Energy Proceedings
Vol 60, 2025
ISSN 2004-2965

Occupant-Feedback-Integrated Model Predictive Control for Air Conditioning: An Experimental Study in a Commercial Building#

Sreenivasan Shiva [1,2], Shiyu Yang [3*], Shakya Pradeep [4], Krishnamoorthy Baskaran [4], Loh Wai Soong [5], Wan Man Pun [2]

1 Energy Research Institute @ NTU, Interdisciplinary Graduate Program, Nanyang Technology University, Singapore

2 Mechanical and Aerospace Engineering, Nanyang Technology University, Singapore

3 State Key Laboratory of Electrical Insulation and Power Equipment, Xi'an Jiaotong University, China

4 Energy Research Institute @ NTU, Singapore

5 JTC Corporation, Singapore

(Corresponding Author: yangshiyu@xjtu.edu.cn)

ABSTRACT

This paper presents the implementation and analysis of a ML-based MPC system integrated with a real-time occupant feedback mechanism through field tests in a commercial building. We demonstrate the efficacy of the occupant-centric control strategy that dynamically adjusts thermal comfort setpoints based on direct user input. Furthermore, we propose a comprehensive framework for analyzing the resulting data, offering insights into system performance, energy consumption, and occupant satisfaction. The findings illustrate that incorporating a human-in-the-loop approach can enhance building energy efficiency without compromising occupant comfort.

Keywords: Model Predictive Control, Machine Learning, AI, PMV, ACMV, Occupant Feedback

NONMENCLATURE

Abbreviations	
BAC	Building Automation and Control
MPC	Model Predictive Control
PID	Proportional-Integral-Derivative
ACMV	Air-Conditioning and Mechanical Ventilation
PMV	Predictive Mean Vote
BMS	Building Management System
ML	Machine Learning
NARX	Nonlinear Autoregressive Exogenous
PAU	Primary Air Handling Unit
FCU	Fan Coil Unit
Symbols	
W_{Term}	Weight for respective Term
COP	Coefficient of Performance
Q_{cool}	Cooling Power
ϵ	Slack Variable

1. INTRODUCTION

The building sector is a primary target for energy reduction initiatives in the global pursuit of carbon neutrality [1]. Advanced control strategies for BAC systems, such as MPC, offer a promising path to improving energy efficiency without sacrificing occupant comfort [2]. Unlike traditional reactive methods (e.g., PID control), MPC leverages predictive models to proactively optimize building operations, forecasting the impact of disturbances like weather and occupancy changes [3].

While the literature confirms MPC’s ability to yield substantial energy savings [4,5], a critical gap remains: standard MPC optimizes for static comfort setpoints, failing to account for the heterogeneous and time-varying thermal preferences of actual occupants. This often leads to comfort complaints and energy waste [6, 7].

This paper addresses this gap by developing and validating an MPC framework that incorporates a real-time "human-in-the-loop" feedback mechanism. By allowing occupant preferences to dynamically adjust control targets, our approach aims to enhance comfort, reduce energy consumption, and guide future implementations of smarter, occupant-centric building controls.

This is a paper for the 5th Applied Energy Symposium and Forum: Renewable Energy Matrix (REM2025), Oct. 29-31, 2025, Yancheng, China.

2. LITERATURE REVIEW

2.1 MPC for Building Automation

Conventional BAC systems use reactive strategies that respond to disturbances only after they occur. In contrast, MPC's predictive nature and ability to manage competing objectives—such as energy use and occupant well-being—make it a superior alternative for optimizing the complex thermal dynamics of buildings [3, 8].

Numerous studies have demonstrated MPC's effectiveness. Researchers have reported energy savings between 15% and 42% by integrating weather forecasts and optimizing chiller plants and radiant systems [10-12]. While widespread adoption faces hurdles like computational load and model fidelity, recent work addresses these challenges through distributed computation, linearization and event-triggered optimization schemes [5, 9, 13].

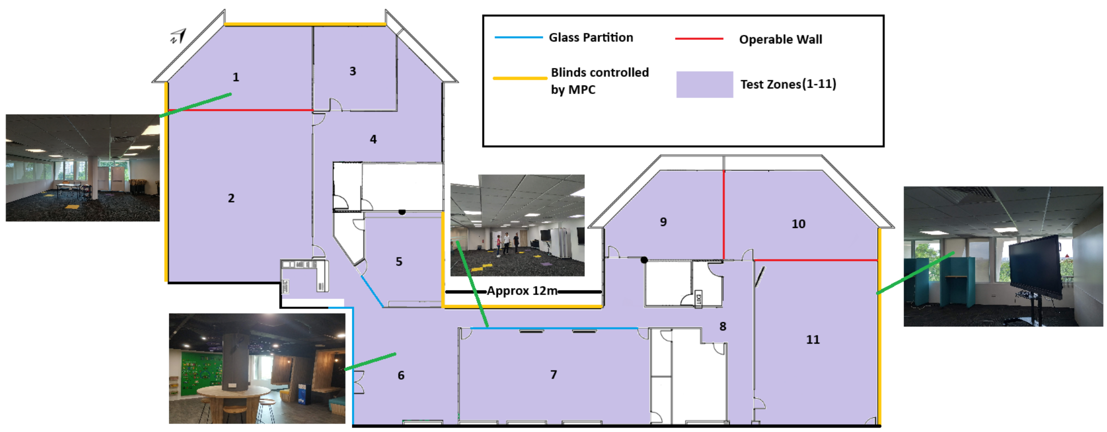

Fig. 1 Layout of spaces in the test zone with a few site pictures to show some of the features such as operable wall, glass partition and controllable blinds.

2.2 Research Gap: From Static Setpoints to Dynamic Comfort

Despite these advances, a fundamental limitation of conventional MPC is its reliance on a predefined, static comfort target (e.g., a Predicted Mean Vote of 0). This approach struggles to capture the subjective, dynamic nature of human comfort, often resulting in conditions where building operation misaligns with occupant needs [8, 14].

This research directly confronts this gap. Our work is situated within the context of an advanced, machine learning-based MPC controller that uses a NARX neural network to forecast indoor conditions. The controller coordinates ACMV, lighting, and shading systems by optimizing against inputs from weather stations, indoor sensors, and occupancy detectors [3, 8]. The primary innovation is the integration of a web-based occupant feedback system. This allows occupants to submit preferences ("Too cold," "OK," "Too hot"), which directly bias the MPC's reference PMV setpoint. This human-in-the-loop model allows the system to move beyond generic targets and adapt control strategies to the collective, real-time preference of its occupants.

3. METHODOLOGY AND TEST BED

3.1 Testbed and Instrumentation

The study was conducted in an 850 m² multi-use space within a commercial building in Singapore, operating on weekdays from 7 AM to 7 PM. The test space consists of 11 zones served by a central chiller plant.

The ACMV system uses PAUs for fresh air and zone-level FCUs for cooling. The MPC system's primary control action is to regulate the chilled water flow to each FCU via motorized valves. For comparison, the baseline BMS uses a PID controller with a fixed 22°C temperature setpoint.

Each zone is instrumented with sensors for globe/ambient temperature, humidity, and CO_2. A rooftop weather station provides outdoor conditions. All data is collected at one-minute intervals.

3.2 MPC Formulation

The MPC system's objective function, J, is formulated to minimize cooling power consumption while driving the PMV towards its reference setpoint (PMVref), with a penalty term for constraint violations:

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

where the cooling load and PMV are subject to the following hard constraints:

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

3.3 Occupant Feedback System

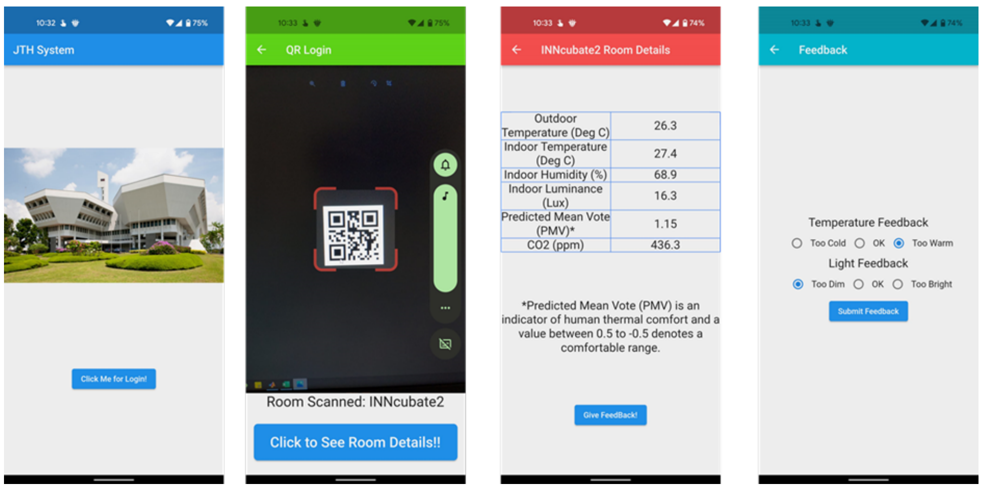

Fig. 2 Screenshots of the user feedback application

A key innovation is the integration of a web-based occupant feedback system with the MPC controller (Figure 2). Occupants in each zone can scan a unique QR code to access a web app, view real-time conditions, and submit their thermal preference by selecting "Too cold", "OK", or "Too hot".

The system uses a voting mechanism to translate this feedback into control actions. The logic is as follows:

- Tallying Window: Feedback is tallied every 18 minutes.
- Participation Threshold: An adjustment is only considered if at least 20% of the zone's estimated occupants have voted within the window.
- Dynamic Adjustment: If a clear majority indicates discomfort (e.g., "Too hot" votes are more than double the "Too cold" votes), the MPC dynamically adjusts the reference PMV (PMVref) by ±0.2, bounded within [-0.5, +0.5].
- Inactivity Rules: To prevent reactions to sparse or biased data, no adjustments are made if no votes have been recorded in the past hour or during specific quiet periods (e.g., lunch hour).

This rule set ensures the controller makes bounded, incremental adjustments only when there is sufficient and consistent evidence of collective occupant discomfort, thereby enhancing comfort personalization while mitigating system instability.

3.4 Data Filtering and Day Selection Methodology

To ensure a valid comparison, a systematic methodology was established for identifying and matching feedback and non-feedback days based on occupancy and environmental conditions.

3.4.1 Feedback Day Identification

A two-stage filtering process was used to isolate days with meaningful occupant feedback:

Occupancy Pre-filtering: Days were first filtered to include only periods with significant occupancy, defined by an average CO2 concentration exceeding 420 ppm during business hours (9 AM to 6 PM). This step ensures that the analysis focuses on periods when feedback is relevant.

Feedback Event Detection: Within this set of occupied days, any day containing at least one thermal comfort feedback event (e.g., "Too hot" or "Too cold") was classified as a "feedback day".

3.4.2 Non-Feedback Day Matching

To create a valid baseline for comparison, each feedback day was matched with a non-feedback day exhibiting similar conditions. A weighted similarity score was developed to facilitate this matching, combining three key factors: CO2 (60% weight), Outdoor Temperature (20% weight), and Solar Radiation (20% weight).

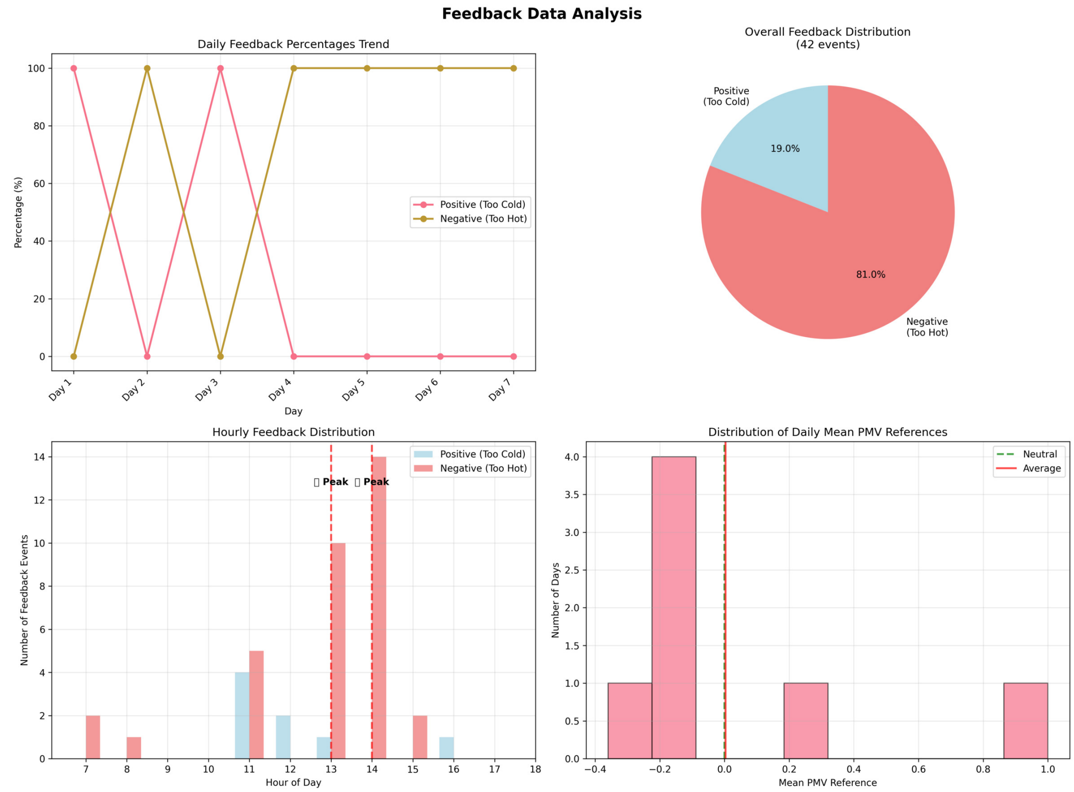

Fig. 3 Comprehensive analysis of feedback trends, showing the distribution of feedback types and the resulting daily mean PMV reference

4. RESULTS

To evaluate the system, we compared "feedback days" against "non-feedback days" that were carefully matched for similar occupancy (using CO_2 levels as a proxy) and weather conditions. This methodology isolates the impact of the occupant feedback mechanism. The analysis revealed three key findings:

Occupant Feedback Patterns: The majority of feedback was "Too hot," peaking midday, which is expected in Singapore's climate. Despite requests for cooler conditions, the average PMV reference derived from feedback remained close to 0, validating PMV as a robust index for thermal neutrality (Figure 3).

Energy Impact: Energy consumption on feedback days was approximately 12% lower than on matched non-feedback days (Figure 4). This suggests that adapting to occupant requests does not incur an energy penalty. However, we note that slightly lower average CO_2 levels on feedback days indicate that minor variations in occupancy might be a confounding variable (Figure 5).

System Responsiveness: The controller demonstrated a high degree of effectiveness, responding appropriately to feedback in 82.4% of all cases by adjusting the indoor temperature in the requested direction.

5. DISCUSSION

The results yield two primary conclusions. First, the integration of a “human-in-the-loop” feedback system into an MPC framework is not only feasible but also effective, allowing for dynamic, occupant-centric control without compromising energy efficiency. Second, the convergence of the occupant-driven PMV reference towards thermal neutrality reinforces the validity of PMV as a foundational metric for thermal comfort in advanced control strategies. The primary limitation encountered in this study was the inherent difficulty in finding perfectly matched feedback and non-feedback days from the available data pool. Minor variations in occupancy and micro-climatic conditions make it challenging to completely isolate the energy impact of feedback events. Furthermore, while the system's 82.4% response effectiveness is strong, it also highlights an area for further improvement. The speed and precision of the system's reaction could be enhanced. Future work should focus on refining the control logic. This could involve exploring more advanced algorithms to reduce the MPC's execution time or developing a more sophisticated method for translating raw feedback into an adjusted PMV reference. Such improvements would further enhance the system's ability to provide a personalized, comfortable, and energy-efficient indoor environment.

6. CONCLUSIONS

This paper presented the successful implementation and evaluation of a machine learning-based Model Predictive Control system enhanced with a real-time occupant feedback mechanism. We introduced a systematic framework for analyzing this "human-in-the-loop" system, demonstrating its practical feasibility and its ability to create a highly responsive indoor environment. The system achieved 82.4% effectiveness in reacting to occupant feedback. Crucially, this was achieved without an energy penalty; in fact, an energy reduction of 12% was observed on feedback days compared to matched non-feedback days, though this may be partially attributed to slight occupancy variations. The study also validated the PMV index, as occupant-driven comfort preferences collectively gravitated towards thermal neutrality.

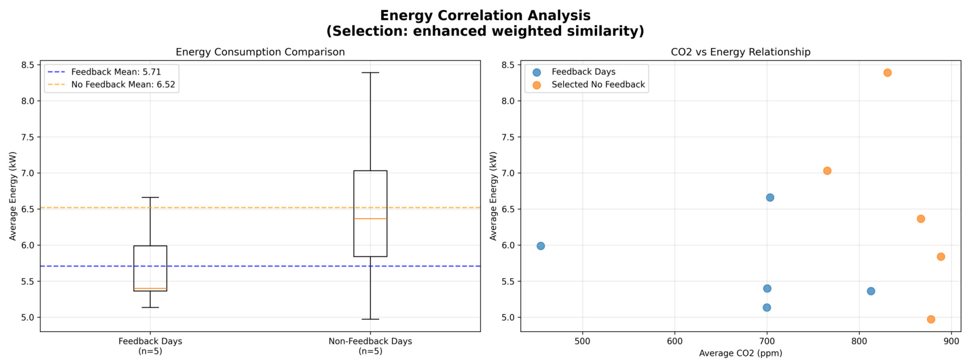

Fig. 4 Energy correlation analysis comparing total energy consumption on feedback days versus matched non-feedback days

While this work establishes a viable framework, future enhancements could address several key areas for improvement. The control logic could be refined by replacing the fixed ±0.2 PMV adjustment with an adaptive step-size mechanism that reflects the intensity of occupant consensus (e.g., larger adjustments for stronger majorities). To move beyond the current "one-size-fits-all" approach, personalization could be deepened by implementing user profiles to weight votes based on individual thermal sensitivities. Finally, to

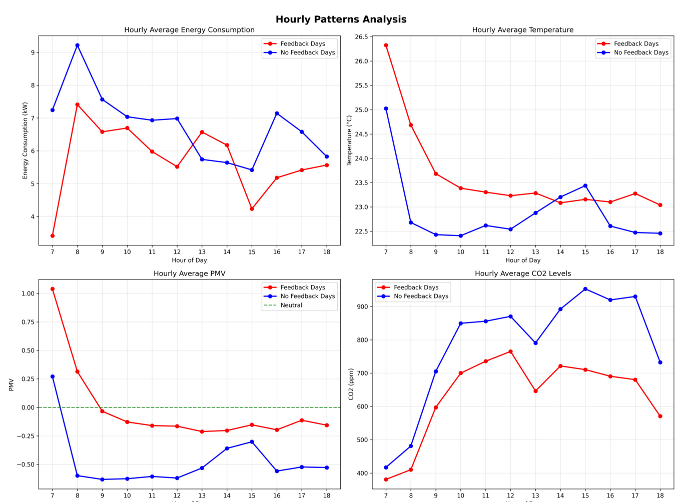

Fig. 5 Comparison of hourly patterns for key variables (Energy, Temperature, PMV, CO2) between feedback and non-feedback days

overcome the limitations of manual submissions and low participation, the system could be augmented with passive sensing data from wearable or environmental sensors to create a richer, more continuous feedback stream.

Ultimately, integrating these advancements will pave the way for future building automation systems that are not only energy-efficient but also truly adaptive and personalized to the needs of their inhabitants.

ACKNOWLEDGEMENT

This research is jointly supported by JTC Corporation of Singapore through contract no. N190107T00 and National Research Foundation (NRF) of Singapore through grant no. TRANS2021-TU002. Technical support from Dr. Loh Wai Soong (JTC Corporation) is highly appreciated.

REFERENCE

[1] M. Company, A strategic approach to climate action in cities—focused acceleration, Tech. rep., McKinsey & Company (2017).
[2] A. Afram, F. Janabi-Sharifi, Theory and applications of hvac control systems – a review of model predictive control (mpc), Building and Environment 72 (2014) 343–355.
[3] S. Yang, M. P. Wan, B. F. Ng, S. Dubey, G. P. Henze, W. Chen, K. Baskaran, Model predictive control for integrated control of air-conditioning and mechanical ventilation, lighting and shading systems, Applied Energy 297 (2021) 117112.
[4] 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, Applied Energy 271 (2020) 115147.
[5] Y. Yang, S. Srinivasan, G. Hu, C. J. Spanos, Distributed control of multizone hvac systems considering indoor air quality, IEEE Transactions on Control Systems Technology 29 (6) (2021) 2586–2597.
[6] S. Yang, M. P. Wan, W. Chen, B. F. Ng, S. Dubey, Experiment study of machine-learning-based approximate model predictive control for energy-efficient building control, Applied Energy 288 (2021) 116648.
[7] S. Yang, M. P. Wan, W. Chen, B. F. Ng, D. Zhai, An adaptive robust model predictive control for indoor climate optimization and uncertainties handling in buildings, Building and Environment 163 (2019) 16106326.
[8] P. Shakya, S. Sreenivasan, B. Krishnamoorthy, S. Yang, M. P. Wan, Machine learning-based model predictive control for multizone building automation: A case study, International Journal of AI for Materials and Design 2 (1) (2025) 39–53.
[9] S. Yang, M. P. Wan, Machine-learning-based model predictive control with instantaneous linearization – a case study on an air-conditioning and mechanical ventilation system, Applied Energy 306 (2022) 118041.
[10] Široký, J., Oldewurtel, F., Cigler, J., & Prívara, S. (2011). Experimental analysis of model predictive control for an energy efficient building heating system. *Applied energy*, *88*(9), 3079-3087.
[11] Y. Ma, F. Borrelli, B. Hencey, B. Coffey, S. Bengea, P. Haves, Model predictive control for the operation of building cooling systems, IEEE Transactions on control systems technology 20 (3) (2011) 796–803.
[12] X. Pang, C. Duarte, P. Haves, F. Chuang, Testing and demonstration of model predictive control applied to a radiant slab cooling system in a building test facility, Energy and Buildings 172 (2018) 432–441.
[13] S. Yang, W. Chen, M. P. Wan, A machine-learning-based event-triggered model predictive control for building energy management, Building and Environment 233 (2023) 110101.
[14] A. Beltran, A. E. Cerpa, Optimal hvac building control with occupancy prediction, in: Proceedings of the 1st ACM conference on embedded systems for energy-efficient buildings, 2014, pp. 168–171.