, and n is the number of observations.

3. Results

This section presents a comprehensive analysis of the sensors' data and the development of predictive models for temperature forecasting. We begin by examining the diurnal patterns and correlations between various environmental parameters and occupancy. We then progress through the optimization of machine learning models, comparing their performance, and finally delve into the interpretability of our best-performing model. This structured approach allows us to build a holistic understanding of the complex interactions within indoor environments and the effectiveness of our predictive modeling techniques.

3.1. Temperature precision and range of measurements

The effectiveness of our approach of sampling and aggregation methods of the Hibou sensors is evidenced by the low Standard Error of Mean (SEM) values achieved across all monitored rooms. Table 2 presents a comprehensive overview of the average SEM, minimum temperature, and maximum temperature for each room in our study.

The average SEM values range from $\pm 0.012\,°\text{C}$ to $\pm 0.026\,°\text{C}$, which is significantly lower than the $\pm 1\,°\text{C}$ accuracy of individual sensors. Furthermore, the data reveals a wide temperature range across all rooms,

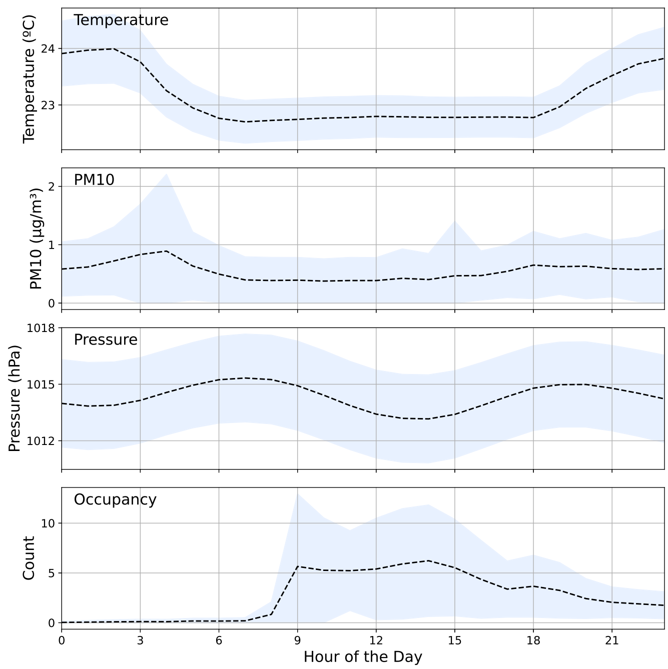

Fig. 3. Diurnal patterns of micro-climate parameters and occupancy in the building. Hourly aggregates of Temperature, PM10, Pressure, and Occupancy count across 10 datasets over a 24-hour period. Data points represent means with 95% confidence intervals. The x-axis shows the hour of the day in 24-hour format, while y-axes depict the respective units for each parameter. Aggregates are derived from sensor measurements and collected building occupancy data.

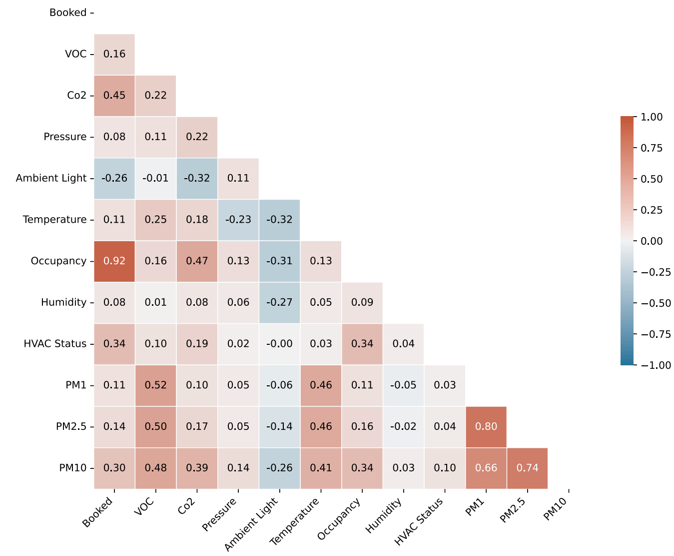

Fig. 4. Micro-climate correlation chart for university room 03.017. Correlation matrix including all environmental parameters. Color intensity represents the strength of correlation, with blue indicating negative correlations and red indicating positive correlations. Room volume is 219.94m^3.

spanning from a minimum of 16.13 °C to a maximum of 31.90 °C. This extensive range underscores the system's capability to detect and record significant temperature variations, further validating the reliability of our measurements.

3.2. Data analysis

Data investigation begins with an in-depth exploration of the temporal patterns and interrelationships of key environmental parameters. This analysis sets the foundation for understanding the dynamic nature of indoor environments and informs our subsequent modeling efforts.

3.2.1. Hourly aggregates

Fig. 3 examines the diurnal patterns of micro-climate parameters and occupancy. This analysis reveals the daily rhythms of indoor environments, highlighting the interplay between human activity and environmental conditions. Below, the distinct diurnal patterns of micro-climate parameters and occupancy are discussed separately for each parameter.

Temperature: Mean indoor temperature shows a clear diurnal cycle, ranging from approximately 23 °C to 24 °C. The lowest temperatures are observed in the early morning (around 07:00), followed by a steady increase throughout the day. Peak temperatures occur in the evening (between 21:00-03:00), after which they gradually decline. This pattern aligns with the HVAC operating hours. The narrow confidence intervals suggest consistent temperature control across the sampled building.

PM_{10}: Particulate matter concentrations (PM_{10}) show subtle variations throughout the day. Levels are highest during the early morning hours (00:00-06:00), likely due to reduced activity within the rooms. A slight decrease is observed during typical working hours (08:00-18:00), associated with active hours HVAC system. The wider confidence intervals during daytime hours suggest greater variability in PM_{10} levels across building spaces during periods of occupancy.

Pressure: Atmospheric pressure within the buildings demonstrates minor fluctuations, ranging from approximately 1012 hPa to 1018 hPa. A slight diurnal pattern is discernible, with lower pressures observed in the afternoon hours and late evening. This pattern may be influenced by outdoor atmospheric conditions. The narrow confidence intervals indicate consistent pressure readings across the sampled sensors.

Occupancy: The occupancy data reveals a strong correlation with typical working hours. Near-zero occupancy is observed from midnight to early morning (00:00-06:00). A sharp increase occurs between 07:00 and 09:00, corresponding to the arrival of building occupants. Peak occupancy is maintained during standard working hours (09:00-17:00), followed by a gradual decline in the evening. The wider confidence intervals during peak hours suggest variability in maximum occupancy levels across different rooms or days.

These findings provide insights into the dynamic interplay between building occupancy patterns and micro-climate parameters, highlighting the potential influence of human activity on indoor environmental conditions. The data underscores the importance of considering temporal variations in building management strategies for optimizing energy efficiency and occupant comfort.

3.3. Correlation analysis

This section explores the correlations between indoor environmental parameters, occupancy, and building HVAC system status across two university spaces (Figs. 4, 5). The analysis reveals complex interactions between various Indoor Environmental Quality (IEQ) parameters, providing insights into building performance and occupant-environment interactions [32].

3.3.1. Particulate matter dynamics

Analysis of PM correlations reveals moderate to strong interconnections between different particle size fractions. The correlations between PM1, PM2.5, and PM10 (r = 0.66 to 0.80 in 03.017 and r = 0.41 to 0.82 in 03.019) align with established indoor aerosol behavior theories [30]. The observed correlation patterns suggest common sources and removal mechanisms affecting different particle size fractions. The relationship between PM levels and temperature shows moderate correlation in 03.017 (r = 0.41 to 0.46) and weak correlation in 03.019 (r

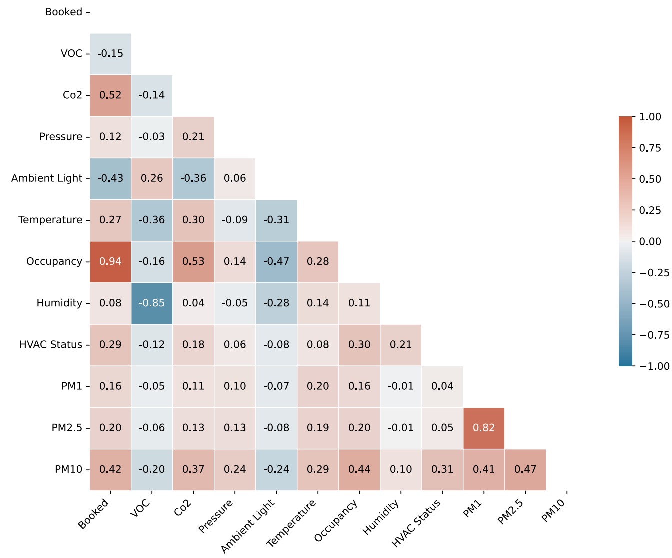

Fig. 5. Micro-climate correlation chart for university room 03.019. Correlation matrix including all environmental parameters. Color intensity represents the strength of correlation, with blue indicating negative correlations and red indicating positive correlations. Room volume is 297.35m^3.

= 0.19 to 0.29), suggesting that temperature may have some influence on particle behavior, though other factors likely play important roles.

3.3.2. *Occupancy-related environmental dynamics*

The strong positive correlations between occupancy and booking status (r = 0.92 for 03.017 and r = 0.94 for 03.019) demonstrate robust space utilization monitoring. CO2 concentrations show moderate positive correlations with occupancy in both spaces (r = 0.45 and r = 0.52, respectively), consistent with established metabolic CO2 generation rates in indoor environments [35]. The moderate correlations between occupancy and PM10 levels (r = 0.34 to 0.44) reflect the impact of human activity on particle resuspension, a phenomenon widely documented in indoor air quality research [50]; [52].

3.3.3. *Thermal-humidity interactions*

The correlation analysis reveals minimal to no relationship between temperature and humidity across the spaces (r = 0.05 in 03.017, r = 0.14 in 03.019). These near-zero correlations suggest that temperature and humidity levels vary independently, indicating effective decoupling of temperature and humidity control in the HVAC systems [13]. Similarly, the pressure-temperature correlations (r = -0.23 and r = 0.09) show negligible relationships, suggesting that pressure variations occur largely independently of temperature changes in these spaces.

3.3.4. *VOC and Co2 behavior*

The analysis of VOC distributions reveals distinct patterns across three monitored spaces (Fig. 6). Room 03.017 exhibits the highest probability density peak (approximately 0.08) for VOC concentrations between 10^0 and 10^1 ppm. Similarly, room 03.019 shows a comparable distribution pattern but with a slightly lower density peak (approximately 0.07). In contrast, the communal space with sensors placed in the front and rear of the room demonstrates markedly different behavior, with a significantly lower density peak (approximately 0.002) and a relatively narrow and uniform distribution at 10^3 ppm. These distribution patterns, when considered alongside the correlation analysis, provide deeper insights into indoor air quality dynamics. In 03.017, VOC levels show moderate correlations with PM levels (r = 0.48 to 0.52), suggesting possible common sources or transport mechanisms. 03.019 shows different relationships, with a strong negative correlation between VOC and humidity (r = -0.85) and a moderate negative correlation with temperature (r = -0.36). Given the proximity to the common space with its

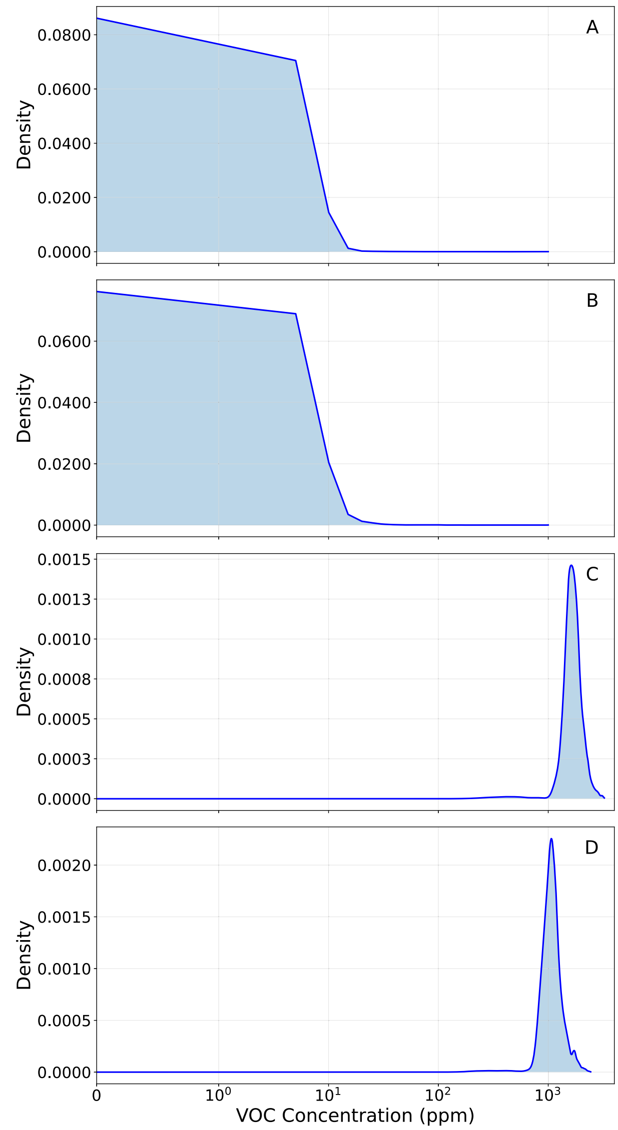

Fig. 6. Probability Density Functions across A) classroom 03.017, B) classroom 03.019, C) northern sensor in university communal space (rear area), and D) southern sensor in university communal space (near entrance). The figure presents VOC concentration distributions using probability density functions plotted on a symmetrical logarithmic scale. Measurements were taken from two rooms (03.017 and 03.019) and their adjacent communal space, with vertical axes showing density values and the horizontal axes displaying VOC concentrations in parts per million (ppm).

higher VOC levels, these contrasting correlations might reflect different air exchange patterns between the classroom and the adjacent communal area.

The contrasting VOC correlations between 03.017 and 03.019 (r = 0.48–0.52 with PM versus r = -0.85 with humidity, respectively) may be attributed to several room-specific dynamics. In 03.017, the positive correlation with PM levels could indicate shared indoor sources, such as human activities that simultaneously generate both particles and VOC (e.g., use of personal care products, cleaning activities, or movement stirring up settled particles). The moderate strength of this correlation suggests that while these sources exist, other factors also influence VOC concentrations independently of PM. In contrast, 03.019's strong negative correlation between VOC and humidity points to potentially different ventilation patterns or room usage characteristics. This inverse relationship might be explained by room orientation and proximity to building exhaust points, local air exchange patterns influenced by door opening frequencies, and distinct usage patterns that affect both VOC sources and humidity levels differently from 03.017. The moderate negative correlation with temperature in 03.019 (r = -0.36) further supports the possibility of ventilation-driven differences, as temperature gradients can influence air movement patterns between spaces. This could be particularly relevant given the room's interaction with the communal space's higher VOC concentrations.

The spatial distribution of VOC in the communal space, as indicated by the different concentrations between the entrance (C) and rear (D) sensors, suggests the presence of a concentration gradient. This gradient could differentially affect rooms 03.017 and 03.019 based on their position relative to the communal space and local air flow patterns. Such concentration gradients can create different diffusion pressures and mixing patterns when doors are opened, potentially contributing to the observed differences in VOC correlations between the two classrooms. This phenomenon has been documented in several studies examining indoor pollutant transport patterns in connected spaces [33]; [35].

CO2 shows weak correlations with temperature across spaces (r = 0.18 to 0.30), suggesting that CO2 levels vary largely independently of temperature. While both CO2 and VOC show distinct distribution patterns in these spaces, their weak correlations with temperature indicate that other factors, such as occupancy patterns and source characteristics, may be more important in determining their concentrations.

This correlation analysis provides valuable insights into the complex interactions governing indoor environmental quality in university buildings. The observed variations in parameter relationships demonstrate the importance of understanding local building physics, HVAC operation, and indoor air quality dynamics [42]. The findings suggest that effective indoor environment management requires comprehensive consideration of ventilation effectiveness, mechanical system performance, building envelope characteristics, and occupant behavior patterns. These results contribute to the growing body of knowledge on indoor environmental quality in educational facilities and underscore the importance of integrated approaches to building systems analysis and control.

3.4. Bayesian optimization for LSTM model configurations

With insights from our data analysis, we proceed to develop predictive models. We employ Bayesian optimization to fine-tune the hyperparameters of our LSTM models, systematically exploring the model configuration space to identify optimal settings for temperature prediction. Due to its efficiency in exploring high-dimensional spaces and capacity to balance exploration and exploitation, this optimization method leads to more effective model performance with fewer iterations compared to traditional grid or random search methods.

Fig. 7 presents parallel coordinates plots for hyperparameter tuning of an LSTM-based model, with purple lines indicating better-performing configurations (lower validation loss). The plots reveal several key trends in the optimal hyperparameter settings.

The activation function analysis showed that high-performing models predominantly use Rectified Linear Unit (ReLU) and Gaussian Error Linear Unit (GELU) activation functions. While some successful models use Hyperbolic Tangent Function (Tanh), it appears less frequently among the top performers. Regarding batch size, a value of 32 is most common among the best-performing models, with some successful configurations also using a batch size of 64. The smaller batch size of 16 is rarely associated with top performance.

Learning rate optimization revealed that the most effective models cluster around values of 0.001 to 0.01. Very low (0.0001) or high (0.1, 1.0) learning rates are generally absent from the best-performing configurations. In terms of LSTM layers, top-performing models tend to have 2 or 3 layers. Deeper networks with 4 or 5 layers appear less frequently among the best configurations, suggesting that moderate depth is preferred for this particular task.

The number of LSTM units also played a crucial role, with high-performing models showing a preference for larger numbers. Many successful configurations use 64 or 128 units, while fewer top performers use 32 units or less, indicating that increased model capacity is beneficial. Optimizer choice proved to be significant, with Adam and Nesterov-accelerated Adaptive Moment Estimation (Nadam) optimizers strongly associated with the best-performing models. Root Mean Square Propagation (RMSprop) also appears among some top configurations, but less frequently than Adam variants. Other optimizers like SGD, Adaptive Gradient Algorithm (Adagrad), and Adaptive Delta (Adadelta) are rarely seen in the best-performing models.

The multivariate plot (Fig. 7b) reveals that the best models often combine ReLU or GELU activation, a batch size of 32, learning rate around 0.001 to 0.01, 2 or 3 LSTM layers, 64 or 128 LSTM units, and Portmanteau of “Adam” and “Maximum” (Adamx) or Nadam optimizer. These findings suggest that while individual hyperparameters show clear trends, their interactions are crucial for achieving optimal performance. The consistent patterns among top performers indicate that focused hyperparameter tuning around these configurations could yield further improvements. Additionally, the results highlight the importance of using adaptive optimizers and sufficient model capacity (through LSTM units) while maintaining a moderate network depth for this specific task.

3.5. Univariate vs. multivariate comparison

Leveraging the optimized model configurations, we now compare the performance of univariate and multivariate approaches. This comparison allows us to assess the value of incorporating multiple environmental parameters in our predictive models.

Fig. 8 presents a comparison of temperature predictions from three LSTM model variants against actual values over a period of approximately 200 hours (timesteps 800-1000 of the test set). As shown in Fig. 8a, the naive model demonstrated the poorest performance among the three. Its predictions (red dashed line) consistently lagged behind the actual temperature trends (blue solid line), often missing both peaks and troughs. The residuals (gray dotted line) exhibited large fluctuations, indicating significant prediction errors throughout the time series.

In contrast, Fig. 8b illustrates that the univariate model showed a marked improvement over the naive approach. The green dashed line, representing the univariate model's predictions, captured the general temperature trends more accurately, more closely following the actual values (blue solid line). However, it still displayed some lag in predicting rapid temperature changes, particularly at troughs. The residuals (gray dotted line) were notably smaller than those of the naive model but still showed some systematic patterns.

Fig. 8c demonstrates that the multivariate model exhibited the best performance among the three. Its predictions (purple dashed line) aligned most closely with the actual temperature values (blue solid line), accurately capturing both the overall trends and the short-term fluctuations. The model successfully predicted most peaks and troughs with minimal lag. The residuals (gray dotted line) were the smallest and most uniform among the three models, indicating consistently low prediction errors across the time series.

The progression from naive to univariate to multivariate models, as visualized in subplots a, b, and c respectively, showed a clear improvement in prediction accuracy. The naive model's poor performance suggested that simple persistence-based predictions were inadequate for the temperature forecasting task. The univariate model's improved accuracy indicated that leveraging historical temperature data significantly enhanced prediction quality. The multivariate model's superior performance demonstrated the value of incorporating additional relevant variables beyond just historical temperature in the prediction process.

The multivariate model's ability to accurately predict temperature fluctuations, including sudden changes, suggested that it effectively captured complex relationships between temperature and other environ-

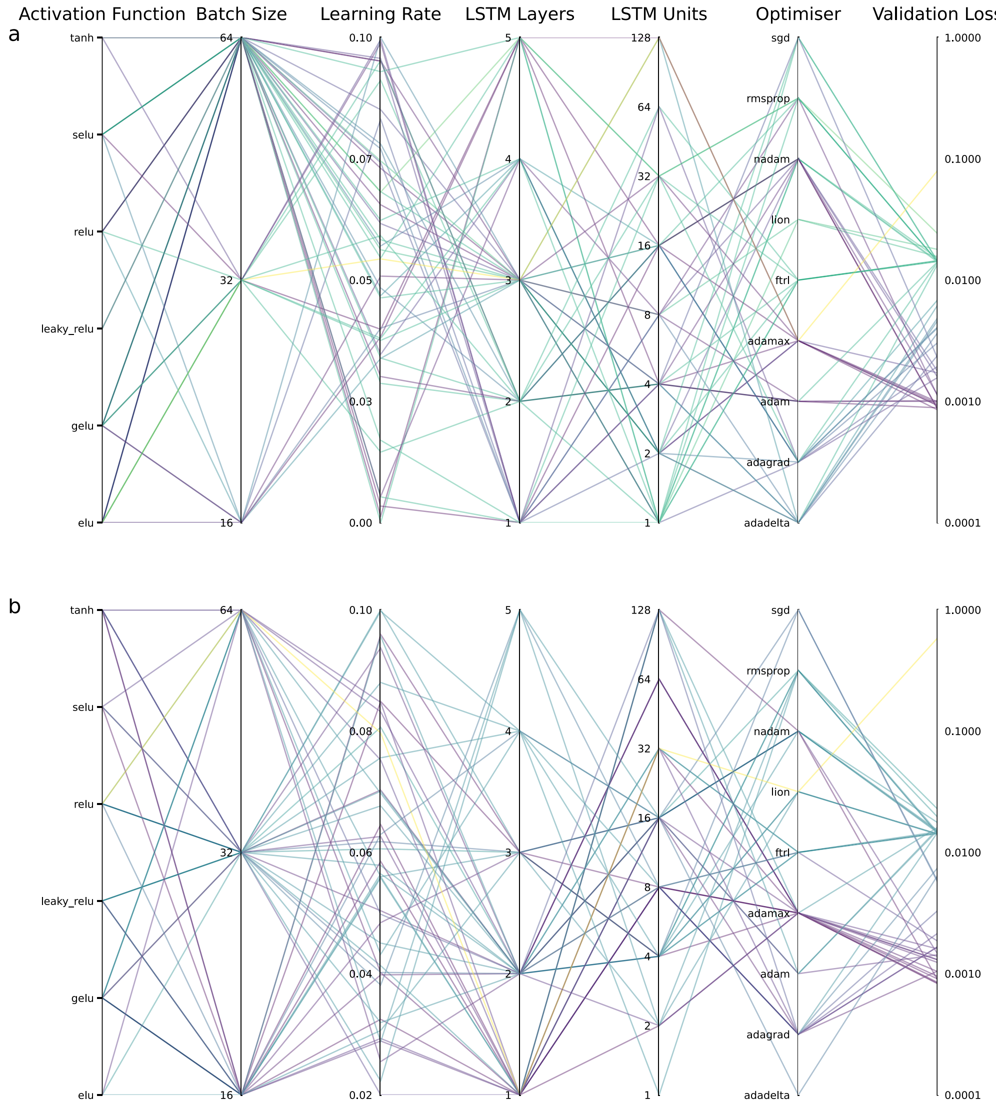

Fig. 7. Parallel coordinates plot for Bayesian optimization of LSTM model configurations. The results of Bayesian optimization for hyperparameter tuning of univariate (a) and multivariate (b) LSTM models. Each line represents a different model configuration, with color indicating the validation loss (purple lines represent better-performing models with lower loss). The plots display the relationships between various hyperparameters including activation function, batch size, learning rate, number of layers, number of units, and optimizer choice. The y-axis for validation loss is shown on the rightmost column. Color intensity corresponds to the logarithmic scale of validation loss.

mental or contextual factors. This improved accuracy, evident from the close alignment of the purple dashed line with the blue solid line in Fig. 8c, could be particularly valuable for applications requiring precise temperature control or forecasting, such as in building energy management systems or climate control in sensitive environments.

These results underscore the importance of model complexity and input richness in time series prediction tasks, particularly for dynamic environmental variables like temperature. The multivariate LSTM approach appears to be the most suitable for this specific temperature prediction task, offering a balance of accuracy and responsiveness to short-term variations, as clearly demonstrated by the comparison in Fig. 8.

3.6. Model performance comparison

Expanding on our univariate-multivariate comparison, we conduct a comprehensive evaluation of different model architectures. This analysis helps us identify the most effective approach for temperature prediction in indoor environments amongst common deep learning architectures.

Fig. 9 presents a comparison of three performance metrics (SMAPE, MAPE, and RMSE) across different model architectures and input types. The boxplots provide a visual representation of the distribution of these metrics.

Examining the SMAPE results in Fig. 9a, we observe that multivariate input consistently outperforms univariate input across all model types. This is particularly evident for the CNN model, where the multivariate median SMAPE (1.38%) is lower than its univariate counterpart (1.42%). The LSTM model shows the most significant improvement with multivariate input, with the median SMAPE decreasing from 1.85% (univariate) to 1.59% (multivariate).

The MAPE results, depicted in Fig. 9b, reveal a similar trend. The CNN model demonstrates the best performance, with a multivariate median MAPE of 1.39%, compared to 1.42% for univariate input. The CNN-LSTM and LSTM models also benefit from multivariate input, with notable reductions in median MAPE values.

Fig. 9c illustrates the RMSE results, which align with the patterns observed in SMAPE and MAPE. The CNN model again exhibits the lowest median RMSE (0.40°C for multivariate, and univariate), indicating

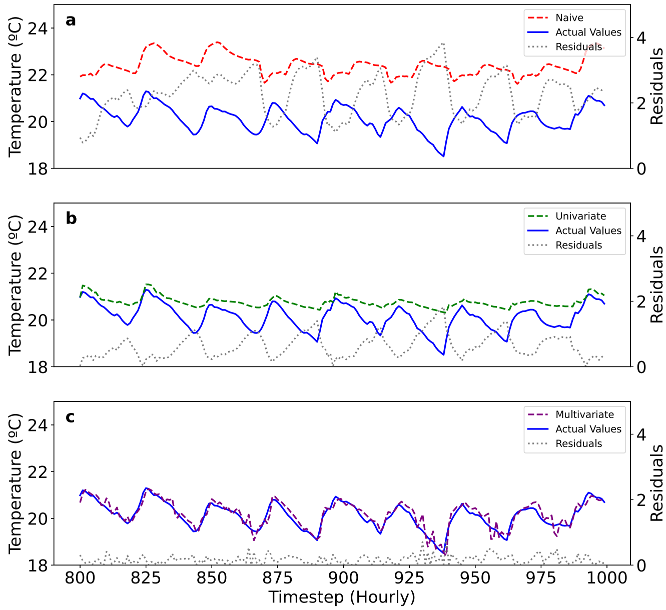

Fig. 8. Comparison of LSTM model temperature predictions against actual values. Each LSTM model uses 12 historical hourly timepoints to predict one hour into the future. Subplots show: (a) Naive model, (b) Univariate LSTM model, and (c) Multivariate LSTM model. Blue lines represent actual temperatures, orange dashed lines show model predictions, and gray areas depict prediction residuals. The x-axis represents hourly timesteps, showcasing approximately 200 hours of predictions.

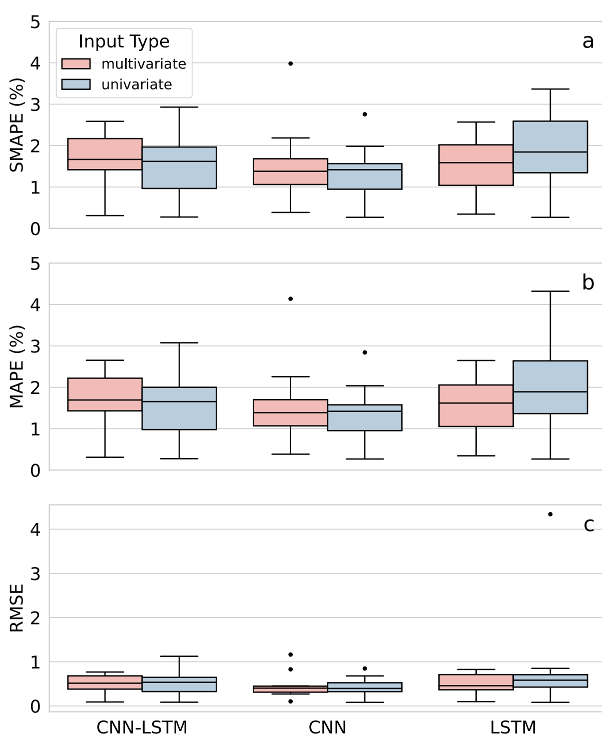

Fig. 9. Performance Comparison of univariate and multivariate Models for temperature prediction. The boxplots show: (a) SMAPE, (b) MAPE, and (c) RMSE for each model type. Blue boxes represent univariate models, while red boxes indicate multivariate models. Each model type (CNN-LSTM, CNN, LSTM) is compared across these three error metrics. Lower values indicate better performance. The boxes show the interquartile range (IQR), with the median marked by a horizontal line. Whiskers extend to 1.5 times the IQR from the edges of the box. Points beyond the whiskers represent outliers.

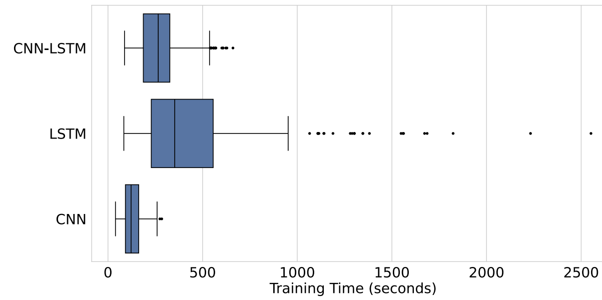

Fig. 10. Training time across model architectures. The boxplots show the distribution of training time for each model type. The boxes show the interquartile range (IQR), with the median marked by a horizontal line. Whiskers extend to 1.5 times the IQR from the edges of the box. Points beyond the whiskers represent outliers.

its superior predictive accuracy. A tighter interquartile range of these models also demonstrates greater reliability across indoor conditions. The LSTM model shows the largest improvement when switching to multivariate input, with the median RMSE decreasing from 0.58ºC to 0.46ºC.

Key findings from this analysis include:

1. Multivariate superiority: Across all metrics and model types, multivariate inputs consistently lead to better performance than univariate inputs. This suggests that incorporating additional relevant variables enhances prediction accuracy. This aligns with previous research on forecasting indoor temperature using multivariate analysis in conditioned indoor environments, which has demonstrated high forecasting accuracy [53].

2. CNN model performance and efficiency: The CNN architecture appears to be the most effective, consistently showing the lowest error rates and highest consistency across all metrics whilst maintaining the lowest average training time (Fig. 10). This indicates that convolutional layers alone perform well for this particular prediction task.

3. Model ranking: Based on these results, the models can be ranked in order of decreasing performance as follows: CNN > CNN-LSTM > LSTM. This ranking holds true for both multivariate and univariate inputs. Similar models architectures were previously used for temperature prediction modeling based on data collected from a room in a university building in Belgium [15]. The results of their study confirmed that the CNN-LSTM model outperformed other models and showed a better robustness against error accumulation. However, it is important to note that their study focused on a single room and primarily utilized HVAC information, whereas our study includes a broader range of environmental metrics.

4. Variance in performance: The LSTM model shows the largest variance in performance across all metrics, as evidenced by the wider boxes and longer whiskers in Fig. 9. This suggests that its predictions are less consistent compared to the CNN-based models.

These results highlight the importance of both model architecture and input complexity in achieving accurate and consistent predictions. The superior performance of the multivariate CNN model suggests that this approach is particularly well-suited for capturing complex spatial relationships in the data, leading to more accurate predictions. The consistent improvement seen with multivariate inputs underscores the value of incorporating diverse, relevant data streams in predictive modeling tasks for indoor temperature forecasting.

3.7. History and prediction length experiments

To further refine our modeling approach, we investigate the impact of historical input length and prediction horizon on model performance, with timesteps on an hourly scale. This exploration informs the optimal configuration for balancing prediction accuracy and computational efficiency, considering the granularity of hourly data intervals. Fig. 11

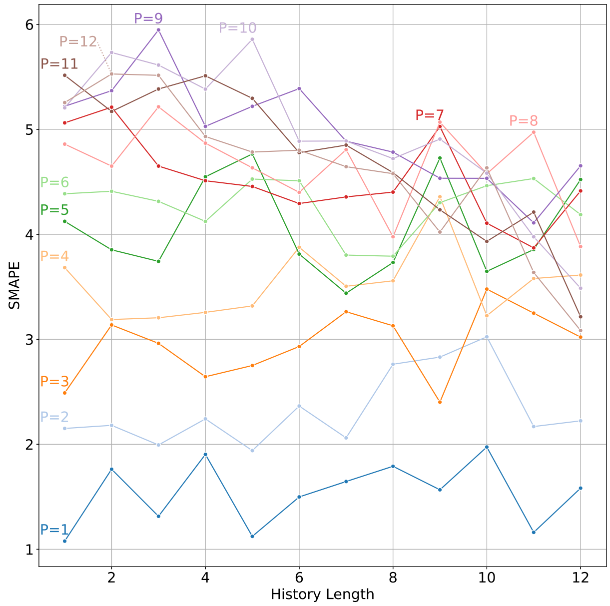

Fig. 11. Performance of CNN multivariate models with varying historical input and prediction lengths. The x-axis represents the historical input length (1 to 12 time steps), while the y-axis shows the SMAPE. Each line represents a different prediction length (P) from 1 to 12 time steps, as indicated by the color legend. Lower SMAPE values indicate better model performance.

illustrates the performance of the CNN multivariate model across different combinations of historical input length and prediction length (P), as measured by SMAPE.

Historical input length The plot reveals a general trend of decreasing SMAPE as the historical input length increases, particularly for longer prediction lengths (P = 7 to P = 12). This suggests that providing more historical context generally improves the model's predictive accuracy, especially for long-term forecasts. For instance, with P = 12, the SMAPE decreases from 5.25% at 1 time step to 3.08% at 12 time steps, a substantial improvement of 41.3%.

However, the impact of historical input length varies across different prediction horizons. For shorter prediction lengths (P = 1 to P = 6), the effect is less pronounced and more erratic. For example, with P = 1, the SMAPE fluctuates between 1.08% and 1.97% across different historical input lengths, showing no clear trend of improvement.

Prediction length (P) Since room temperature has a highly recurrent dynamics, previous values of the variable significantly impact future predictions. Therefore, it was anticipated that performance would decrease with longer prediction horizons. As the prediction length increases, we observe:

1. Short-term predictions (P = 1 to P = 4) show the lowest SMAPE values across all historical input lengths, indicating higher accuracy for near-term forecasts. The mean SMAPE for P = 1 is 1.54%, while for P = 4 it increases to 3.49%.

2. Medium-term predictions (P = 5 to P = 8) exhibit moderate SMAPE values, with performance degrading as P increases. The mean SMAPE for P = 5 is 4.06%, rising to 4.58% for P = 8.

3. Long-term predictions (P = 9 to P = 12) consistently show the highest SMAPE values, reflecting the increasing difficulty of accurate long-range forecasting. The mean SMAPE for P = 9 is 4.97%, reaching 4.61% for P = 12. Interestingly, there's a slight improvement in performance from P = 11 to P = 12, with mean SMAPE decreasing from 4.69% to 4.61%.

Optimal configuration Based on this analysis, the optimal configuration for the CNN multivariate model appears to be:

1. Historical input length: 10 - 12 time steps provides good accuracy for long-term predictions (P = 10-12). For example, with P = 12, the SMAPE decreases from 4.63% at 10 time steps to 3.08% at 12 time steps, a significant improvement of 33.5%. For short term predictions (P = 1-4), there are varying historical input lengths that can be considered.

2. Prediction length: The models perform best for short-term predictions (P = 1 to P = 4), with a gradual decrease in performance as P increases. If longer-term predictions are required, users should expect a notable decrease in accuracy, particularly beyond P = 8. The mean SMAPE increases by 196% from P = 1 (1.54%) to P = 12 (4.61%).

These findings highlight the importance of carefully tuning the historical input length and managing expectations for prediction accuracy as the forecast horizon extends. While longer historical contexts generally improve performance, there are diminishing returns, especially for short-term predictions.

The analysis reveals that the relationship between historical input length, prediction length, and model performance is complex and non-linear. For instance, the best performance for P = 12 is achieved with a historical input length of 12 (SMAPE = 3.08%), while for P = 1, the optimal historical input length is 1 (SMAPE = 1.08%). This suggests that the optimal configuration may need to be tailored to specific prediction horizons.

This detailed examination can guide users in selecting the most appropriate model configuration based on their specific forecasting needs and accuracy requirements, balancing the trade-offs between prediction accuracy, computational resources, and the desired forecast horizon.

3.8. Model interpretability: SHAP analysis

To gain deeper insights into our best-performing model's decision-making process, we conducted a SHAP analysis [41]. For this analysis, we focused on a model configuration with a 12-timestep historical input and a 1-timestep prediction horizon. This specific configuration was chosen to elucidate the temporal impact of various features on a simple, short-term prediction timeframe. By examining how the model weighs information from different historical timesteps for a single-step future prediction, we can better understand the relative importance of recent versus older data across different features. This approach allows us to dissect the model's internal logic and reveal how it integrates information over time to make its predictions, providing valuable insights into the dynamics of indoor environmental forecasting.

The SHAP analysis provides a deeper understanding of feature importance and temporal relevance, highlighting key drivers of indoor temperature dynamics. By quantifying the contribution of each feature at each historical timestep to the final prediction, we can identify which factors are most crucial for accurate short-term forecasting and how their influence varies over the recent past. Similar analyses have previously been used to investigate the impact of borehole field data input parameters on the forecasting accuracy of multivariate hybrid deep learning models for building heating and cooling [3].

Fig. 12 presents the mean absolute SHAP values for all features, providing an aggregate view of feature importance in our CNN multivariate model. The analysis reveals a clear hierarchy of feature importance, with temperature emerging as the most influential factor by a significant margin (SHAP value of 0.021063). This is followed by humidity (0.003982) and HVAC status (0.002960), underscoring the critical role of these parameters in indoor climate prediction. Pressure (0.001750) and VOC levels (0.001317) show moderate importance, while booking status (0.000853) demonstrates notable influence despite being a non-environmental factor. The remaining features, including PM2.5, CO2 levels, people count, and PM10, exhibit relatively low importance with SHAP values ranging from 0.000445 to 0.000211.

Examining the temporal importance of building utilization features (Fig. 13), HVAC status demonstrates high importance, particularly in recent timesteps (T-6 to T-1). Its SHAP value increases from -0.0017 at

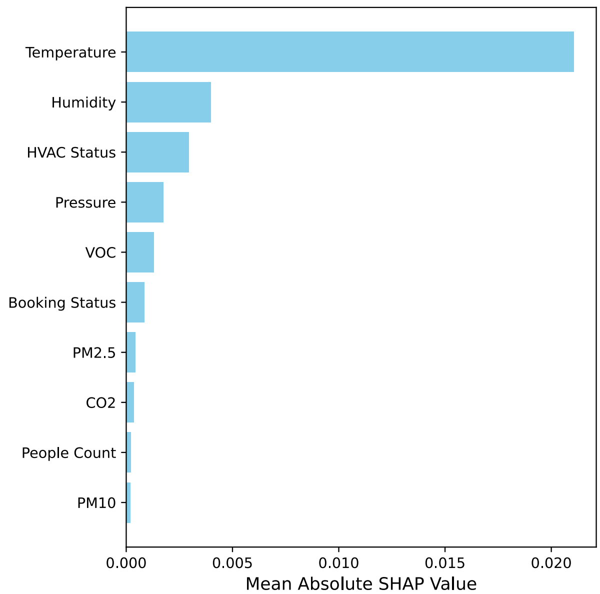

Fig. 12. Mean Absolute SHAP values for all features in the CNN multivariate model. Features are ranked from most to least important based on their impact on model predictions. Higher SHAP values indicate greater feature importance.

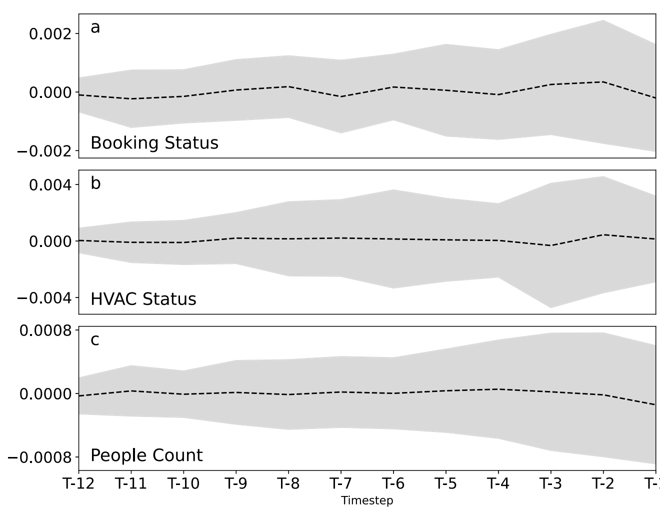

Fig. 13. Temporal SHAP values for building utilization features over 12 historical timesteps (T-12 to T-1). The plot shows the importance of Booking Status, People Count, and HVAC Status at each timestep (mean ± 95% CI), with larger magnitude of values indicating greater influence on model predictions.

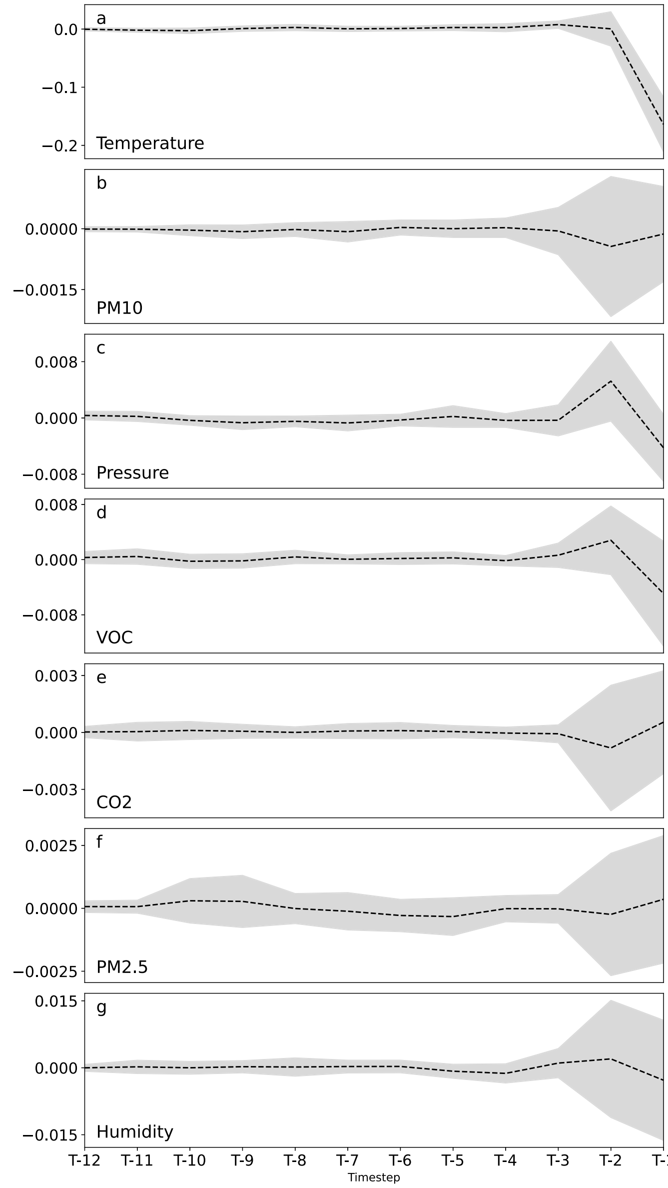

Fig. 14. Temporal SHAP values for micro-climate features over 12 historical timesteps (T-12 to T-1). The plot illustrates the importance of Temperature, Pressure, VOC Levels, Humidity, PM10, PM2.5, and CO2 Levels at each timestep (mean ± 95% CI). Larger magnitude of values indicating greater influence on model predictions.

T-12 to -0.0037 at T-6, indicating that recent HVAC operations have a significant and immediate impact on predictions. Booking status shows relatively consistent importance across all timesteps, with values ranging from -0.0007 to -0.0020, suggesting a steady influence of room reservation patterns on predictions. People count exhibits the lowest importance among these features, with values between -0.0003 and -0.0009, showing minimal variation across timesteps.

The temporal importance of micro-climate features (Fig. 14) further emphasizes the critical role of temperature. Its SHAP values increase dramatically from -0.0031 at T-3 to -0.2110 at T-1, indicating its paramount importance in recent timesteps. Humidity demonstrates moderate importance, with values ranging from -0.0011 to -0.0162, showing increased relevance in recent timesteps (T-3 to T-1). Pressure exhibits moderate importance with values between -0.0004 and -0.0090, reflecting relatively consistent influence across all timesteps. VOC levels show low to moderate importance, with values from -0.0006 to -0.0125, while PM2.5, PM10, and CO2 levels all demonstrate relatively low importance with absolute SHAP values generally below 0.005.

This SHAP analysis offers valuable insights into the CNN multivariate model's decision-making process. The dominance of temperature as a predictor, particularly in recent timesteps, suggests that recent temperature data is crucial for accurate forecasting. The model's emphasis on recent information (especially T-6 to T-1) for most features indicates that immediate past conditions are more relevant for predictions. The high importance of HVAC status, especially in recent timesteps, highlights the direct impact of climate control systems on indoor conditions.

While some features like PM10 and people count show low overall importance, their inclusion still contributes to the model's multivariate approach, potentially capturing subtle interactions or edge cases. The varying temporal importance of features suggests that the CNN architecture effectively captures both short-term and longer-term dependencies in the data.

These findings can guide feature selection, data collection priorities, and potential areas for model refinement in future iterations. The clear importance of recent temperature and HVAC status data suggests that real-time monitoring and responsive control systems could significantly improve prediction accuracy and, consequently, indoor climate management.

4. Discussion

The results of our comprehensive analysis of indoor environmental prediction models provide valuable insights into the performance, optimization, and interpretability of various machine learning approaches within building automation systems. Due to the safety-critical nature of HVAC systems in buildings, conducting experiments on the actual systems is not feasible because of the associated risks. Therefore, predictive models serve as digital twins of the buildings, enabling experimental implementations and performance validation. This discussion section synthesizes our findings and explores their implications for future research and practical applications in smart building management.

4.1. Model performance and architecture comparison

Our analysis revealed that the CNN model consistently outperformed both the LSTM and hybrid CNN-LSTM models for temperature prediction, as evidenced by lower SMAPE, MAPE, and RMSE values. A previous study investigated Multi-Layer Perceptron (MLP), LSTM, and CNN-LSTM models, evaluating and comparing them across 1, 30, 60, and 120-minute horizons using a closed-loop prediction scheme [15]. The results demonstrated that the CNN-LSTM outperformed all other models across all prediction horizons and exhibited better robustness against error accumulation. However, they did not investigate the CNN model. Additionally, that study only analyzed data from one room over 19 days, whereas our study collected and analyzed data from multiple rooms over a longer period. Furthermore, unlike that study, our variables were independent of the HVAC system, with the HVAC schedule being only one of the influencing factors.

The superior performance of the CNN model suggests that the spatial dependencies and local patterns captured by convolutional layers are particularly effective in modeling indoor temperature dynamics. This aligns with recent studies, such as Zhao et al. [54], which explored the optimal control of heat network in residential district using hybrid model based on CNN. The relatively poorer performance of the LSTM and hybrid CNN-LSTM models in this context is intriguing. It may indicate that for temperature prediction in indoor environments, the long-term dependencies captured by LSTM layers are less critical than the local patterns extracted by convolutional operations. This could be due to the strong influence of recent conditions and control actions (e.g., HVAC operations) on indoor temperature.

4.2. Optimization of model parameters

Our investigation into the effects of historical input length and prediction length on model performance revealed several key insights. Increasing historical input length generally improved model performance, particularly for long-term predictions (7-12 timesteps ahead). However, the benefits of longer historical contexts diminished beyond 8-10 timesteps, suggesting an optimal range for balancing performance and computational efficiency. Prediction accuracy consistently decreased as the prediction length increased, with a notable performance drop for predictions beyond 8 timesteps ahead. This is aligned with all previous studies [31,6,49,15], which a significant drop in accuracy was observed as the prediction horizon widened. This behavior is expected since models use their own predictions to forecast further into the future, causing prediction errors to accumulate over time.

These findings highlight the challenge of long-term predictions in dynamic indoor environments and suggest that frequent model updates with recent data may be more effective than attempting to forecast far into the future. This aligns with previous study [25] on adaptive HVAC control systems, which emphasized the importance of short-term predictions for real-time optimization.

4.3. Feature importance and temporal dynamics

The SHAP analysis provided crucial insights into the relative importance of different features and their temporal dynamics. Temperature emerged as the most influential feature, followed by humidity and HVAC status. This hierarchy of importance aligns with fundamental principles of indoor climate control and human comfort, as discussed in Fanger's seminal work on thermal comfort [17]. It also aligns with a review study on measuring factors impacting thermal comfort and indoor air quality [28].

The high importance of recent timesteps (particularly T-6 to T-1) for most features underscores the rapid dynamics of indoor environments and the need for real-time or near-real-time monitoring and prediction systems. This also highlights the importance of CNN models and their capabilities in short-horizon predictions.

The varying temporal importance patterns across features (e.g., consistent importance for booking status vs. increasing importance for temperature in recent timesteps) reflect the complex interactions between building usage, environmental factors, and climate control systems. The relatively low importance of some features (e.g., PM10, people count) does not necessarily negate their value. In specific contexts or edge cases, these features may play crucial roles. This is the first time these factors have been included in a building temperature prediction model. Therefore, future work could explore the significance of these features in greater detail under various conditions or in specific types of buildings.

4.4. Implications for smart building management

Our findings have several important implications for the development and implementation of smart building management systems. The superior performance of multivariate models suggests that comprehensive sensor networks capturing a wide range of environmental parameters are crucial for accurate predictions and efficient building management. The effectiveness of the CNN model indicates that investment in this neural network architecture can yield significant improvements in prediction accuracy.

The optimal historical input length of 8-10 timesteps provides guidance for data storage and processing requirements in real-time prediction systems. The high importance of temperature, humidity, and HVAC status underscores the need for high-quality, reliable sensors for these parameters and suggests that they should be prioritized in data collection and model development efforts.

The proposed CNN architecture could potentially make a significant contribution to Model Predictive Control (MPC) HVAC systems [1]. These models enable controllers to explore the consequences of their actions without interacting with the real environment, allowing them to optimize their strategies to achieve desired control objectives. These insights can guide the design of more effective and efficient smart building systems, potentially leading to improved energy efficiency, occupant comfort, and overall building performance.

4.5. Limitations and future research directions

While our study provides valuable insights, it also has limitations that point to future research directions. Our analysis focused on a specific set of buildings and environmental conditions, using only within-room sensor data. Various modeling factors, such as the employed architecture, software framework, and data processing, could influence prediction capability. Therefore, results from the methods developed in this paper should be considered as the baseline for comparison and not a generalized solution.

Future work should explore the generalizability of these findings across different building types, climates, and usage patterns. The relatively low importance of some features (e.g., PM levels, CO2) in our models may not hold true for all environments. Studies in more polluted areas or densely occupied spaces might yield different results.

Binary representation of HVAC status has proved effective for the current study, yet we acknowledge its limitations in capturing the full complexity of HVAC operations. Future work could explore more nuanced representations, potentially incorporating variable setpoints or multiple operational states, should such data become available and demonstrate significant improvements in model performance.

A key objective for subsequent studies would be to expand the spatial and temporal scope of our models. Specifically, we aim to investigate the spatio-temporal impacts of these deep learning systems by incorporating information across multiple rooms and, crucially, including outdoor measurements. This expanded dataset would allow us to better understand and model temperature transfer across different areas of the built environment. By integrating outdoor climate data such as temperature, humidity, and solar radiation, we could potentially tease apart the complex interactions between internal and external factors influencing indoor thermal conditions.

This paper was primarily focused on point predictions. Exploring probabilistic forecasting methods could provide valuable uncertainty estimates for decision-making in building management. The integration of external factors (e.g., weather forecasts, energy prices) into the prediction models could potentially improve long-term forecasting capabilities.

Future research should address these limitations and explore emerging technologies, such as federated learning for privacy-preserving multi-building models or the integration of reinforcement learning for adaptive building control strategies. Additionally, investigating the potential of transfer learning techniques could enhance model adaptability across different building types and environmental conditions.

Such an expanded research direction aligns with our overarching goal of developing more comprehensive and robust predictive models for indoor environmental management, ultimately contributing to improved energy efficiency and occupant comfort in diverse building settings.

5. Conclusion

In this study, we compared the performance of CNN, LSTM, and a CNN-LSTM hybrid model architecture for indoor temperature modeling. The dataset, collected from multiple rooms on one level of a university building in Sydney's CBD, includes various environmental and building management measurements.

The results showed multivariate inputs consistently lead to better performance than univariate inputs. The CNN architecture appears to be the most effective, consistently showing the lowest error rates and highest consistency across all metrics. Additionally, the models perform best for short-term predictions ($P = 1$ to $P = 4$). The SHAP analysis provides a deeper understanding of feature importance and temporal relevance, highlighting temperature, humidity and HVAC status as the key drivers of indoor temperature dynamics.

Our comprehensive analysis of indoor environmental prediction models provides a solid foundation for the development of more accurate, efficient, and interpretable smart building management systems. By leveraging multivariate data, advanced neural network architectures, and insights from feature importance analysis, HVAC systems can better optimize energy usage, enhance occupant comfort, and contribute to more sustainable built environments. The superiority of CNN models in this context opens new avenues for research in time series forecasting for indoor environmental prediction, challenging researchers to reconsider traditional approaches and explore innovative architectures tailored to the unique characteristics of smart building data.

CRediT authorship contribution statement

Roupen Minassian: Writing – review & editing, Writing – original draft, Visualization, Software, Investigation, Formal analysis. **Adriana-Simona Mihăiţă:** Writing – review & editing, Writing – original draft, Supervision, Project administration, Methodology, Investigation. **Arezoo Shirazi:** Writing – review & editing, Supervision, Project administration, Investigation, Conceptualization.

Code availability

The complete source code for training and evaluating the models presented in this study is available at https://github.com/Future-Mobility-Lab/DSI-DAB-smart-building-env-prediction.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

References

[1] A. Afram, F. Janabi-Sharifi, Theory and applications of hvac control systems–a review of model predictive control (mpc), Build. Environ. 72 (2014) 343–355, https://doi.org/10.1016/j.buildenv.2013.11.016.

[2] Z. Afroz, G.M. Shafiullah, T. Urmee, M.A. Shoeb, G. Higgins, Predictive modelling and optimization of hvac systems using neural network and particle swarm optimization algorithm, Build. Environ. 209 (2022) 108681, https://doi.org/10.1016/j.buildenv.2021.108681.

[3] N. Ahmed, M. Assadi, Q. Zhang, Investigating the impact of borehole field data's input parameters on the forecasting accuracy of multivariate hybrid deep learning models for heating and cooling, Energy Build. 301 (2023) 113706, https://doi.org/10.1016/j.enbuild.2023.113706.

[4] P. Aparicio-Ruiz, E. Barbadilla-Martín, J. Guadix, P. Cortés, Knn and adaptive comfort applied in decision making for hvac systems, Ann. Oper. Res. 303 (2021) 217–231, https://doi.org/10.1007/s10479-019-03489-4.

[5] ASHRAE, ASHRAE Handbook: Fundamentals, American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Inc., 2017.

[6] N. Attoue, I. Shahrour, R. Younes, Smart building: use of the artificial neural network approach for indoor temperature forecasting, Energies 11 (2018), https://doi.org/10.3390/en11020395.

[7] A. Ben Said, A. Erradi, A. Ghari Neiat, A. Bouguettaya, A deep learning spatiotemporal prediction framework for mobile crowdsourced services, Mob. Netw. Appl. 24 (2018) 1120–1133, https://doi.org/10.1007/s11036-018-1105-0.

[8] K. Benidis, S.S. Rangapuram, V. Flunkert, Y. Wang, D. Maddix, C. Turkmen, J. Gasthaus, M. Bohike-Schneider, D. Salinas, L. Stella, F.X. Aubet, L. Callot, T. Januschowski, Deep learning for time series forecasting: tutorial and literature survey, ACM Comput. Surv. 55 (2022) 1–36, https://doi.org/10.1145/3533382.

[9] L. Biewald, Experiment tracking with weights and biases, https://www.wandb.com/, 2020, software available from wandb.com.

[10] S. Brandi, M.S. Piscitelli, M. Martellacci, A. Capozzoli, Deep reinforcement learning to optimise indoor temperature control and heating energy consumption in buildings, Energy Build. 224 (2020) 110225, https://doi.org/10.1016/j.enbuild.2020.110225.

[11] M.V. Cano, L. Dufour, A. Skarmeta, A.J. Jara, D. Genoud, B. Ladevie, J.J. Bézian, Big data: the key to energy efficiency in smart buildings, Soft Comput. 20 (2015), https://doi.org/10.1007/s00500-015-1679-4.

[12] S. Dimitroulopoulou, M.R. Dudzińska, L. Gunnarsen, L. Hägerhed, H. Maula, R. Singh, O. Toyinbo, U. Haverinen-Shaughnessy, Indoor air quality guidelines from across the world: an appraisal considering energy saving, health, productivity, and comfort, Environ. Int. 178 (2023) 108127, https://doi.org/10.1016/j.envint.2023.108127.

[13] N. Djongyang, R. Tchinda, D. Njomo, Thermal comfort: a review paper, Renew. Sustain. Energy Rev. 14 (2010) 2626–2640, https://doi.org/10.1016/j.rser.2010.07.040.

[14] B. Dong, B. Andrews, K.P. Lam, M. Höynck, R. Zhang, Y.S. Chiou, D. Benitez, An information technology enabled sustainability test-bed (itest) for occupancy detection through an environmental sensing network, Energy Build. 42 (2010) 1038–1046, https://doi.org/10.1016/j.enbuild.2010.01.016.

[15] F. Elmaz, R. Eyckerman, W. Casteels, S. Latré, P. Hellinckx, Cnn-lstm architecture for predictive indoor temperature modeling, Build. Environ. 206 (2021) 108327, https://doi.org/10.1016/j.buildenv.2021.108327.

[16] C. Faloutsos, J. Gasthaus, T. Januschowski, Y. Wang, Classical and contemporary approaches to big time series forecasting, in: Proceedings of the 2019 International Conference on Management of Data, Association for Computing Machinery, New York, NY, USA, 2019, pp. 2042–2047.

[17] P.O. Fanger, Thermal Comfort: Analysis and Applications in Environmental Engineering, Danish Technical Press, 1970.

[18] J.C.B. Gamboa, Deep learning for time-series analysis, arXiv (Cornell University), https://arxiv.org/abs/1701.01887, 2017.

[19] G.A. Ganesh, S.L. Sinha, T.N. Verma, S.K. Dewangan, Investigation of indoor environment quality and factors affecting human comfort: a critical review, Build. Environ. 204 (2021) 108146, https://doi.org/10.1016/j.buildenv.2021.108146.

[20] R. Garg, S. Barpanda, G.R.S. N S, S. Ramya, Machine learning algorithms for time series analysis and forecasting, in: Proceedings of the 6th International Conference on Microelectronics, Computing and Communication Systems (MCCS-2021), 2021, pp. 1–9, https://doi.org/10.48550/arXiv.2211.14387, in press.

[21] M. González-Torres, L. Pérez-Lombard Martín de Oliva, J. Francisco Coronel Toro, I. Rodríguez Maestre, D. Yan, A review on buildings energy information: trends, end-uses, fuels and drivers, Energy Rep. (2022), https://doi.org/10.1016/j.egyr.2021.11.280.

[22] A. Graves, J. Schmidhuber, Framewise phoneme classification with bidirectional lstm and other neural network architectures, in: iJCNN 2005, Neural Netw. 18 (2005) 602–610, https://doi.org/10.1016/j.neunet.2005.06.042.

[23] P. Hameed Shaikh, N. Bin Mohd Nor, P. Nallagownden, I. Elamvazuthi, T. Ibrahim, A review on optimized control systems for building energy and comfort management of smart sustainable buildings, Renew. Sustain. Energy Rev. 34 (2014) 409–429, https://doi.org/10.1016/j.rser.2014.03.027.

[24] Y. LeCun, B. Boser, J.S. Denker, D. Henderson, R. Howard, W. Hubbard, L.D. Jackel, Handwritten digit recognition with a back-propagation network, in: D. Touretzky (Ed.), Advances in Neural Information Processing Systems, Morgan-Kaufmann, 1989, https://proceedings.neurips.cc/paper_files/paper/1989/file/53c3bce66e43be4f209556518c2fcb54-Paper.pdf.

[25] W. Li, C. Koo, T. Hong, J. Oh, S.H. Cha, S. Wang, A novel operation approach for the energy efficiency improvement of the hvac system in office spaces through real-time big data analytics, Renew. Sustain. Energy Rev. 127 (2020) 109885, https://doi.org/10.1016/j.rser.2020.109885.

[26] B. Lim, S. Zohren, Time-series forecasting with deep learning: a survey, Philos. Trans. - Royal Soc., Math. Phys. Eng. Sci. 379 (2021) 20200209, https://doi.org/10.1098/rsta.2020.0209.

[27] S.M. Lundberg, S.I. Lee, A unified approach to interpreting model predictions, Adv. Neural Inf. Process. Syst. 30 (2017).

[28] N. Ma, D. Aviv, H. Guo, W.W. Braham, Measuring the right factors: a review of variables and models for thermal comfort and indoor air quality, Renew. Sustain. Energy Rev. 135 (2021) 110436, https://doi.org/10.1016/j.rser.2020.110436.

[29] L. Mba, P. Meukam, A. Kemajou, Application of artificial neural network for predicting hourly indoor air temperature and relative humidity in modern building in humid region, Energy Build. 121 (2016) 32–42, https://doi.org/10.1016/j.enbuild.2016.03.046.

[30] L. Morawska, T. Salthammer (Eds.), Indoor Environment: Airborne Particles and Settled Dust, Wiley-VCH Verlag GmbH & Co. KGaA, Weinheim, Germany, 2003, https://eprints.qut.edu.au/238625/.

[31] F. Mtibaa, K.K. Nguyen, M. Azam, A. Papachristou, J.S. Venne, M. Cheriet, Lstm-based indoor air temperature prediction framework for hvac systems in smart buildings, Neural Comput. Appl. 32 (2020) 17569–17585, https://doi.org/10.1007/s00521-020-04926-3.

[32] K. Mui, L. Wong, Cooling load calculations in subtropical climate, Build. Environ. 42 (2007) 2498–2504, https://doi.org/10.1016/j.buildenv.2006.07.006.

[33] W.W. Nazaroff, Indoor particle dynamics, Indoor Air 14 (2004) 175–183, https://doi.org/10.1111/j.1600-0668.2004.00286.x.

[34] Y. Peng, A. Rysanek, Z. Nagy, A. Schlüter, Using machine learning techniques for occupancy-prediction-based cooling control in office buildings, Appl. Energy 211 (2018) 1343–1358, https://doi.org/10.1016/j.apenergy.2017.12.002.

[35] A. Persily, L. de Jonge, Carbon dioxide generation rates for building occupants, Indoor Air 27 (2017) 868–879, https://doi.org/10.1111/ina.12383.

[36] J. Reynolds, Y. Rezgui, A. Kwan, S. Piriou, A zone-level, building energy optimisation combining an artificial neural network, a genetic algorithm, and model predictive control, Energy 151 (2018) 729–739, https://doi.org/10.1016/j.energy.2018.03.113.

[37] J. Runge, S. Bathiany, E. Bollt, G. Camps-Valls, D. Coumou, E. Deyle, C. Glymour, M. Kretschmer, M.D. Mahecha, J. Muñoz-Marí, E.H. van Nes, J. Peters, R. Quax, M. Reichstein, M. Scheffer, B. Schölkopf, P. Sprites, G. Sugihara, S. Jie, K. Zhang, J. Zscheischler, Inferring causation from time series in Earth system sciences, Nat. Commun. 10 (2019) 1–13, https://doi.org/10.1038/s41467-019-10105-3.

[38] J. Runge, R. Zmeureanu, A review of deep learning techniques for forecasting energy use in buildings, Energies 14 (2021), https://doi.org/10.3390/en14030608.

[39] O. Seppanen, W.J. Fisk, M.J. Mendell, Association of ventilation rates and co2 concentrations with health and other responses in commercial and institutional buildings, Indoor Air 9 (1999) 226–252, https://doi.org/10.1111/j.1600-0668.1999.00003.x.

[40] S. Seyedzadeh, F.P. Rahimian, I. Glesk, M. Roper, Machine learning for estimation of building energy consumption and performance: a review, Vis. Eng. 6 (2018) 1–20, https://doi.org/10.1186/s40327-018-0064-7.

[41] L.S. Shapley, A value for n-person games, in: Contributions to the Theory of Games, vol. 2, Princeton University Press, 1953, pp. 307–317.

[42] J. Spengler, J. McCarthy, J. Samet, Indoor Air Quality Handbook, McGraw-Hill Education, 2001.

[43] A. Tzempelikos, A.K. Athienitis, The impact of shading design and control on building cooling and lighting demand, Sol. Energy 81 (2007) 369–382, https://doi.org/10.1016/j.solener.2006.06.015.

[44] J. Vesterberg, S. Andersson, Achieved energy and climate goals in project Ålidhem: an evaluation of a refurbishment of 21 Swedish multifamily buildings, in: 11th Nordic Symposium on Building Physics, NSB2017, 11-14 June 2017, Trondheim, Norway, Energy Proc. 132 (2017) 51–56, https://doi.org/10.1016/j.egypro.2017.09.630.

[45] S. Wang, J. Cao, P.S. Yu, Deep learning for spatio-temporal data mining: a survey, arXiv (Cornell University), https://arxiv.org/abs/1906.04928, 2019.

[46] Y. Wei, X. Zhang, Y. Shi, L. Xia, S. Pan, J. Wu, M. Han, X. Zhao, A review of data-driven approaches for prediction and classification of building energy consumption, Renew. Sustain. Energy Rev. 82 (2018) 1027–1047, https://doi.org/10.1016/j.rser.2017.09.108.

[47] C.J. Weschler, W.W. Nazaroff, Semivolatile organic compounds in indoor environments, Atmos. Environ. 42 (2008) 9018–9040, https://doi.org/10.1016/j.atmosenv.2008.09.052.

[48] WHO, WHO Guidelines for Indoor Air Quality: Selected Pollutants, World Health Organization, 2010.

[49] C. Xu, H. Chen, J. Wang, Y. Guo, Y. Yuan, Improving prediction performance for indoor temperature in public buildings based on a novel deep learning method, Build. Environ. 148 (2019) 128–135, https://doi.org/10.1016/j.buildenv.2018.10.062.

[50] D. Yang, C.M. Mak, Relationships between indoor environmental quality and environmental factors in university classrooms, Build. Environ. 186 (2020) 107331, https://doi.org/10.1016/j.buildenv.2020.107331.

[51] S. Yang, M.P. Wan, W. Chen, B.F. Ng, S. Dubey, Model predictive control with adaptive machine-learning-based model for building energy efficiency and comfort optimization, Appl. Energy 271 (2020) 115147, https://doi.org/10.1016/j.apenergy.2020.115147.

[52] F. Yuan, R. Yao, W. Yu, S. Sadrizadeh, H. Awbi, P. Kumar, Dynamic characteristics of particulate matter resuspension due to human activities in indoor environments: a comprehensive review, J. Build. Eng. 79 (2023) 107914, https://doi.org/10.1016/j.jobe.2023.107914.

[53] F. Zamora-Martínez, P. Romeu Guallart, P. Botella-Rocamora, J. Pardo Albiach, Towards energy efficiency: forecasting indoor temperature via multivariate analysis, Energies 6 (2013) 4639–4659, https://doi.org/10.3390/en6094639.

[54] A. Zhao, L. Mi, X. Xue, J. Xi, Y. Jiao, Heating load prediction of residential district using hybrid model based on cnn, Energy Build. 266 (2022) 112122, https://doi.org/10.1016/j.enbuild.2022.112122.

[55] Z. Zheng, Y. Yang, J. Liu, H.N. Dai, Y. Zhang, Deep and embedded learning approach for traffic flow prediction in urban informatics, IEEE Trans. Intell. Transp. Syst. 20 (2019) 3927–3939, https://doi.org/10.1109/TITS.2019.2909904.