

ORIGINAL RESEARCH ARTICLE

Machine learning-based model predictive control for multizone building automation: A case study

Pradeep Shakya¹ , Shiva Sreenivasan² , Baskaran Krishnamoorthy¹ , Shiyu Yang³ , and Man Pun Wan^{2*} 

¹Sustainable Built Environments, Energy Research Institute @ Nanyang Technological University (ERI@N), 1 Cleantech Loop, Singapore

²Thermal and Fluids Division, School of Mechanical and Aerospace Engineering (MAE), Nanyang Technological University (NTU), 50 Nanyang Avenue, Singapore

³Department of Fluid Dynamics, Institute of High Performance Computing (IHPC), Agency for Science, Technology, and Research (A*STAR), 1 Fusionopolis Way, Singapore

Abstract

In Singapore's hot and humid climate, air-conditioning and mechanical ventilation (ACMV) systems account for over 60% of commercial building energy consumption, driving efforts to enhance energy efficiency through predictive control strategies such as model predictive control (MPC) to overcome the limitations of conventional reactive building automation systems. This paper presents a multizone MPC system designed to optimize energy consumption and thermal comfort in a commercial building's ACMV system in Singapore. The system was implemented in a multi-use test building with real occupancy and a deployment area of approximately 850 m², partitioned into six learning zones, two office spaces, and three open spaces. The ACMV system serving the deployment area consisted of two primary air-handling units and 16 fan coil units, where chilled water was supplied to the cooling coils, and conditioned air was distributed through motorized diffusers. To facilitate predictive control, data-driven thermal prediction models were developed for each zone using a non-linear autoregressive exogenous network with exogenous inputs trained on historical data and disturbances. Thermal comfort optimization was guided by the predictive mean vote, which was targeted at 0, representing thermal neutrality (as per ASHRAE 55 standards), and constrained within a range of -0.5 - 0.5. Performance comparisons demonstrated that the MPC system achieved over 42% energy savings compared to the original thermostat-based control while enhancing thermal comfort. Despite its advantageous control performances, challenges for large-scale deployment remain, including implementation costs, scalability, and model accuracy. Future work can address these challenges by developing comfort models that leverage existing building sensors.

Keywords: Model predictive control; Coordinated multisystem control; Energy saving; Thermal comfort; Visual comfort; High-performance building

*Corresponding author:

Man Pun Wan
 (mpwan@ntu.edu.sg)

Citation: Shakya P, Sreenivasan S, Krishnamoorthy B, Yang S, Wan MP. Machine learning-based model predictive control for multizone building automation: A case study. *Int J AI Mater Design*. 2025;2(1):39-53.
 doi: 10.36922/ijamd.8161

Received: December 24, 2024

Revised: February 7, 2025

Accepted: February 20, 2025

Published online: March 5, 2025

Copyright: © 2025 Author(s). This is an Open-Access article distributed under the terms of the Creative Commons Attribution License, permitting distribution, and reproduction in any medium, provided the original work is properly cited.

Publisher's Note: AccScience Publishing remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

1. Introduction

In Singapore, owing to its hot and humid climate throughout the year, commercial buildings typically spend more than 60% of their electricity consumption on air conditioning, mechanical, and ventilation (ACMV) systems.¹ On March 31, 2020, Singapore submitted its Long-Term Low Emissions Development Strategy to the United Nations Framework Convention on Climate Change, pledging to halve its emissions from their peak to 33 MTCO₂e and to achieve net-zero emissions by 2050.² At present, the Building Construction Authority (BCA) has reported that high-performance buildings in Singapore have achieved more than 65% improvement in energy efficiency compared to 2005 levels. The BCA further aims to raise the energy efficiency improvement of Singapore's buildings to 80% relative to 2005 levels by 2030 through its Green Building Innovation Cluster program.³

The conventional approach in current building automation and control systems is predominantly "reactive," generating control signals based on deviations of previously measured information from a control setpoint (thermostat-based). Consequently, due to the different thermal characteristics of buildings and the non-linear operation of their ACMV systems, reactive control strategies are unable to achieve optimal efficiency and occupant comfort. These strategies only respond to transient disturbances – changes in occupancy, internal heat load profiles, and external weather conditions – as well as to building dynamics after the disturbances occur. This limitation highlights the necessity for predictive control strategies that can anticipate and mitigate such disturbances before they adversely impact building performance.

Studies on predictive control solutions for optimizing ACMV operation in recent decades can broadly be categorized based on model architecture into (i) model predictive control (MPC) and (ii) reinforced learning (RL).⁴ MPC uses prediction models to forecast future system behavior and proactively adjust control actions, making it particularly well-suited for complex systems with non-linear dynamics. Within the MPC framework, three primary techniques are used to develop prediction models: physics-based (white box), data-driven (black box), and hybrid (gray box) methods.⁵ Physics-based models rely on fundamental principles such as mass and energy conservation and can be implemented using software tools such as EnergyPlus,⁶ TRNSYS,⁷ eQuest,⁸ and Modelica.⁹ However, these models require significant expertise, resulting in higher implementation costs, particularly in complex real-world buildings. In contrast, data-driven models apply regression techniques to large

datasets from real buildings or simulations, necessitating extensive data collection. Hybrid models combine elements of both approaches, estimating certain physical parameters through data-driven methods.

Data-driven MPC strategies have demonstrated considerable promise in achieving both optimal thermal comfort and high energy efficiency. Over the years, researchers such as Yang *et al.*,¹⁰ Yang and Wan¹¹ have enhanced MPC by simplifying building models through linearization techniques. Široký *et al.*¹² reported energy savings of 15 – 30% in a university building by integrating MPC with weather forecasting capabilities, whereas Ma *et al.*¹³ achieved a 19% efficiency improvement using a hierarchical MPC system for a chiller plant. In experimental studies, Pang *et al.*¹⁴ demonstrated that MPC reduced chilled water consumption by 42% in a radiant slab system compared to heuristic methods. Despite these promising results, two primary obstacles hinder the widespread adoption of MPC: the need for highly accurate predictive models and the significant computational cost associated with solving complex optimization problems across multiple zones. To address these issues, Yang *et al.*¹⁵ proposed a two-level distributed computation scheme that optimizes thermal comfort and internal air quality across multiple zones using MPC. In addition, an event-triggered optimization mechanism has been shown to reduce the computational load by 12–22% while achieving over 9% energy savings,¹⁶ further highlighting MPC's potential in real-world applications.

Alongside MPC, recent advancements in machine learning (ML) have introduced powerful tools for predicting building energy performance. A review by Fathi *et al.*¹⁷ examined 14 ML methods relevant to this domain. Among these, support vector machines (SVM) are particularly effective for regression tasks involving non-linear relationships, especially when using the radial basis function kernel, which captures non-linear interactions more effectively compared to sigmoid and polynomial kernels.^{18,19} Random forest (RF), an ensemble learning method based on decision trees, is adept at managing non-linear interactions and offers improved prediction stability, albeit at the expense of increased computational time.²⁰ In the realm of artificial neural networks (ANN), the non-linear autoregressive network with exogenous inputs (NARX) is well-suited for dynamic systems, leveraging both historical and external inputs to forecast future states.²¹ In addition, long short-term memory (LSTM) networks, a type of recurrent neural network known for capturing long-term dependencies, are particularly effective for modeling the time-series dynamics of building energy consumption.²² Integrating these ML-based forecasting methods with MPC can further enhance the

predictive capabilities and overall performance of building automation systems.

Despite the promise of MPC, challenges remain in accurately modeling the complex thermal dynamics of buildings, which are influenced by varying weather conditions, occupancy patterns, and building envelope characteristics. In addition, the computational burden associated with optimization across multiple zones presents further difficulties.²³ To address these challenges, model-free approaches based on RL (such as Q-learning²⁴ and deep reinforcement learning²⁵) have been explored, with several studies demonstrating notable energy savings. For instance, Ding *et al.*²⁶ reported approximately 14% energy savings using deep RL with inputs from EnergyPlus simulations. However, the present study focuses on MPC due to its structured approach and its ability to directly integrate system constraints into the optimization process.

In this paper, we present a multizone MPC framework for the coordinated control of ACMV systems, aiming to optimize both occupant comfort and energy efficiency. To evaluate performance, we compare the energy consumption and thermal comfort achieved using our coordinated MPC approach with those obtained from a conventional reactive control strategy employing proportional-integral-derivative (PID) control. The key novelties of our MPC system are as follows:

- (i) Large-scale, multizone implementation: This study demonstrates a real-world deployment of MPC in an 11-zone building (850 m²) with actual occupancy. This contrasts with previous studies, which were limited to smaller-scale implementations (e.g., two zones²⁷), simulated environments,²⁸ or unoccupied buildings²⁹
- (ii) Integration of ML-based weather forecasting: By incorporating ML-based weather forecasting into the MPC framework, our system proactively adjusts control actions in anticipation of weather changes. This represents an advancement over previous studies, which typically relied on existing or historical weather data.³⁰⁻³²

This work provides valuable insights into both the challenges and benefits of implementing MPC in complex, real-world settings, particularly within tropical environments. Consequently, it contributes significantly to the advancement of scalable solutions for building energy management.

This paper is organized as follows:

- Section 1 (Introduction): Provides background, motivation, and objectives of the study
- Section 2 (Testbed setup): Describes the heating, ventilation, and air conditioning (HVAC) system and sensor configuration of the testbed

- Section 3 (MPC controller for ACMV): Details the design and implementation of the MPC controller for the ACMV system
- Section 4 (ML model development for MPC): Outlines the development of the ML models used for forecasting and their integration within the MPC framework
- Section 5 (Results and discussion): Presents the performance evaluation, key findings, and a comparative analysis with conventional control strategies and discusses the challenges encountered
- Section 6 (Conclusion): Summarizes the study's contributions, implications, and suggestions for future research.

2. Testbed setup

For the testbed, a commercial building located in Jurong East, Singapore, was selected. The building is served by a central chiller plant consisting of two chillers, each with a capacity of 438 refrigeration tons. The ACMV system in the building operates from 7 AM to 7 PM on weekdays. The experiments were conducted in an 850 m² multiuse test space, partitioned into 11 zones, including two office spaces (OS1 – 2), six learning zones (LZ1 – 6), and three common areas (CA1 – 3), as depicted in [Figure 1](#). MPC was implemented to coordinate and control the ACMV systems, as illustrated in [Figure 2](#). The test space includes operable partition walls between LZ1 and LZ2, as well as among LZ4, LZ5, and LZ6. These partitioning walls can be adjusted so that two zones (LZ1 and LZ2) or up to three zones (LZ4, LZ5, and LZ6) can be merged into a single zone. The open/closed state of these walls is determined by proximity sensors, which output “1” if the walls are closed (i.e., zones remain separate, as in LZ1 and LZ2) or “0” if the walls are open (zones merge into one larger room).

The schematic diagram of the ACMV system installed in the test space is shown in [Figure 2](#). The ACMV system serving the test space consists of two primary air-handling units (PAUs) for pre-conditioned fresh air supply and 16 fan coil units (FCUs) for cooling and dehumidification. The cooling coils in the PAUs and FCUs are supplied with chilled water from the building's central chiller plant (not shown in [Figure 2](#)). Since the MPC system of this study controls only the airside equipment of the ACMV system within the test space – a fraction of the entire building – the net effect of MPC on the overall chiller plant is not considered. Instead, cooling energy consumption, measured by British thermal unit (BTU) meters (or energy meters) installed at the cooling coils of the PAUs and FCUs, is used to evaluate the cooling energy consumed by the ACMV system in the test space. Each PAU is equipped with an on/off damper, a cooling coil, and a supply fan. The PAUs




Figure 1. Layout of spaces within the test zone




Figure 2. Sensors, actuators, and other hardware installed in the test space for model predictive control

Abbreviations: ACMV: Air conditioning and mechanical ventilation; AHU: Air handling unit; BMS: Building management system; BTU: British thermal unit; FCU: Fan coil unit; IoT: Internet of Things; PAU: Primary air-handling unit; VSD: Variable speed drive.

operate at a constant off-coil temperature setpoint of 24°C and a constant supply airflow rate of 2.2 m³/s. The 16 FCUs provide terminal air conditioning to the test space, each equipped with a cooling coil and a supply fan. The FCUs draw in a mixture of fresh air from the PAUs and return air from the test zones for cooling and dehumidification, and then supply the conditioned air back to the test zones through motorized air diffusers. Each FCU operates at a constant supply fan speed, with chilled water flowing

through its cooling coil regulated by a motorized water valve. The cooling energy consumed by each FCU is determined by converting thermal energy measured using BTU meters to electrical energy consumption. Similarly, the electrical energy consumed by each FCU’s fan is measured using an electrical power meter installed at the fan. The total cooling energy consumed by the ACMV system is calculated as the sum of cooling energy consumed by each FCU and electrical energy consumed by their fans.

When the test space is controlled by the existing building management system (referred to as “BMS” in this study), cooling power delivered to the test space is regulated by controlling the chilled water flow rate through each FCU’s cooling coil through a motorized water valve using a PID controller. The control is based on a temperature setpoint of 21°C, measured by a thermostat located at the diffuser outlet.

Each zone is equipped with a set of sensors, including globe temperature, ambient temperature, and humidity sensors for measuring thermal comfort, as well as CO₂ sensors for measuring occupancy within each zone. For the evaluation of the predicted mean vote (PMV) as per ASHRAE 55,³³ ambient air velocity, metabolic rate, and clothing insulation factor for each occupant were assumed constant at 0.1 m/s, 1.2 met, and 0.5 CLO, respectively. Duct air temperature, humidity, and flow sensors were installed at both the supply and return sides of the PAUs and FCUs to enable precise measurement of the cooling capacity of supply air for each zone. In addition, duct CO₂ sensors were installed on both the supply and return sides of FCUs. By integrating data from the duct CO₂ and flow sensors in the FCUs and PAUs with the indoor CO₂ sensors in each zone, the occupancy of each zone was determined.

A weather station, equipped with outdoor temperature, humidity, and solar radiation sensors, was installed on the rooftop of the building. The specifications of these sensors, installed in the test space, are summarized in Table 1.

3. MPC controller for ACMV

The MPC controller comprises an ML-based building dynamics predictive model and an optimization solver.

Table 1. Specifications of the sensors installed in the test space

Location	Sensors	Range	Accuracy
Outdoor	Air temperature	-40 – 120°C	±0.25°C
	Relative humidity	0 – 100%	±1.5%
	Solar radiation	0 – 2000 W/m ²	±10 W/m ²
Indoor room space	Air temperature	0 – 50°C	0.2°C
	Globe temperature	0 – 50°C	0.5°C
	Relative humidity	0 – 100%	±1.7%
ACMV system	CO ₂	0 – 2000 ppm	±30 ppm
	BTU meter	0 – 50 kW	±1%
	Elec. power meter	0 – 5 kW	±0.5%
	Duct air temperature	0 – 50°C	±0.2°C
	Duct RH	0 – 100%	±1.5%
	Duct airflow	0 – 16 m/s	±5% of the reading
	Duct CO ₂	0 – 2000 ppm	±30 ppm

Abbreviations: ACMV: Air conditioning and mechanical ventilation; BTU: British thermal unit; RH: Relative humidity.

The ML model for predicting indoor PMV is developed using a NARX feedback neural network, a recurrent neural network commonly used for time-series modeling. A detailed description of the development of the ML network for MPC is provided in Section 4. For the MPC controller, the cooling power supplied by the FCU cooling coil serves as the manipulated variable and is the control signal to be optimized. The current cooling power is measured as feedback using the BTU meter installed in each FCU within the zone. The number of occupants, representing the internal heat load, is measured using the duct CO₂ and flow sensors at the FCU and PAU, along with the indoor CO₂ sensors in each zone. Outdoor temperature and solar radiation data are obtained from the rooftop weather station. The solar heat gain from the windows (Q_{win}) is derived as a function of solar radiation, as shown in Equation I. This equation accounts for heat gain from both the shaded (by blinds) and unshaded regions of the window. Here, SR represents the ratio of the area shaded by blinds to the total window area, A_{win} is the total window area, E_{inc} is the incident solar radiation, $SHGC$ is the solar heat gain coefficient, and IAC is the indoor attenuation coefficient for blinds, which are made of light translucent fabric. The $SHGC$ values of the windows used in the study are obtained from the as-built drawings of the test building. For a given window type and climate, $SHGC$ is generally considered a constant parameter³⁴ and is commonly used in evaluating building energy performance.³⁵ As per the ASHRAE Handbook, $SHGC$ and IAC are assumed constant, with values of 0.287 and 0.75 (for the shaded region), respectively.³³

$$Q_{win} = \underbrace{SR * A_{win} * E_{inc} * SHGC * IAC}_{\text{region shaded by blinds}} + \underbrace{(1 - SR) * A_{win} * E_{inc} * SHGC}_{\text{unshaded region}} \quad (I)$$

The MPC framework for ACMV has two primary objectives: (i) minimizing the cooling power (Q_{cool}) consumed by the ACMV system and (ii) maximizing thermal comfort by maintaining the indoor PMV as close as possible to thermal neutrality (i.e., $PMV_{ref} = 0$). Feedback from room air temperature, humidity, and globe temperature sensors is used to determine PMV.³³ The objective function for ACMV is formulated as shown in Equation II:

$$J = \text{Minimize} \left(\sum_{k=0}^N \frac{W_{cool} * Q_{cool, t+k|t}}{COP} + \sum_{k=0}^N W_{PMV} * (PMV_{t+k|t} - PMV_{ref})^2 + \sum_{k=0}^N W_{\epsilon} * (\epsilon_{t+k|t})^2 \right) \quad (II)$$

This objective function is subject to the following constraints, as expressed in Equations III–IV:

$$Q_{cool,lb} \leq Q_{cool} \leq Q_{cool,ub} \tag{III}$$

$$-0.5 \leq PMV \leq 0.5 \tag{IV}$$

where Q , W , N , and ϵ refer to the cooling power, weighting factor, number of control intervals within one prediction horizon, and slack variable, respectively. The subscripts *cool*, *t*, *k*, *ref*, *lb*, and *ub* refer to cooling, current time, index of the time step, preferred PMV (i.e., $PMV = 0$), lower bound, and upper bound, respectively. The slack variable (ϵ) is used for constraint relaxation. In addition, the coefficient of performance (COP) of the cooling system is assumed to be constant at 3.7 as per the specification of the ACMV system. The three terms on the right-hand side of Equation II correspond to the cost functions related to cooling energy consumption, thermal discomfort, and constraint violation. The variables associated with the multizone MPC are summarized in Table 2. A more detailed explanation of Equations I-IV is provided in the Appendix.

4. ML model development for MPC

ML model training plays a pivotal role in developing predictive models for MPC. This section outlines the approach used to select, train, and evaluate neural network models for predicting PMV. This study frames the prediction of PMV as a regression task, given the continuous nature of this comfort index. An overview of the model training process and its integration into the MPC framework is presented in Figure 3.

Table 2. Summary of variables used in multizone MPC

Variable type in MPC	Variable type in ML models	Variables	Description
Manipulated variable	Inputs	Q_{cool}	Cooling power supplied by FCU
		Temperature	Outdoor ambient temperature
Measured disturbances		Solar heat gain	Measured from solar irradiance
		Adjacent PMV	PMV of adjacent/connected zones
		Occupancy	Occupancy measured from CO ₂ concentration in return air
		Mean radiant temperature	Mean radiant temperature
State variable		PMV	Current indoor PMV
Output variable	Output	PMV	Indoor PMV

Abbreviations: FCU: Fan coil unit; ML: Machine learning; MPC: Model predictive control; PMV: Predicted mean vote.

For this study, four ML algorithms – SVM, RF, NARX, and LSTM – were selected for evaluation. Following the testbed setup, historical operational data from the building was collected over 2 months. This dataset was then preprocessed and divided into training, validation, and testing subsets. Each ML model underwent training, prediction, and evaluation processes. The model demonstrating the best prediction accuracy was selected for the construction of the MPC system in this study. Below, we elaborate on dataset creation, model evaluation criteria, model selection (including hyperparameter tuning), and the results of the training process.

4.1. Data collection and model evaluation

4.1.1. Dataset creation

The dataset used for model training was derived from historical data collected from the baseline measurement phase from August to October 2021 as part of the project's test plan. Given Singapore's consistently low month-to-month climatic variations throughout the year,³⁴ this three-month baseline dataset, inclusive of outdoor conditions, can be considered representative of the entire year. However, it is important to acknowledge that the model should ideally be trained on a dataset capturing a full range of weather conditions and building operational characteristics. In regions experiencing significant seasonal climate variations, a more extensive dataset covering a longer period would be necessary. Furthermore, certain model parameters – such as the lag length for the NARX model or the input depth for the LSTM model, defining the context window for learning temporal patterns – may need adjustment to accommodate these variations. Such adaptation could introduce challenges related to computational costs, overfitting, and vanishing gradients, which require careful consideration. Future research will focus on validating the model's performance across diverse climatic zones and exploring strategies to enhance its generalizability.

In the current study, the input variables include cooling power, outdoor temperature, solar irradiance, occupancy, mean radiant temperature, and the current indoor PMV, whereas the output is future indoor PMV prediction. A summary of the collected data is provided in Table 2.

4.1.2. Handling historical data

The dataset comprises 15,606 time steps, recorded at 6-min intervals. Data preprocessing included dimensionality reduction and Z-score normalization to minimize scale effects and improve computational efficiency. The data were partitioned into training (70%), validation (15%), and testing (15%) subsets to facilitate model development and evaluation. The Z-score standardization formula³⁶ used is shown in Equation V:




Figure 3. Overview of ML model training and integration loop
Abbreviations: ML: Machine learning; MPC: Model predictive control.

$$z_i = \frac{x_i - \bar{x}}{s} \tag{V}$$

where z_i is the Z-score for the i^{th} data point, x_i is the i^{th} data point, s is the standard deviation, and \bar{x} is the mean.

4.1.3. Model evaluation

The performance of the models was evaluated based on the mean absolute error (MAE) and computational time. MAE is calculated using Equation VI:

$$MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| \tag{VI}$$

where n is the number of data points, y_i is the actual value of the i^{th} data point, and \hat{y}_i is the predicted value of the i^{th} data point. MAE is commonly used as a performance metric for ML model training in building control applications and has demonstrated good results in similar studies.^{37,38} Therefore, this study adopts this metric following established literature.

4.2. Training results

Before training the models for PMV prediction, their hyperparameters were optimized. Table 3 summarizes

the hyperparameters employed and their respective optimization techniques.

Table 4 compares the performance of non-ANN ML models (SVM and RF) and ANN models (NARX and LSTM) using MAE as a measure of model accuracy and maximum prediction error as an indication of generalization capability. The best results are indicated. Computational cost is represented as the sum of training and prediction times.

As seen in Table 4, SVM showed a comparable MAE to RF but had slightly poorer generalization, as indicated by its higher maximum prediction error. However, SVM significantly outperformed RF in terms of computational cost. Nonetheless, both non-ANN models demonstrated inferior MAE compared to ANN models (LSTM and NARX).

For RF, Out-of-Bag feature importance analysis can intuitively demonstrate the influence of different input variables on the output variables. This feature sensitivity analysis, presented in Figure 4, indicated that cooling power and current PMV were the most influential predictors for PMV prediction. The higher MAE observed for RF, compared to the ANN models, could be attributed to RF's lack of a time-memory structure, limiting its

Table 3. Hyperparameters optimization methods for the PMV prediction models

Model name	Hyperparameters tuned
SVM	1. Box constraint (C): Balances model complexity and training error. Higher values lead to better fit but risk overfitting. 2. Kernel scale (σ): Determines the shape of the RBF kernel, controlling the decision boundary's flexibility. 3. Epsilon (ϵ): Specifies the tolerance level for prediction errors, influencing model robustness. Bayesian optimization identified the optimal hyperparameters, with the kernel scale showing the most significant impact on accuracy.
RF	1. Number of trees: Determines the number of decision trees in the ensemble. More trees generally increase accuracy but also computational cost. 2. Minimum leaf size: Regulates the minimum number of samples in a leaf node to control model complexity and prevent overfitting. 3. Number of predictors for splitting: Adjusted to optimize feature selection at each decision split. Bayesian optimization was used for hyperparameter tuning, revealing that increasing the number of trees enhanced prediction accuracy, while smaller leaf sizes improved granularity but increased training time.
NARX neural network	The NARX model structure involved determining the number of neurons in the hidden layer and lag lengths for inputs and outputs. Bayesian optimization was applied to identify the optimal combination of hyperparameters.
LSTM	LSTM model training involved tuning the number of LSTM units and dropout rates to prevent overfitting and learning rates. Bayesian optimization and trial-and-error methods were used to achieve optimal model configurations.

Abbreviations: LSTM: Long short-term memory; NARX: Non-linear autoregressive network with exogenous inputs; PMV: Predicted mean vote; RBF: Radial basis function; RF: Random forest; SVM: Support vector machine.

Table 4. Summary of ML model training

ML model	MAE (Test dataset)	Maximum prediction error	Computing cost (s)		
			Training	Prediction	Total
SVM	0.1041	0.4810	235	1 ^a	236
RF	0.1040	0.4200	18 ^a	450	468
NARX	0.0623 ^a	0.3140 ^a	96	2	98 ^a
LSTM	0.0843	0.4830	840	26	866

Note: ^aBest results.

Abbreviations: LSTM: Long short-term memory; MAE: Mean absolute error; ML: Machine learning; NARX: Non-linear autoregressive network with exogenous inputs; RF: Random forest; SVM: Support vector machine.

applicability in scenarios characterized by strong temporal dependencies.²⁰

The NARX model exhibited the best accuracy and generalization among all models, along with the lowest computational cost. Potential improvements in NARX's accuracy could be achieved by increasing the number of hidden layers, but this would increase computational time.

The LSTM model significantly outperformed the non-ANN models in PMV prediction, displaying substantially better prediction accuracy. However, it did not show an advantage in generalization, as indicated by its maximum prediction error. The LSTM model required the highest training times due to the computational complexity of LSTM cells. Thus, its superior prediction accuracy came at a higher computational cost.

In summary, ANN models outperformed non-ANN models in terms of accuracy. Among the ANN models, the choice between NARX and LSTM depends on specific application requirements. For real-time implementation




Figure 4. Feature sensitivity analysis of predictors
Abbreviation: PMV: Predicted mean vote.

with limited computational resources, NARX is preferable. Conversely, for complex systems where accuracy is paramount, LSTM is more appropriate. Given that NARX achieved the lowest test and prediction errors, along with the lowest combined training and prediction cost, it was selected for integration into the MPC framework in this study.

5. Results and discussion

In this section, the performance of the MPC system is compared against the baseline BMS mode of operation. Data collected over a 7-day period for both the BMS and MPC systems indicated that the statistical variations of outdoor temperature and solar radiation were similar.

As shown in Figure 5A and B, the median outdoor air temperatures during the test days were 30°C for the BMS and 30.1°C for the MPC tests, while median solar radiation was approximately 309 and 274 W/m², respectively. These comparable outdoor conditions ensure that subsequent performance comparisons are based on similar heat load scenarios.

Figure 6A-C presents the time series data of measured CO₂ levels, indoor air temperature, and indoor PMV in the test room for a selected typical day. This day was chosen to illustrate the performance of BMS and MPC, as depicted in Figure 6, through a daily time series graph. The selection criterion for the typical day was based on the daily median outdoor temperature that most closely matched the median outdoor temperature of the entire 7-day measurement period. In addition, the comparison of CO₂ levels in Figure 6A indicates similar room occupancy

patterns for BMS and MPC, except for the lunch hour between 12 PM and 1 PM.

Figure 6B suggests that under BMS, the room is consistently overcooled throughout the day due to the fixed fan speed settings in the PAUs and FCUs, as well as the indoor air temperature setpoint of 21°C. This occurs regardless of occupancy or outdoor conditions. In contrast, MPC accounts for multiple inputs, including room occupancy and external heat load, to dynamically adjust the cooling power delivered to the room. As a result, MPC maintains a more stable indoor temperature and PMV profile throughout the day compared to the test building's original BMS. The variation in ACMV cooling power, measured using a BTU meter and fan power meter for a typical day under both BMS and MPC, is shown in Figure 6D. Through its predictive capability, MPC effectively regulates cooling power while enhancing indoor thermal comfort, as suggested by Figure 6C.

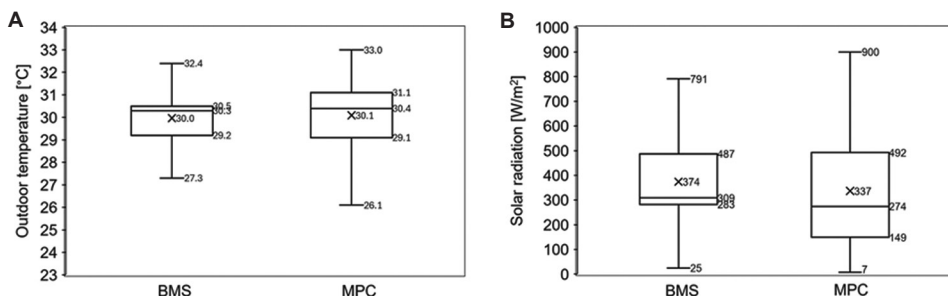


Figure 5. Weather conditions of test days. (A) Variation in outdoor temperature, (B) Variation in solar radiation.




Figure 6. Typical variations of (A) indoor CO₂ in a test space, (B) indoor temperature, (C) indoor PMV, and (D) Variation of ACMV Cooling Power for a day
Abbreviations: BMS: Building management system; MPC: Model predictive control; PMV: Predicted mean vote.

Figure 7A-E summarizes the control performance of BMS and MPC over the 1-week measurement period. Figure 7A illustrates the distribution of indoor CO₂ levels for the week under both BMS and MPC. The median CO₂ levels for BMS and MPC were found to be 610.3 and 645.4 ppm, respectively, suggesting comparable occupancy levels. This, along with the outdoor data, establishes similar heat load conditions for the ACMV system, allowing for a fair thermal comfort comparison. Figure 7B presents the temperature distribution over the seven-day test period, where the median temperature under BMS was 20.8°C compared to 22.8°C under MPC. In terms of thermal comfort, Figure 7C clearly demonstrates that BMS overcools the room, resulting in <1% of its operational time within the thermally comfortable PMV ranges (-0.5 - 0.5). Conversely, MPC is designed to maintain PMV within the comfortable zone (as defined by Equation IV), ensuring acceptable thermal comfort for more than 98% of its operation time. In addition, the small interquartile range in PMV, as seen in Figure 7C, indicates that MPC maintains stable indoor thermal comfort for a long period of time

with minimal variation. As a consequence, MPC significantly reduces cooling energy waste, which is a common issue with BMS due to overcooling. Figure 7D illustrates the reduction in total cooling power consumption, measured through BTU and fan power meters. Given Singapore’s consistently hot and humid climate throughout the year, the daily cooling energy consumption recorded over the 7-day test period was extrapolated to estimate annual consumption. Figure 7E shows that by leveraging predictive capabilities and multiobjective optimization, MPC reduces energy consumption by over 42% compared to BMS. Long-term performance sustainability of MPC can be achieved through periodic online learning, which calibrates MPC to account for changes in occupancy patterns, climatic conditions, and/or building dynamics. In our previous study, we successfully deployed such an adaptive MPC system incorporating online learning.³⁹

A further breakdown of the total cooling energy consumption associated with FCUs was conducted. The FCU fans installed in the test space were constant-speed fans, meaning that electrical power consumption by these

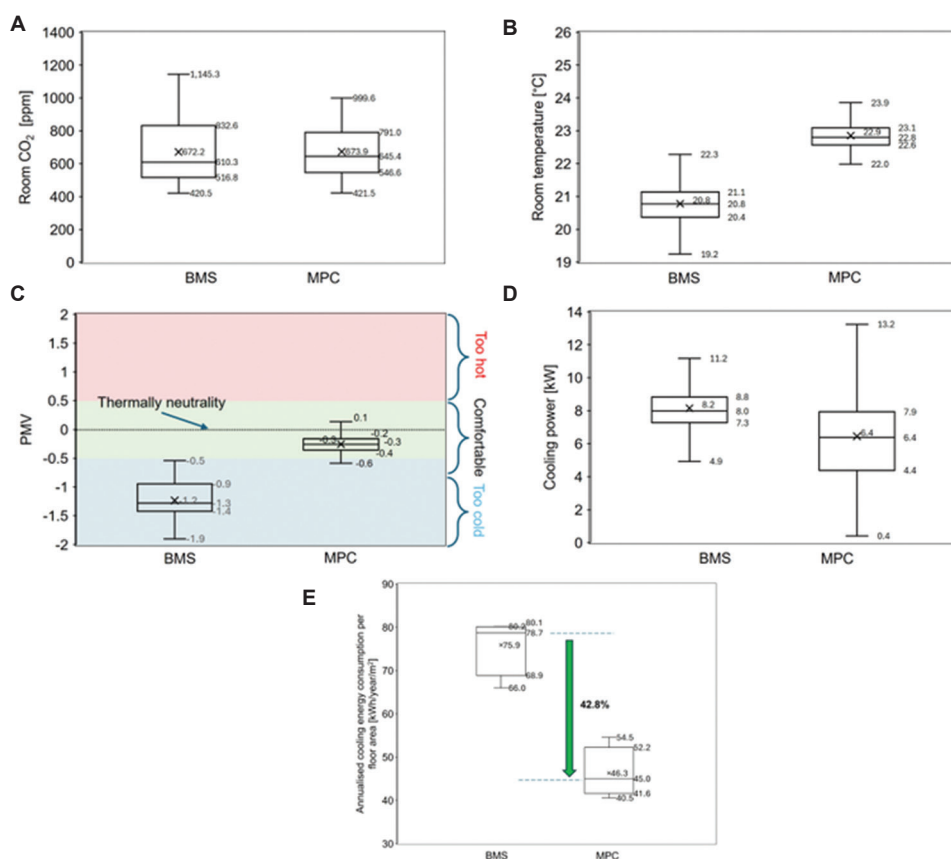


Figure 7. Statistical distributions of indoor conditions on test days: (A) indoor CO₂, (B) indoor temperature, (C) indoor PMV, (D) ACMV cooling power, and (E) annualized cooling energy consumption per floor area
Abbreviations: ACMV: Air conditioning and mechanical ventilation; BMS: Building management system; MPC: Model predictive control; PMV: Predicted mean vote.

fans remained approximately the same for both baseline BMS and MPC. During MPC duration, it was found that approximately 32% of the total cooling energy consumption was attributable to electric fan power. If FCU fans equipped with variable-speed drives were installed instead of constant-speed fans, MPC would have the opportunity to further optimize fan power consumption. For example, MPC could be configured to reduce FCU fan speed by over 70% compared to constant-speed operation when no occupants are present in the test rooms. By adopting such control logic for optimizing fan operation, a conservative reduction of 10 – 15% in fan power consumption was estimated. This estimation excludes the reduction in cooling energy produced by chilled water (which can be measured by the BTU meter) during the heat exchange process with the reduced airflow in the FCU. Such a detailed impact of reduced fan power on cooling energy consumption will be investigated in future studies. Excluding the potential effect of airflow reduction on chilled water cooling energy, an additional 3 – 5% reduction in total cooling energy consumption was estimated.

A key advantage of our MPC implementation is its modular and adaptable framework, involving: (i) collecting building data; (ii) developing a control-oriented model using ML tools; (iii) designing the optimization algorithm; and (iv) integrating the MPC with the BMS. While this framework is generally applicable to buildings in various climates and with diverse operational profiles, building-specific retraining is necessary to accurately capture unique thermal dynamics. In environments with significant seasonal or regional variations, model parameters (e.g., lag length for NARX or input depth for LSTM) may require adjustment, posing challenges such as increased computational cost and potential overfitting. Future studies will further validate the framework across different climatic zones.

A key challenge in MPC deployment is balancing model complexity against computational cost. In our current implementation, the NARX model provided superior accuracy and efficiency (Table 4). For more complex systems, LSTM architectures – with increased input depth and larger hidden layers – could offer enhanced prediction accuracy, albeit with higher computational demands. In such cases, employing artificial intelligence/ML-optimized edge hardware and multithreaded computing can help manage computational loads. In our previous work,¹¹ we introduced an approach using instantaneous linearization within ML-based MPC to further address these challenges. Although a detailed exploration of this topic is beyond the scope of the present paper, it represents a promising direction for future research.

6. Conclusion

This study demonstrated the effectiveness of an ML-based MPC system in optimizing energy consumption and

occupant comfort in a multizone commercial building. The implemented system achieved energy savings exceeding 42% compared to conventional thermostat-based control while simultaneously improving thermal comfort and stability. The novelty of our approach lies in its real-world, large-scale deployment in an 11-zone building, integration of ML-based weather forecasting for proactive control adjustments, and the use of a NARX neural network for accurate PMV prediction.

While initial implementation costs and model development efforts present challenges for widespread adoption, future research endeavors aim to mitigate these limitations. Such efforts include developing comfort models that rely exclusively on existing building sensors, exploring cost-effective sensor technologies, and examining the integration of additional building systems, such as lighting and shading, for more holistic optimization strategies. The present study contributes to advancing intelligent building automation systems, with significant potential to enhance energy efficiency and sustainability in the built environment.

Acknowledgments

Technical support from Dr. Wai Soong Loh from JTC for this project is much appreciated. We also appreciate the patience, understanding, and logistical support of the administrative team from Civil Service College, Singapore, for this project.

Funding

This research is jointly supported by Jurong Town Corporation (JTC) of Singapore (NTU REF 2019-0607) and Smart Nation and Digital Government Office, SNDGO of Singapore (NRF2016IDM-TRANS001-031).

Conflict of interest

The authors declare they have no competing interests.

Author contributions

Conceptualization: Pradeep Shakya, Man Pun Wan

Data Curation: Shiva Sreenivasan

Formal analysis: Pradeep Shakya, Man Pun Wan

Funding acquisition: Shiyu Yang, Man Pun Wan

Investigation: Shiva Sreenivasan, Pradeep Shakya

Methodology: Pradeep Shakya

Resources: Baskaran Krishnamoorthy, Shiyu Yang

Writing – original draft: Pradeep Shakya, Shiva Sreenivasan

Writing – review & editing: Man Pun Wan, Pradeep Shakya, Shiva Sreenivasan

Ethics approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Availability of data

The datasets presented in this article are not readily available due to a confidentiality agreement with the funding agency.

Further disclosure

Part of or the entire set of findings have been presented at 2024 Herrick Conferences in Compressor Engineering, Refrigeration and Air Conditioning and High-Performance Buildings (held on July 14, 2024, at Purdue University, West Lafayette, Indiana, United States of America) hosted by the Ray W. Herrick Laboratories and the Center for High-Performance Buildings in West Lafayette, Indiana, United States of America.

References

1. Building and Construction Authority (BCA). *Super Low Energy Building Technology Roadmap*. Building and Construction Authority; 2018. Available from: https://www1.bca.gov.sg/docs/default-source/docs-corp-buildsg/sustainability/sle-tech-roadmap-report--published-ver1-1.pdf?sfvrsn=f2df22ed_0 [Last accessed on 2024 Dec 05].
2. Singapore Green Building Council (SGBC). *Singapore Green Building Masterplan*. Singapore Green Building Council; 2020. Available from: <https://www.sgbc.sg/about-green-building/sgbmp> [Last accessed on 2024 Dec 05].
3. Building and Construction Authority (BCA). *Green Building Masterplan*. Building and Construction Authority; 2021. Available from: https://www1.bca.gov.sg/docs/default-source/docs-corp-buildsg/sustainability/sgbmp-80-80-80-in-2030-infographic.pdf?sfvrsn=57172d48_2 [Last accessed on 2024 Dec 05].
4. Das HP, Lin YW, Agwan U, *et al.* Machine learning for smart and energy-efficient buildings. *Environ Data Sci*. 2024;3:e1. doi: 10.1017/eds.2023.43
5. Qiang G, Tang S, Hao J, Di Sarno L, Wu G, Ren S. Building automation systems for energy and comfort management in green buildings: A critical review and future directions. *Renew Sustain Energy Rev*. 2023;179:113301. doi: 10.1016/j.rser.2023.113301
6. Pinheiro S, Wimmer R, O'Donnell J, *et al.* MVD based information exchange between BIM and building energy performance simulation. *Autom Constr*. 2018;90:91-103. doi: 10.1016/j.autcon.2018.02.009
7. Mauri L, Vallati A, Ocloń P. Low impact energy saving strategies for individual heating systems in a modern residential building: A case study in Rome. *J Clean Prod*. 2019;214:791-802. doi: 10.1016/j.jclepro.2018.12.320
8. Mathews Roy A, Prasanna Venkatesan R, Shanmugapriya T. Simulation and analysis of a factory building's energy consumption using eQuest software. *Chem Eng Technol*. 2021;44(5):928-933. doi: 10.1002/ceat.202000489
9. Andriamamonjy A, Saelens D, Klein R. An automated IFC-based workflow for building energy performance simulation with Modelica. *Autom Constr*. 2018;91:166-181. doi: 10.1016/j.autcon.2018.03.019
10. Yang S, Wan MP, Ng BF, *et al.* A state-space thermal model incorporating humidity and thermal comfort for model predictive control in buildings. *Energy Build*. 2018;170:25-39. doi: 10.1016/j.enbuild.2018.03.082
11. Yang S, Wan MP. Machine-learning-based model predictive control with instantaneous linearization-A case study on an air-conditioning and mechanical ventilation system. *Appl Energy*. 2022;306:118041. doi: 10.1016/j.apenergy.2021.118041
12. Široký J, Oldewurtel F, Cigler J, Prívará, S. Experimental analysis of model predictive control for an energy efficient building heating system. *Appl Energy*. 2011;88:3079-3087. doi: 10.1016/j.apenergy.2011.03.009
13. Ma Y, Borrelli F, Hency B, Coffey B, Bengesa S, Haves P. Model predictive control for the operation of building cooling systems. *IEEE Trans Control Syst Technol*. 2012;20:796-803. doi: 10.1109/TCST.2011.2124461
14. Pang X, Duarte C, Haves P, Chuang F. Testing and demonstration of model predictive control applied to a radiant slab cooling system in a building test facility. *Energy Build*. 2018;172:432-441. doi: 10.1016/j.enbuild.2018.05.013
15. Yang Y, Srinivasan S, Hu G, Spanos CJ. Distributed control of multizone HVAC systems considering indoor air quality. *IEEE Trans Control Syst Technol*. 2021;29:2586-2597. doi: 10.1109/TCST.2020.3047407
16. Yang S, Chen W, Wan MP. A machine-learning-based event-triggered model predictive control for building energy management. *Build Environ*. 2023;233:110101. doi: 10.1016/j.buildenv.2023.110101
17. Fathi S, Srinivasan R, Fenner A, Fathi S. Machine learning applications in urban building energy performance forecasting: A systematic review. *Renew Sustain Energy Rev*. 2020;133:110287. doi: 10.1016/j.rser.2020.110287
18. Seyedzadeh S, Rahimian FP, Rastogi P, Glesk I. Tuning machine learning models for prediction of building energy loads. *Sustain Cities Soc*. 2019;47:101484.

- doi: 10.1016/j.scs.2019.101484
19. Zhou X, Xu L, Zhang J, *et al.* Data-driven thermal comfort model via support vector machine algorithms: Insights from ASHRAE RP-884 database. *Energy Build.* 2020;211:109795.
doi: 10.1016/j.enbuild.2020.109795
 20. Chaudhuri T, Zhai D, Soh YC, Li H, Xie L. Random forest based thermal comfort prediction from gender-specific physiological parameters using wearable sensing technology. *Energy Build.* 2018;166:391-406.
doi: 10.1016/j.enbuild.2018.02.035
 21. Koschwitz D, Frisch J, Van Treeck C. Data-driven heating and cooling load predictions for non-residential buildings based on support vector machine regression and NARX Recurrent Neural Network: A comparative study on district scale. *Energy.* 2018;165:134-142.
doi: 10.1016/j.energy.2018.09.068
 22. Zhang C, Li J, Zhao Y, Li T, Chen Q, Zhang X. A hybrid deep learning-based method for short-term building energy load prediction combined with an interpretation process. *Energy Build.* 2020;225:110301.
doi: 10.1016/j.enbuild.2020.110301
 23. Beltran A and Cerpa AE. Optimal HVAC Building Control with Occupancy Prediction. In: *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-efficient Buildings*; 2014. p. 168-171.
doi: 10.1145/2674061.2674072
 24. Li B, Xia L. A Multi-grid Reinforcement Learning Method for Energy Conservation and Comfort of HVAC in Buildings. In: *2015 IEEE International Conference on Automation Science and Engineering (CASE)*; 2015. p. 444-449.
doi: 10.1109/CoASE.2015.7294119
 25. Zhang Z and Lam KP. Practical Implementation and Evaluation of Deep Reinforcement Learning Control for a Radiant Heating System. In: *Proceedings of the 5th Conference on Systems for Built Environments*; 2018. p. 148-157.
doi: 10.1145/3276774.3276775
 26. Ding X, Cerpa A, Du W. Exploring deep reinforcement learning for holistic smart building control. *ACM Trans Sens Netw.* 2024;20(3):1-28.
doi: 10.1145/3656043
 27. Shamachurn H, Seebaruth M, Kowlessur NS, Hassen SS. Real-time model predictive control of air-conditioners through IoT-results from an experimental setup in a tropical climate. *Adv Control Appl Eng Ind Syst.* 2024;6:e232.
doi: 10.1002/adc2.232
 28. Hu G, You F. Multi-zone building control with thermal comfort constraints under disjunctive uncertainty using data-driven robust model predictive control. *Adv Appl Energy.* 2023;9:100124.
doi: 10.1016/j.adapen.2023.100124
 29. Joe J, Im P, Cui B, Dong J. Model-based predictive control of multi-zone commercial building with a lumped building modelling approach. *Energy.* 2023;263:125494.
doi: 10.1016/j.energy.2022.125494
 30. Oldewurtel F, Parisio A, Jones CN, *et al.* Use of model predictive control and weather forecasts for energy efficient building climate control. *Energy Build.* 2012;45:15-27.
doi: 10.1016/j.enbuild.2011.09.022
 31. Hou J, Li H, Nord N, Huang G. Model predictive control under weather forecast uncertainty for HVAC systems in university buildings. *Energy Build.* 2022;257:111793.
doi: 10.1016/j.enbuild.2021.111793
 32. Mazar MM, Rezaeizadeh A. Adaptive model predictive climate control of multi-unit buildings using weather forecast data. *J Build Eng.* 2020;32:101449.
doi: 10.1016/j.jobe.2020.101449
 33. ASHRAE Handbook. *Heating, Ventilating, and Air-Conditioning Systems and Equipment.* Vol. 39. Atlanta, GA, USA: American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.; 1996.
 34. Meteorological Service Singapore (MSS). *Climate of Singapore.* Meteorological Service Singapore; 2020. Available from: <https://www.weather.gov.sg/climate-climate-of-singapore> [Last accessed on 2024 Dec 05].
 35. Yang S, Wan MP, Ng BF, *et al.* Model predictive control for integrated control of air-conditioning and mechanical ventilation, lighting and shading systems. *Appl Energy.* 2021;297:117112.
doi: 10.1016/j.apenergy.2021.117112
 36. Freedman DA. *Statistical Models: Theory and Practice.* United Kingdom: Cambridge University Press; 2009.
 37. Chifu VR, Pop CB, Chifu ES, Barleanu H. Deep Learning for Forecasting the Energy Consumption in Public Buildings. In: *2021 20th RoEduNet Conference: Networking in Education and Research (RoEduNet)*; 2021. p. 1-6.
doi: 10.48550/arXiv.2207.11953
 38. Yang R, Hao J, Jiang H, Jin X. *Machine-Learning-Driven, 2020, Site-Specific Weather Forecasting for Grid-Interactive Efficient Buildings.* Golden, CO, United States: National Renewable Energy Lab. (NREL); 2020. Available from: <https://www.osti.gov/biblio/1669587> [Last accessed on 2025 Feb 06].
 39. Yang S, Wan MP, Chen W, Ng BF, Dubey S. Model predictive control with adaptive machine-learning-based model for building energy efficiency and comfort optimization. *Appl Energy.* 2020;271:115147.
doi: 10.1016/j.apenergy.2020.115147
 40. Zhao S, Cajo R, De Keyser R, Liu S, Ionescu CM. Nonlinear predictive control applied to steam/water loop in large scale ships. *IFAC PapersOnLine.* 2019;52(1):868-873.
doi: 10.1016/j.ifacol.2019.06.171

Appendix

1. Elaboration on MPC Equations

This appendix provides a comprehensive explanation of the equations presented in Section 3, which govern the MPC controller for the ACMV system.

Equation 1: Solar heat gain

Solar heat gain through windows is a critical component of building thermal load modeling. The formulation accounts for dynamic shading effects, which modulate solar radiation entering space. The total solar heat gain, Q_{win} , is expressed as:

$$Q_{win} = \underbrace{SR * A_{win} * E_{inc} * SHGC * IAC}_{\text{region shaded by blinds}} + \underbrace{(1 - SR) * A_{win} * E_{inc} * SHGC}_{\text{unshaded region}} \quad (I)$$

Description of key variables in Equation I:

- Shading ratio (SR): Represents the fraction of the window area covered by blinds. SR can vary dynamically based on occupant preferences or automated shading control strategies
- Solar heat gain coefficient (SHGC): Typically ranges between 0.3 and 0.9 depending on glazing properties (e.g., low-emissivity coatings, double/triple glazing). SHGC values are often derived from standardized tests (ASHRAE 90.1) or manufacturer data
- Indoor attenuation coefficient (IAC): Quantifies the reduction in solar radiation due to blinds. For example, horizontal slat blinds with high reflectivity may have an IAC of 0.5 – 0.7, while blackout shades could reduce IAC to near zero. For this study, as per ASHRAE Handbook,³³ IAC of 0.75 was assumed for blinds made of light translucent fabric
- Incident radiation (E_{inc}): Includes direct, diffuse, and reflected solar radiation, which depends on window orientation, time of day, and geographic location.

Expanded discussion for Equation I:

The equation assumes uniform shading distribution across the window, which may simplify real-world scenarios where partial shading or non-uniform blind deployment occurs. The model does not explicitly account for spectral properties of solar radiation or transient thermal storage in glazing materials, though these effects are often negligible for hourly or sub-hourly MPC timescales. For further validation, SHGC and IAC values should align with empirical measurements or established databases (e.g., Lawrence Berkeley National Laboratory's Window Module).

Equation II: Objective function

The MPC controller's objective function, J , balances energy efficiency, thermal comfort, and operational feasibility. Building on the framework from,³⁹ the cost function is defined as:

$$J = \text{Minimize} \left(\sum_{k=0}^N \frac{W_{cool} * Q_{cool, t+k|t}}{COP} + \sum_{k=0}^N W_{PMV} * (PMV_{t+k|t} - PMV_{ref})^2 + \sum_{k=0}^N W_{\epsilon} * (\epsilon_{t+k|t})^2 \right) \quad (II)$$

Description of key variables in Equation II:

- Normalized cooling power (Q_{cool}): The first term uses cooling power normalized by the system's maximum capacity to ensure scalability across different HVAC systems
- Normalized PMV (PMV): The second term penalizes deviations from the neutral setpoint ($PMV_{ref} = 0$), normalized by the acceptable comfort range (± 0.5). This normalization ensures equitable weighting between thermal comfort and energy use.
- Weighting factors:
 - $W_{cool} = 1/10$: Prioritizes thermal comfort over energy savings, reflecting occupant-centric design principles
 - $W_{PMV} = 4$: Prioritizes thermal comfort over energy savings, reflecting occupant-centric design principles
 - $W_{\epsilon} = 10,000$: Strongly penalizes constraint violations to enforce strict adherence to comfort and equipment limits.
- Prediction horizon: Set to 12 control intervals (60 minutes), chosen to match the thermal response time of the building (~40 minutes for PMV stabilization during morning start-up). A longer horizon would increase computational complexity without improving performance.^{39,40}

Practical implications for Equation II:

The weighting scheme reflects a design philosophy where occupant comfort is prioritized, but energy efficiency remains a secondary goal. The slack variable ϵ ensures feasibility under unexpected disturbances (e.g., occupancy spikes), though its high penalty weight (W_ϵ) minimizes its use in practice.

Equations III and IV

The optimization is bounded by operational and comfort limits:

$$Q_{cool,lb} \leq Q_{cool} \leq Q_{cool,ub} \quad (III)$$

$$-0.5 \leq PMV \leq 0.5 \quad (IV)$$

Contextual details for Equations III and IV:

- Cooling power bounds ($Q_{cool, lb}$, $Q_{cool, ub}$) are determined by the HVAC system's capacity
- The PMV range (± 0.5) corresponds to the ASHRAE Standard 55 "neutral" comfort zone, ensuring 80 – 90% occupant satisfaction.

Additional notes:

- Control interval: The MPC operates at 5-min intervals, balancing responsiveness with computational tractability
- Slack V=variable (ϵ): Introduced as a soft constraint to avoid infeasible solutions during transient disturbances (e.g., abrupt weather changes)
- Thermal comfort: The PMV model assumes steady-state conditions and standardized clothing/metabolic rates; dynamic adjustments for occupant activity are not modeled here.