




Title	Power-constrained VRF system optimization using symbolic regression for multiple zones environment
Author(s)	Theint Thu, Theint; Kato, Kenshiro; Zhao, Dafang et al.
Citation	Energy and Buildings. 2025, 347, p. 116264
Version Type	VoR
URL	https://hdl.handle.net/11094/102866
rights	This article is licensed under a Creative Commons Attribution 4.0 International License.
Note	

The University of Osaka Institutional Knowledge Archive : OUKA

<https://ir.library.osaka-u.ac.jp/>

The University of Osaka



Power-constrained VRF system optimization using symbolic regression for multiple zones environment

Theint Theint Thu ^{*}, Kenshiro Kato, Dafang Zhao, Hiroki Nishikawa, Ittetsu Taniguchi , Takao Onoye

The University of Osaka, 1-5 Yamadaoka, Suita, Osaka, Japan

ARTICLE INFO

Keywords:

Model predictive control
Symbolic regression
Variable refrigerant flow
Thermal comfort
Energy efficiency
Power constraints

ABSTRACT

The rapid growth in global energy consumption highlights the urgency of doubling energy efficiency improvements by 2030. Heating, ventilation, and air-conditioning (HVAC) systems, which account for nearly half of building energy use, represent a critical target for optimization. Conventional HVAC control strategies, however, often suffer from inefficient power allocation, high peak demand, and compromised thermal comfort, especially under dynamic occupancy and environmental conditions. Existing multi-zone control methods often overlook peak power constraints and are not designed to optimize energy use under variable occupancy conditions, resulting in suboptimal energy performance. This study proposes a symbolic regression-based model predictive control (MPC) framework to address these challenges. The framework optimizes energy consumption and thermal comfort for multi-zone variable refrigerant flow (VRF) systems while addressing peak power constraints to reduce energy costs and improve thermal comfort. The method is evaluated under three operating priorities, $\omega = 0.1, 0.5$, and 0.9 , across varying power constraints. Simulation results demonstrate that the proposed method consistently outperforms a decentralized MPC state-of-the-art (SOTA) baseline, achieving up to 16 % energy savings under a 30 % power constraint, with average temperature deviations (ATD) remaining within comfortable bounds ($< 2^\circ\text{C}$). Even under tight energy constraints, the framework maintains stable control performance, outperforming existing methods that fail to adequately manage peak loads. Compared to rule-based and model-based MPC approaches, the proposed method is more flexible and robust, as it does not require detailed system identification or extensive training data. These results highlight the method's potential as a scalable and energy-efficient solution for contributing to global energy efficiency goals.

1. Introduction

Doubling global energy efficiency by 2030 is critical to achieving environmental and economic sustainability targets [1]. Buildings, which account for nearly 40 % of global energy consumption, play a pivotal role in this effort [2]. Within buildings, heating, ventilation, and air conditioning (HVAC) systems are the largest energy consumers [3,4]. According to the International Energy Agency (IEA), HVAC systems account for a significant portion of total energy consumption in buildings, underscoring their critical role in improving energy efficiency (Fig. 1). This highlights the urgent need for more efficient HVAC solutions. Among advanced technologies, variable refrigerant flow (VRF) systems have emerged as a promising solution due to their flexibility and energy efficiency [5,6]. VRF systems enable precise temperature control in individual zones, improving occupant comfort while reducing energy consumption compared to traditional HVAC systems [7,8]. Furthermore,

their compatibility with smart control strategies makes them well-suited for sustainable building operations [9,10].

However, managing multi-zone VRF systems remains challenging. Lack of coordinated control among VRF systems can lead to energy inefficiencies, higher operational costs, and thermal comfort inconsistent across zones [11]. Without considering the interactions between different units and zones, systems may overcool or overheat in certain areas, which negatively impacts overall system performance [12]. Moreover, uncoordinated operation of VRF systems can easily increase peak load during peak hours, causing higher utility costs and reduced reliability of the power grid [13]. Conventional control approaches, such as bang-bang control and rule-based methods, often fail to handle the complexities of multi-zone VRF systems, where thermal loads and occupancy patterns vary over time [14,15].

Advanced control strategies have been widely explored to address these challenges, including fuzzy logic control [16], heuristic

^{*} Corresponding author.

E-mail address: theint.thu@ist.osaka-u.ac.jp (T.T. Thu).

<https://doi.org/10.1016/j.enbuild.2025.116264>

Received 18 February 2025; Received in revised form 29 July 2025; Accepted 6 August 2025

Available online 14 August 2025

0378-7788/© 2025 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).




Fig. 1. Global energy consumption by sector and the contribution of HVAC systems to building energy use. Source: International Energy Agency (IEA), 2022 [2,4].

optimization [17], reinforcement learning [18], and model predictive control (MPC) [19,20]. Among them, MPC is notable for its ability to predict future states and optimize HVAC operations in real time [21] and [22] demonstrated MPC's effectiveness in minimizing HVAC energy use. Recent works such as [23] and [24] further highlight MPC's potential for real-time coordination under operational constraints. Moreover, MPC can adjust in real-time to these changes, providing a more efficient and precise solution to manage multiple VRF systems.

While MPC can handle power constraints theoretically, its application to multi-zone VRF systems under lower peak power limits remains largely unexplored. Achieving the trade-off between energy efficiency and thermal comfort under power constraints is crucial in real-world applications. In particular, the dynamic and interdependent power consumption patterns of VRF units under such constraints require accurate modeling and control strategies.

To address these issues, this paper proposes an operation scheduling method for VRF systems that integrates symbolic regression (SR) with MPC for multi-zone VRF systems. The proposed method tackles two major issues: (1) accurate thermal modeling and (2) coordinated power allocation across zones under peak power constraints. SR is used to derive mathematical models of indoor thermal dynamics directly from data. Unlike traditional regression methods, SR automatically discovers nonlinear, interpretable mathematical expressions, offering a flexible and data-driven way to capture room thermal dynamics [25]. This capability is particularly valuable for modeling complex interactions between indoor temperature, outdoor conditions, occupancy, and HVAC operation. By enhancing predictive accuracy through SR, the proposed framework enables MPC to schedule energy use more efficiently while maintaining occupant comfort. The novelty of the approach lies in its ability to manage power allocation flexibility and improve energy efficiency. The main contributions of this paper are as follows:

1. Development of an operational scheduling approach that integrates symbolic regression with MPC to optimize energy consumption in multi-zone VRF systems under peak power constraints.
2. Comprehensive analysis of the impact of power constraints on energy consumption, peak power demand, and thermal comfort through simulations reflecting real-world VRF system scheduling.

The remainder of this paper is organized as follows: Section 2 introduces the proposed data-driven multi-zone VRF system management framework and formulation of the energy optimization problem for multi-zone VRF systems. Section 3 presents the simulation evaluation, and Section 4 presents the simulation results. Section 5 discusses the main findings, and finally, Section 6 concludes with suggestions for future research.

1.1. Related work

Optimizing HVAC systems has been a significant focus of research, driven by the need to balance energy efficiency and occupant comfort.

Early approaches mainly used rule-based or heuristic methods, which are simple to implement but lack flexibility and scalability. These methods often rely on fixed schedules or historical trends and struggle to adapt to real-time changes in occupancy or weather [26]. For example, setpoint controls based on average conditions may work in stable environments but perform poorly in dynamic, multi-zone settings. Paulo et al. addressed energy consumption during peak periods by designing a time-of-use pricing model for residential thermostats, showing the benefit of incorporating electricity pricing into HVAC control [27]. However, their method did not adapt to real-time variations in occupancy or indoor climate conditions. While effective for incorporating network constraints and pricing, their method did not explicitly model HVAC dynamics or support real-time operational control, limiting its responsiveness in practical HVAC applications.

Model predictive control (MPC) has emerged as a prominent solution for HVAC systems due to its ability to forecast future conditions and optimize control strategies accordingly. Drgoña et al. provided a comprehensive review on MPC applications in buildings, highlighting its benefits for balancing comfort and energy efficiency [28]. However, MPC can face scalability and computational issues, particularly in large systems [29]. To address these issues, researchers have proposed hierarchical and distributed MPC frameworks that divide buildings into smaller, manageable subsystems [30], improving scalability and coordination. Xie et al. implemented such a distributed MPC approach across university campus buildings, achieving notable reductions in energy consumption while maintaining thermal comfort [31]. Hybrid methods that combine MPC with simpler control rules have also been introduced to reduce computational burden while keeping good performance.

Recently, data-driven techniques including supervised learning and reinforcement learning (RL), have gained significant attention in HVAC optimization. Zhang et al. reviewed machine learning and reinforcement learning approaches integrated with MPC, highlighting the growing use of learning-based control in complex environments [32]. Yang et al. developed a learning-based adaptive MPC method to deal with uncertainty in building dynamics [33], while Kathirgamanathan et al. showed that predictive control can improve energy flexibility and manage peak loads [34]. However, these methods require extensive training data and frequent model updates, leading to implementation challenges in dynamic environments.

To address these challenges, recent studies have focused on incorporating real-time data and online learning techniques to reduce the need for manual model retraining. Zhao et al. demonstrated significant energy cost reductions and peak shaving through the integration of symbolic regression into MPC-based energy management frameworks in on-site experiments [35]. Symbolic regression is useful for capturing HVAC behavior as it discovers interpretable models directly from data. While previous studies have explored various optimization strategies, few studies have explicitly addressed peak power constraints. High peak power demand contributes to increased operational costs and grid instability, emphasizing the need for constraint-aware control strategies.

Table 1
Comparison of existing approaches and the proposed method.

Approach	Ref.	Control Objective(s)	Pros	Cons	Gap Filled by Proposed Method
Rule-based / Heuristic	[16,17,27]	Energy efficiency, peak demand reduction	Simple implementation, fast response, low computational cost	Static logic, lacks real-time adaptability to occupancy/weather	Introduces dynamic, responsive control for real-time multi-zone VRF coordination
Model-based MPC	[19–24,28]	Energy and thermal comfort optimization	Predictive, constraint handling, suitable for multi-zone control	High model complexity, computational burden, scalability issues	Uses SR to reduce model complexity and enable real-time scalability
Data-driven MPC	[33,34]	Peak power reduction, energy flexibility	Adaptive to system behavior, learns from data, flexible	Requires large datasets, retraining, generalizability issues	Symbolic models reduce need for retraining and work with limited data
Reinforcement Learning (RL)	[18,32]	Energy efficiency, thermal comfort optimization	Learns optimal policies, adaptive to uncertainty, no explicit model needed	Requires extensive training and exploration; may lack interpretability	Proposed method offers interpretable, reliable alternative with lower data dependency
SR + MPC (Single-zone)	[35]	Energy cost reduction, peak shaving	Interpretable, accurate nonlinear models, adaptable	Single-zone only; lacks coordination, increases energy cost in multi-zone systems	Extends symbolic MPC to multi-zone VRF with coordinated control under power limits
Our Method (SR + MPC for Multi-zone VRF)	-	Peak power optimization, real-time response	Accurate, interpretable, flexible control of shared VRF systems	Symbolic models require careful selection and tuning	First scalable SR + MPC method with coordinated real-time multi-zone VRF scheduling under power constraints

Incorporating peak power constraints into HVAC control is becoming increasingly important for reducing energy costs and strain on the grid. Existing methods often fail to manage peak loads effectively and tend to overlook the impacts of peak demand on electricity costs, peak demand, and grid stability in multi-zone buildings. A detailed comparison of related approaches and the proposed method is summarized in Table 1.

Our method integrates symbolic regression with MPC in a multi-zone VRF system, incorporating peak power constraints to enable interpretable, scalable, and efficient control. Symbolic Regression (SR) is chosen in this study because it generates interpretable and compact equations that can be directly applied in model predictive control (MPC). Theoretically, SR represents a data-driven technique that does not require prior domain knowledge to construct an accurate thermal dynamics model. Practically, SR requires only readily available building operational data such as indoor temperature, outdoor temperature, and HVAC status, enabling fast deployment without the need for a detailed physics-based model. This makes it highly scalable across zones with diverse occupancy patterns and usage behaviors. Leprince et al. [36] revealed the effectiveness of SR in building energy modeling by accurately predicting temperature variations under diverse environmental conditions. Further research has applied SR to predict energy usage and thermal dynamics in various HVAC optimization scenarios [37,38].

Traditional regression techniques, such as linear and polynomial regression, are commonly used in MPC frameworks due to their simplicity and interpretability. However, these methods assume fixed functional forms and struggle to model the complex, nonlinear, and dynamic behavior typical of HVAC systems. In contrast, symbolic regression (SR) offers a data-driven approach that simultaneously discovers both the structure and parameters of governing equations without assuming a predefined model form. This flexibility allows SR to capture complex thermal dynamics and interactions within multi-zone systems more effectively than traditional regression. Moreover, compared to black-box models, SR provides clear expressions that describe how indoor temperature changes with inputs like outdoor temperature and power usage. There are three main reasons why SR fits our setting:

- **Easy integration into MPC:** SR produces mathematical expressions that are continuous and differentiable, making them suitable for use in optimization solvers. This is a key advantage over models like LSTM, which often require additional steps such as numerical differentiation or surrogate modeling to be incorporated into control frameworks [39].

- **Interpretability:** The output of SR provides clear and interpretable relationships between variables such as power usage, outdoor temperature, and indoor conditions. This transparency aids in both system understanding and model validation. In contrast, models like Random Forests, while potentially accurate, lack this level of interpretability [40].
- **Computational efficiency:** SR models are lightweight and quick to evaluate, which is crucial for real-time applications. In comparison, models like Random Forests or deep neural networks involve complex architectures that are computationally intensive and less straightforward to apply in optimization problems [41].

In this study, we prioritize generating low-complexity SR models to balance prediction accuracy with optimization compatibility. While this may slightly reduce predictive performance, previous work [42] has shown SR remains competitive with other ML models, while offering significant advantages in transparency and control integration.

2. Data-driven multi-zone VRF system management

This study employs a coordinated, data-driven multi-zone optimization framework for variable refrigerant flow (VRF) systems based on [43]. The main objective is to minimize the total energy consumption without compromising the desired comfort levels. The proposed method leverages model predictive control (MPC) to optimize the multi-zone VRF systems scheduling, as illustrated in Fig. 2. The process begins by collecting real-time environmental data, such as outdoor and indoor temperatures, from sensors. The collected data is combined with thermal comfort and a symbolic regression model, which are integrated into the MPC framework to improve prediction accuracy. During scheduling periods, the MPC controller calculates optimal temperature set points for each room, ensuring consistent and efficient operation. To improve energy efficiency, this study simply incorporates overall peak power constraints into the VRF scheduling process. This ensures that the system operates within specified peak power limits while maintaining thermal comfort. We formulated a mixed-integer linear programming (MILP) problem integrated with MPC for coordinated load scheduling.

2.1. Symbolic regression-based temperature prediction model

Symbolic Regression (SR) employs a tree-like structure, where nodes represent mathematical operations (e.g., addition, multiplication) and leaves correspond to input variables. The tree structure is optimized




Fig. 2. Overview of the coordinated VRF scheduling approach utilizing model predictive control (MPC) and symbolic regression-based temperature prediction.




Fig. 3. Binary tree representation in the symbolic regression (SR) model.

through an iterative process that minimizes the error between predicted and actual temperatures, ensuring the model adapts dynamically to environmental and operational changes. As shown in Fig. 3, SR constructs predictive equations by combining variables, coefficients, and operators to define the function f^{SR} . The function maps the explanatory variables X to the target variable Y , represented as:

$$Y = f^{SR}(X). \quad (1)$$

In the context of multi-zone VRF systems, the input variables $X_{i,t}$ are selected to capture the key factors influencing temperature dynamics. These include room temperature ($T_{i,t}^{in}$), which reflects the current thermal state; outdoor temperature (T_t^{out}), which represents the external condition driving heat exchange with the indoor space; VRF power consumption ($D_{i,t}$), which indicates system load; and VRF capacity ($Q_{i,t}$), which determines the system's cooling or heating capability. Altogether, these variables enable SR to model the complex relations that derive temperature variations in VRF systems. The variable $X_{i,t}$ can therefore be expressed as:

$$X_{i,t} = \left\{ \begin{array}{c} T_{i,t}^{in} \\ T_t^{out} \\ D_{i,t} \\ Q_{i,t} \end{array} \right\}. \quad (2)$$




Fig. 4. Overview of the symbolic regression (SR) training process for temperature prediction.

To accurately determine the prediction function, the SR model leverages both current and lagged explanatory variables. These variables include the current time step t , one-step behind $t - 1$, further lagged steps $t - 1 \cdot a$, $t - 2 \cdot a$, ..., $n \cdot a$ -step behind $t - n \cdot a$, where a is the lag interval and n is the number of lagged steps considered. By incorporating these lagged terms, the SR model can better capture the dynamics of VRF system behavior, leading to more accurate temperature predictions. Fig. 4 provides an overview of the SR training process for temperature prediction. The prediction function for the indoor temperature at time $t + 1$ is thus expressed as:

$$\hat{T}_{i,t+1}^{in} = f(X_{i,t}, X_{i,t-1}, X_{i,t-1 \cdot a}, \dots, X_{i,t-1 \cdot n \cdot a}), \quad (3)$$

where $\hat{T}_{i,t+1}^{in}$ represents the predicted indoor temperature, and $X_{i,t}$ includes the explanatory variables at time t .

To train the SR model, a dataset containing N samples of explanatory and objective variables is constructed. Each sample is generated by shifting the time-series data step by step, ensuring that the model captures temporal dependencies. For each sample, the current time step is defined by the most recent explanatory variables for each room, where the current time steps for Sample 1, Sample 2, ..., Sample N correspond to $X_{i,t}$, $X_{i,t+1}$, $X_{i,t+N-1}$, respectively. The objective variable for each sample is the one-step-ahead indoor temperature, $T_{i,t+1}^{in}$. During training, the SR

Table 2

List of variables used in the model.

Variable	Description	Unit
$T_{i,t}^{in}$	Indoor temperature in room i at time t	°C
$T_{i,t}^{out}$	Outdoor temperature at time t	°C
$D_{i,t}$	Power consumption of VRF system in room i	kW
$Q_{i,t}$	Cooling/heating capacity of VRF system in room i	kW
$O_{i,t}$	Occupancy status of room i at time t	Binary (0 or 1)
S_t	Total energy consumption at time t	kWh
ϵ_t	Electricity price at time t	currency/kWh
$Stc_{i,t}$	Slack variable for comfort constraint	°C
J^{cost}	Electricity cost function	currency unit
$J^{comfort}$	Thermal comfort cost function	°C ²

model iteratively optimizes its tree structure to minimize the error between predicted and actual temperatures, ensuring robust performance across diverse operating conditions.

2.2. Formulation of multi-zoom VRF system optimization problem

The VRF optimization problem includes temperature prediction and energy cost calculation under power constraints. For indoor thermal dynamics, the SR-derived building thermal model denoted as $f^{SR}(X)$, predicts future indoor temperatures as:

$$T_{i,t+1}^{in} = f^{SR}(X_{i,t}, \dots, X_{i,t-N}), \forall t. \quad (4)$$

where $T_{i,t+1}^{in}$ represents the predicted indoor temperature at time $t+1$, and $X_{i,t}, \dots, X_{i,t-N}$ are explanatory variables capturing current and past environmental and operational conditions. The variables used in the equations, along with their units and descriptions are shown in Table 2. This approach enables reliable temperature prediction based on data, without requiring physical parameter estimation.

To compute the energy cost, the operational status of the VRF system is represented by the binary variable $z_{i,t}$:

$$z_{i,t} = \begin{cases} 1, & \text{if the VRF system is ON at time } t, \\ 0, & \text{if the VRF system is OFF at time } t. \end{cases} \quad (5)$$

A binary variable $z_{i,t}$ is introduced to represent the ON/OFF operational status of the VRF unit, enabling discrete operation control. The power consumption constraint for room i at time t is:

$$z_{i,t} \cdot D_i^{lowest} \leq D_{i,t} \leq z_{i,t} \cdot D_i^{rated}, \quad (6)$$

where D_i^{lowest} and D_i^{rated} denote the minimum and rated power consumption of the VRF system, respectively. This constraint ensures that the power consumption remains within the valid operating range only when the system is ON.

The total energy consumption S_t over a time slot is computed as:

$$S_t = \sum_{i=1}^N \sum_{t=1}^T D_{i,t} \cdot \frac{\Delta t}{3600}, \forall t, \quad (7)$$

where $D_{i,t}$ is the power consumption [kW] of room i at time t , and Δt is the duration of each time step in seconds. The total energy usage is accumulated across all rooms and time steps to track the system's overall consumption profile. The associated electricity cost, J^{cost} is:

$$J^{cost} = \sum_{t=1}^T \epsilon_t \cdot S_t, \quad (8)$$

where ϵ_t is the electricity unit price at time t . This cost term is incorporated with time-varying electricity price to represent the energy consumption during the scheduling period.

To limit total energy demand, a peak power constraint is imposed:

$$\sum_{i=1}^N D_{i,t} \leq D^{max}, \forall t, \quad (9)$$

ensuring that the combined power consumption of all VRF units does not exceed the maximum allowable demand D^{max} . The peak power constraint limits total peak power demand to avoid peak demand penalties.

To enhance occupant comfort, the thermal comfort term $J^{comfort}$ is defined as the squared error between the actual $T_{i,t}^{in}$ and the target temperature $T_{i,t}^{ref}$:

$$J^{comfort} = \sum_{i=1}^N \sum_{t=1}^T O_{i,t} \cdot (T_{i,t}^{in} - T_{i,t}^{ref})^2, \quad (10)$$

where $O_{i,t}$ is a binary variable indicating room occupancy, which equals one when the room is occupied and zero when the room is unoccupied. The thermal comfort penalty is defined as the squared error between the actual and reference temperatures, activated only during occupancy to reflect realistic usage.

$$T_{i,t}^{lower} - Stc_{i,t} \leq T_{i,t}^{in} \leq T_{i,t}^{upper} + Stc_{i,t}, \quad Stc_{i,t} \geq 0 \quad (11)$$

where $T_{i,t}^{lower}$ and $T_{i,t}^{upper}$ are the lower and upper bounds for indoor temperatures, and $Stc_{i,t}$ are the slack variables to accommodate temperature deviations. This allows the indoor temperature $T_{i,t}^{in}$ to slightly deviate from its bounds when necessary. The slack ensures that a solution always exists, avoiding infeasibility in the optimization problem. In this multi-zone VRF scheduling framework, the objective function is designed to jointly minimize the total energy cost, occupant discomfort, and constraint violation penalties. To ensure each term contributes fairly to the optimization, the energy cost and thermal discomfort values are normalized based on their respective maximum and minimum values across the scheduling horizon.

The energy cost term J^{cost} is normalized using rated power limits, ensuring fairness in contribution. Here, $D_{i,t}$ denotes the power consumption of zone i at time t , while $D_{i,t}^{rated}$ and $D_{i,t}^{lowest}$ represent the maximum and minimum power thresholds, respectively.

$$J_{norm}^{cost} = \frac{\sum_{t=1}^T \sum_{i=1}^N D_{i,t}}{\sum_{t=1}^T \sum_{i=1}^N \max(D_{i,t}^{rated} - D_{i,t}, D_{i,t} - D_{i,t}^{lowest})}, \quad (12)$$

The thermal comfort term $J^{comfort}$ quantifies temperature deviation from the desired setpoint $T_{i,t}^{ref}$, weighted by the occupancy indicator $O_{i,t}$. Here, $T_{i,t}^{in}$ is the indoor temperature of zone i at time t , and $T_{i,t}^{upper}$, $T_{i,t}^{lower}$ are the comfort bounds.

$$J_{norm}^{comfort} = \frac{\sum_{t=1}^T \sum_{i=1}^N O_{i,t} \cdot (T_{i,t}^{in} - T_{i,t}^{ref})^2}{\sum_{t=1}^T \sum_{i=1}^N O_{i,t} \cdot (\max(T_{i,t}^{upper} - T_{i,t}^{ref}, T_{i,t}^{ref} - T_{i,t}^{lower}))^2}, \quad (13)$$

This normalization process ensures that both terms contribute comparably to the objective function, avoiding disproportionate influence from one metric over the other due to differences in magnitude or measurement units.

The objective is to minimize both energy cost and thermal comfort deviations, expressed as:

$$\begin{aligned} \text{minimize: } & \omega \cdot \sum_{t=1}^T J_t^{cost} + (1 - \omega) \cdot \sum_{i=1}^N \sum_{t=1}^T J_{i,t}^{comfort} \\ & + P_e \cdot \sum_{i=1}^N \sum_{t=1}^T Stc_{i,t} \end{aligned}$$

subject to (4) – (7), (9), (11)

given $O_{i,t}, \forall i, t,$

decision variables: $D_{i,t}, \forall i, t,$

slack variables: $Stc_{i,t}, \forall i, t,$

The overall objective minimizes a weighted sum of electricity cost, thermal comfort violations, and slack penalties. The cost term J^{cost} is subject to Eqs. (6) and (9), while the comfort term $J^{comfort}$ is subject to Eq. (11). The weight ω allows the trade-off between cost efficiency

and occupant satisfaction. To determine appropriate trade-off weights ω , we analyzed system performance by simulating under various ω values ranging from 0 to 1. This allowed us to observe how the system's energy consumption, average temperature deviation, and peak power demand varied with different priority settings. Based on this analysis, we selected three representative values of $\omega = 0.1, 0.5$, and 0.9 , which reflect distinct operational priorities: comfort-oriented, balanced, and cost-oriented. In this optimization, the decision variables include power consumption $D_{i,t}$ for each room i and each time step t . To address potential infeasibility caused by strict temperature bounds under power constraints, we introduce slack variables $Slc_{i,t}$. These variables allow for the relaxation of thermal comfort constraints and are penalized in the objective function to limit excessive violations. In contrast, power operation and peak power constraints are modeled as hard constraints to strictly enforce system limitations.

The system dynamics are parameterized using a symbolic regression model, which predicts indoor temperature, outdoor temperature, and power consumption based on historical data. These predictions are then used in the MPC optimization, where power constraints (Eq. (6)) are enforced and thermal comfort constraints are relaxed through the slack variable (Eq. (11)). These variables interact with the constraints to balance energy efficiency and occupant comfort while maintaining feasibility. Moreover, the operation of the VRF system involves several key variables. The ON/OFF status and rated power are treated as manipulated variables within the MPC framework. Indoor temperature, energy consumption, and thermal comfort deviations are considered controlled variables. External factors such as outdoor temperature, occupancy, and electricity prices serve as disturbance variables that influence system behavior. Real-time sensor data, along with the predicted indoor temperatures from symbolic regression, provide feedback inputs, enabling dynamic control actions.

For symbolic regression training, we used Optuna [44], an open-source framework for hyperparameter optimization, to enhance model accuracy. For the MPC optimization, the CPLEX solver [45] was employed, relying on its internal parameter tuning.

3. Simulation evaluation

3.1. Simulation setup

We conducted simulations to compare the performance of the proposed (coordinated VRF) approach with the baseline (state-of-the-art, uncoordinated VRF) approach. We evaluate the impact of peak power constraints on total energy consumption, peak power demand, and thermal comfort. These approaches are referred to as “SOTA” and “Ours” throughout this section. The simulation is based on a real-world scenario, including office rooms in an educational building at the University of Osaka, Japan. Fig. 5 illustrates the room layout for four rooms with different cooling capacities, varying sizes, and occupancy patterns.




Fig. 5. Floor plan representing the room layout for the multi-zone VRF simulation.

Table 3

VRF indoor unit capacities and occupancy schedules for each room.

	Room 1	Room 2	Room 3	Room 4
Capacity [kW]	5.6	2.8	5.6	9.0
Area [m ²]	32	22	27	178
Occupied [hrs]	10:00-11:00	9:00-16:00	9:30-20:30	8:00-18:00

In the VRF system setup, summarized in Table 3, all rooms are connected to a central outdoor unit with a rated power of 8.93 kW. During occupied hours, we set a target temperature of 26 °C, with an acceptable fluctuation range between 24 °C and 28 °C to ensure occupant comfort while allowing flexibility for energy optimization.

In this study, the prediction and control horizon are set to 24 h with a time step of 15 min. This setting ensures that temperature set points for all rooms are updated frequently, allowing the system to respond dynamically to changes in room occupancy and outdoor temperature. Key metrics in this study are total energy consumption (kWh), peak power demand (kW), and thermal comfort (°C, measured as the average deviation from the target temperature across the rooms).

To evaluate the impact of peak power constraints, we defined the peak power constraint settings for SOTA and Ours methods as follows:

- Peak Power Constraint (SOTA): Each VRF system operates independently, and applies individual power constraints to each indoor unit, as illustrated in Fig. 6(a).
- Peak Power Constraint (Ours): Our approach applies a global constraint to coordinate power consumption across all units, as shown in Fig. 6(b).

Moreover, we introduced weight coefficients ω to adjust the trade-off between energy cost and thermal comfort. In this simulation, we employed three different ω values to represent different operational priorities:

1. $\omega = 0.1$: which prioritizes thermal comfort to minimize temperature deviation from the target.
2. $\omega = 0.5$: which balances energy cost and thermal comfort equally.
3. $\omega = 0.9$: which prioritizes cost reduction, allowing larger temperature deviations to minimize energy consumption.

In this study, we analyzed ten distinct power constraint settings, ranging from 10 % to 100 % of the maximum peak power. These settings represent the effect of varying peak power on energy consumption, thermal comfort, and the system's performance under both restricted and relaxed peak power. The peak power constraint that applies to SOTA approach is defined by the following equation:

$$D_i \leq D_i^{max} \cdot R, \quad (14)$$

where D_i and D_i^{max} denote the individual power demand and the maximum rated power of the i th indoor unit. R is the percentage rate of the power constraint (10%–100%) in this setting. In this simulation, the rated power of each indoor unit is unknown, and the individual maximum rated power is determined by its capacity ratio relative to the total rated power:

$$D_i^{max} = D^{max} \cdot (Q_i / \sum_{i=1}^N Q_i), \quad (15)$$

where Q_i and $\sum_{i=1}^N Q_i$ are the individual and total capacity of the i th indoor unit. This allocation method ensures that the power constraint for each unit reflects its proportional share of the outdoor unit's capacity, enabling effective power distribution across the system.

In contrast, the peak power constraint for the proposed approach is defined by the following equation:

$$\sum_{i=1}^N D_i \leq D^{max} \cdot R, \quad (16)$$




Fig. 6. Comparison of (a) baseline approach and (b) proposed approach for VRF system optimization.




Fig. 7. Measured and predicted temperature dynamics from SR training for each room: (a) Room 1, (b) Room 2, (c) Room 3, and (d) Room 4.

where $\sum_{i=1}^N D_i$ represents the total power demand of all VRF units, D^{max} is the total allowable peak power. This constraint allows the VRF system to manage power demand collectively across rooms.

3.2. Results of data training using the SR model

The Symbolic Regression (SR) model was trained using actual data collected from the operation of four office rooms in an educational build-

ing at the University of Osaka, Japan. The data, including outdoor temperature, indoor temperature (measured by sensors), and power consumption, was gathered over the period from July 13 to August 15, 2024. The dataset was divided into three phases: training (July 13–17), validation (August 1–5), and testing (August 6–10), with each phase consisting of 5 days of data. This real-world dataset serves as the foundation for the simulation conducted in this study. We implemented this model using the PySR library [46] and to control model complexity and

Table 4
Parameters of symbolic regression using PySR.

Parameter	Description	Value
Loss	Loss function used during regression	L2Distance()
Niterations	Number of iterations of the algorithm to run	120
Populations	Number of populations running in parallel	100
Population_size	Number of individuals in each population	100
Complexity_of_variables	Complexity assigned to each variable	2
Constraints	Max size constraints on arguments of operators	‘‘*’’: (2, 1)

preserve physical interpretability, we limited the allowed operator set to subtraction and multiplication ([‘−’, ‘*’]). Key hyperparameters for training, including the number of iterations, population size, and complexity constraints, are summarized in Table 4. These parameters were selected based on empirical testing to ensure a balance between model accuracy and interpretability. The temperature prediction functions based on three key values: the indoor temperature ($T_{i,t}^{\text{in}}$), outdoor temperature (T_t^{out}), and power consumption ($D_{i,t}$) over the 15 min intervals.

We evaluate the accuracy of the generated equations by comparing the actual indoor temperatures with the predicted values, using mean squared error (MSE), the coefficient of determination R^2 , mean absolute error (MAE) and mean absolute percentage error (MAPE) as performance metrics. The testing results for each room are summarized in Table 5. The symbolic regression (SR) models demonstrated strong predictive performance across Rooms 1–3, with high R^2 values exceeding 0.99 and low error metrics (MSE, MAE, and MAPE) across training, validation, and testing datasets (shown in Table 5). These results indicate that the SR models effectively captured the thermal dynamics in these rooms, likely due to their relatively stable occupancy patterns and less variable environmental conditions.

However, the model performance for Room 4 was notably lower, with a significant drop in R^2 (e.g., 0.0922 in validation) and increased prediction errors (shown in Fig. 7). This discrepancy can be attributed to several factors. First, Room 4 is a large common space (178 m²) characterized by irregular and diverse usage patterns, which introduce non-linear and unpredictable thermal behavior. Second, as a walking and gathering area, the room is subject to frequent disturbances such as door openings, variable ventilation, and unmeasured internal heat gains factors not explicitly captured in the current modeling framework. Lastly, while SR offers compact and interpretable equations, its ability to capture the complexity of highly dynamic environments is limited. Future work will focus on enhancing model accuracy for such spaces by incorporating additional contextual variables and exploring hybrid modeling approaches that can better accommodate complex, high-variance conditions.

The model’s prediction functions for Rooms 1–4 are represented by Eqs. (17)–(20), respectively.

$$T_{t+1}^{\text{in}} = 0.9622545 \cdot T_t^{\text{in}} + 0.018169424 \cdot T_t^{\text{out}} + 0.67211 - 0.28949875 \cdot D_t. \quad (17)$$

$$T_{t+1}^{\text{in}} = 0.94959116 \cdot T_t^{\text{in}} + 0.0127858445 \cdot T_t^{\text{out}} + 1.1754573 - 0.25755432 \cdot D_t. \quad (18)$$

$$T_{t+1}^{\text{in}} = 0.9294086 \cdot T_t^{\text{in}} + 9.4946474 \times 10^{-5} \cdot T_t^{\text{out}} + 2.058031 - 0.36269522 \cdot D_t. \quad (19)$$

$$T_{t+1}^{\text{in}} = 0.92937356 \cdot T_t^{\text{in}} + 7.00316 \times 10^{-5} \cdot T_t^{\text{out}} + 2.0596 - 0.36220685 \cdot D_t. \quad (20)$$

Table 5
Model performance metrics for each room across training, validation, and testing datasets.

Room	Phase	MSE	R^2	MAE	MAPE
Room 1	Training	0.0044	0.9937	0.0488	0.1631
	Validation	0.0051	0.9936	0.0488	0.1899
	Testing	0.0058	0.9951	0.0488	0.2021
Room 2	Training	0.0095	0.9888	0.0769	0.2738
	Validation	0.0056	0.9957	0.0769	0.1910
	Testing	0.0065	0.9950	0.0769	0.2020
Room 3	Training	0.0134	0.9745	0.1062	0.3913
	Validation	0.0068	0.9963	0.1062	0.2404
	Testing	0.0059	0.9933	0.1062	0.2030
Room 4	Training	0.0435	0.8490	0.0866	0.3086
	Validation	0.1678	0.0922	0.0866	1.0485
	Testing	0.0736	0.5211	0.0866	0.6708




Fig. 8. Power consumption in each room on July 13, 2024.

The maximum rated power of the outdoor unit in this simulation is 8.93 kW. However, the actual rated power during the winter reaches approximately 30 % of the maximum power, which is 2.43 kW, as shown in the data from July 13, 2024 (Fig. 8). Although the proposed method could theoretically be proven to consume up to 100 % of the rated power, it practically operates at only 30 % under normal conditions. To reflect real-world power settings, we set 30 % of the rated power (2.68 kW) as the maximum peak power constraint in our simulation. To evaluate the effectiveness of the power constraint on temperature dynamics and energy costs, we varied the power constraint from 10 % (0.27 kW) to 100 % (2.68 kW) of the maximum rated power and adjusted the ω settings. This setting ensured that the system maintained the target temperature of 26 °C while minimizing energy consumption under varying power constraints.

4. Simulation result

This section presents comparison results of the proposed coordinated VRF system (*Ours*) with the baseline uncoordinated VRF system (*SOTA*) across ten power constraint levels of 2.68 kW. We implemented the simulation by three operation scenarios: $\omega = 0.1$, $\omega = 0.5$, and $\omega = 0.9$, representing thermal comfort priority, a balanced approach between cost and comfort, and cost priority, respectively and the results are discussed in different sections. Then, we evaluated the performance by three key metrics: (1) total energy consumption, (2) peak power demand, and (3) average temperature deviation (ATD) from the target temperature of 26 °C.

4.1. Comfort-oriented, $\omega = 0.1$

In this scenario, the primary objective is to minimize temperature deviation from the target setpoint while maintaining energy efficiency. Table 6 and Fig. 9 summarize the performance of the baseline method (*SOTA*) and the proposed optimization strategy across a range of power

Table 6Comparison of baseline and proposed method at $\omega = 0.1$ (comfort priority).


Power constraint (%)	Method	Total peak power [kW]	Energy consumption [kWh]	ATD [$^{\circ}\text{C}$]
10 % (0.27 kW)	SOTA	0.25	4.09	3.65
20 % (0.54 kW)	SOTA	0.50	8.19	3.32
30 % (0.80 kW)	SOTA	0.80	9.60	0.32
40 % (1.07 kW)	SOTA	1.07	8.86	0.15
50 % (1.34 kW)	SOTA	1.34	8.44	0.13
60 % (1.61 kW)	SOTA	1.61	8.10	0.13
70 % (1.88 kW)	SOTA	1.88	8.01	0.13
80 % (2.14 kW)	SOTA	2.13	7.95	0.13
90 % (2.41 kW)	SOTA	2.41	7.87	0.13
100 % (2.68 kW)	SOTA	2.65	7.82	0.13
10 % (0.27 kW)	Ours	0.27	5.23	2.31
20 % (0.54 kW)	Ours	0.54	9.33	0.70
30 % (0.80 kW)	Ours	0.80	8.64	0.27
40 % (1.07 kW)	Ours	1.07	8.18	0.21
50 % (1.34 kW)	Ours	1.34	7.96	0.18
60 % (1.61 kW)	Ours	1.61	7.85	0.17
70 % (1.88 kW)	Ours	1.88	7.78	0.16
80 % (2.14 kW)	Ours	2.14	7.76	0.15
90 % (2.41 kW)	Ours	2.41	7.75	0.15
100 % (2.68 kW)	Ours	2.68	7.75	0.14
Average (SOTA)		1.46	7.89	0.82
Average (Ours)		1.47	7.82	0.44

constraints. As shown in Table 6, the proposed method consistently achieves superior thermal comfort, particularly under more stringent power constraints. For example, at a 30 % power constraint, *Ours* attains an average temperature deviation (ATD) of 0.27 $^{\circ}\text{C}$, outperforming *SOTA*, which records an ATD of 0.32 $^{\circ}\text{C}$. Even under more relaxed constraints, such as 60 % and 100 % the proposed method achieves comparable performance, with ATDs of 0.17 $^{\circ}\text{C}$ and 0.14 $^{\circ}\text{C}$, respectively, closely matching the baseline's 0.13 $^{\circ}\text{C}$ in both cases. On average, across all constraint levels, the proposed method yields an ATD of 0.44 $^{\circ}\text{C}$, nearly half that of *SOTA*'s 0.82 $^{\circ}\text{C}$, indicating improved thermal comfort.


In terms of energy consumption, the proposed method consistently uses less energy than the baseline. Under a 30 % power constraint,

Ours consumes 8.64 kWh which is approximately 10 % less than *SOTA*'s 9.60 kWh. Similar efficiency gains are observed at other levels; at 60 %, energy use is 7.85 kWh compared to 8.10 kWh under *SOTA*, and even at full power availability (100 %), the proposed method consumes slightly less energy, 7.75 kWh versus 7.82 kWh. Both methods adhere to their respective peak power constraints. At 30 %, the peak power reached is 0.80 kW for both methods. At 100 %, the proposed method fully utilizes the available capacity (2.68 kW), slightly more than *SOTA* (2.65 kW), while maintaining greater efficiency and comfort.


To enhance the interpretability of temperature control behavior, shaded regions were added to the temperature plots (Fig. 9(a) and (b)) to indicate room occupancy periods. These regions are color-coded to match each room's temperature trajectory: red for Room 1, orange for




(a) Baseline temperature (weight = 0.1)



(b) Proposed temperature (weight = 0.1)



(c) Baseline power (weight = 0.1)



(d) Proposed power (weight = 0.1)

Fig. 9. Comparison of baseline and proposed methods at weight = 0.1 (comfort priority): (a) baseline temperature, (b) proposed temperature, (c) baseline power, and (d) proposed power.

Table 7Comparison of baseline and proposed method at $\omega = 0.5$ (trade-off between energy cost and comfort).

Power constraint (%)	Method	Total peak power [kW]	Energy consumption [kWh]	ATD [$^{\circ}\text{C}$]
10 % (0.27 kW)	SOTA	0.25	4.09	3.65
20 % (0.54 kW)	SOTA	0.50	8.19	3.32
30 % (0.80 kW)	SOTA	0.80	6.74	0.90
40 % (1.07 kW)	SOTA	1.07	6.43	0.81
50 % (1.34 kW)	SOTA	1.13	6.19	0.81
60 % (1.61 kW)	SOTA	1.61	5.99	0.81
70 % (1.88 kW)	SOTA	1.86	5.94	0.82
80 % (2.14 kW)	SOTA	2.10	5.92	0.82
90 % (2.41 kW)	SOTA	2.26	5.87	0.82
100 % (2.68 kW)	SOTA	2.50	5.83	0.82
10 % (0.27 kW)	Ours	0.27	4.89	2.36
20 % (0.54 kW)	Ours	0.54	5.52	1.39
30 % (0.80 kW)	Ours	0.80	5.69	1.13
40 % (1.07 kW)	Ours	1.07	5.68	1.06
50 % (1.34 kW)	Ours	1.34	5.67	1.03
60 % (1.61 kW)	Ours	1.61	5.70	0.99
70 % (1.88 kW)	Ours	1.88	5.70	0.98
80 % (2.14 kW)	Ours	2.14	5.69	0.98
90 % (2.41 kW)	Ours	2.41	5.69	0.97
100 % (2.68 kW)	Ours	2.68	5.69	0.97
Avg SOTA		1.41	6.12	1.36
Avg Ours		1.47	5.59	1.19

Room 2, green for Room 3, and blue for Room 4. A detailed comparison under a 30 % power constraint further illustrates the benefits of the proposed method. Fig. 9(c) shows the hourly power consumption of Room 1 to 4 using the baseline method. In this graph, solid lines represent power usage during occupied hours, while dotted lines indicate unoccupied periods. Under the baseline strategy, Room 1 initiates cooling as early as 00:30 to reach the target temperature by 10:00, resulting in considerable energy consumption during unoccupied hours.

Similarly, Room 2 begins cooling well ahead of its occupancy period and continues to lower the indoor temperature beyond the set-point by more than 1°C after 11:30 due to constrained cooling capacity. Rooms 3 and 4, which have longer occupied periods (9:30–20:30 and 8:00–18:00, respectively), maintain stable indoor temperatures but also experience extended cooling durations before occupancy begins. The baseline method treats each room independently, without accounting for shared power limitations or coordinated occupancy patterns, which leads to increased energy usage and reduced efficiency.

In contrast, the proposed method adopts a coordinated, occupancy-aware control strategy. As shown in Fig. 9(d), Room 1 initiates pre-cooling at 04:30, Room 2 at 06:15, Room 3 at 08:45, and Room 4 at 07:15—significantly later than the baseline approach. This scheduling aligns closely with each room's occupancy period, reducing unnecessary energy use during unoccupied hours while still ensuring thermal comfort. As a result, the proposed method balances comfort and energy efficiency, particularly under limited power availability.

4.2. Trade-off between cost and comfort, $\omega = 0.5$

The main objective of this simulation is to balance energy consumption and thermal comfort. A performance summary of both the proposed and baseline methods under this trade-off is presented in Table 7 and Fig. 10. The proposed method demonstrates considerable energy savings, especially under tighter power constraints. At a 20 % power constraint, *Ours* consumed 5.52 kWh, 32.6 % less than *SOTA*'s 8.19 kWh. This energy-saving trend persists as the power limit increases, with *Ours* consistently using less energy. At full capacity (100 %), *Ours* consumed 5.69 kWh, compared to *SOTA*'s 5.83 kWh—an overall reduction of 2.4 %.

In addition to improved energy efficiency, the proposed method also maintains better thermal comfort, particularly under constrained conditions. For instance, at 20 % power, *Ours* achieved an ATD of 1.39°C , significantly lower than *SOTA*'s 3.32°C , while also consuming less en-

ergy. This suggests better performance in managing indoor temperature when power is limited. As the power constraint relaxes, yet *Ours* consistently sustains comparable comfort levels above the 60 % constraint. At 100 %, *Ours* maintains an ATD of 0.97°C , while *SOTA* records 0.82°C . Both methods respect the specified peak power limits at all constraint levels. At 100 %, each reaches the maximum allowable peak of 2.68 kW without exceeding it. On average, the proposed method uses 5.59 kWh, while *SOTA* consumes 6.12 kWh, an 8.7 % overall reduction. Furthermore, the proposed method maintains a lower average ATD of 1.19°C , compared to *SOTA*'s 1.36°C , indicating better thermal management.

To further illustrate the effectiveness of both methods, we focus on the scenario with a 30 % power constraint. Fig. 10 provides a detailed hourly breakdown of energy usage in Rooms 1 through 4. Power consumption for each room is shown in red, orange, green, and blue lines (Fig. 10(c)), with solid lines indicating occupied periods, dotted lines indicating unoccupied hours and the shaded areas indicate the occupied hours. To meet the target temperature of 26°C , the *SOTA* method initiates early pre-cooling: Room 1 at 03:15 and Room 2 at 04:00. Rooms 3 and 4 start at 08:00 and 06:30, respectively, ensuring comfortable conditions at occupancy.

While this approach results in a lower ATD of 0.90°C , slightly better than *Ours*, it also leads to higher energy usage due to extended pre-conditioning. In contrast, the proposed method adjusts pre-cooling more responsively to actual occupancy times. Here, the proposed method dynamically updates control decisions at each time step in a coordinated manner, based on real-time inputs such as occupancy schedules and power constraints. This responsiveness is twofold: (1) it adapts to occupancy changes by minimizing unnecessary pre-cooling, and (2) it allocates limited power efficiently across zones to ensure comfort while reducing energy consumption. As depicted in Fig. 10(d), Room 1 begins pre-cooling at 07:00, Room 2 at 07:15, Room 3 at 09:00, and Room 4 at 07:15. This pre-cooling schedule avoids unnecessary conditioning during unoccupied hours, resulting in substantial energy savings while maintaining acceptable comfort levels.

4.3. Energy cost priority case, $\omega = 0.9$

The primary objective of this simulation is to prioritize energy cost savings, allowing for greater deviations from the target temperature. Both methods are evaluated under various power constraint levels. As shown in Table 8 and Fig. 11, the proposed method consistently reduces




Fig. 10. Comparison of baseline and proposed methods at weight = 0.5 (trade-off between energy cost and comfort): (a) baseline temperature, (b) proposed temperature, (c) baseline power, and (d) proposed power.

Table 8

Comparison of baseline and proposed method at $\omega = 0.9$ (energy cost priority).

Power constraint (%)	Method	Total peak power [kW]	Energy consumption [kWh]	ATD [°C]
10 % (0.27 kW)	SOTA	0.25	4.09	3.65
20 % (0.54 kW)	SOTA	0.50	8.19	3.32
30 % (0.80 kW)	SOTA	0.80	4.23	1.98
40 % (1.07 kW)	SOTA	1.02	4.10	1.96
50 % (1.34 kW)	SOTA	1.33	4.05	1.95
60 % (1.61 kW)	SOTA	1.61	4.02	1.94
70 % (1.88 kW)	SOTA	1.77	3.98	1.94
80 % (2.14 kW)	SOTA	2.07	3.97	1.94
90 % (2.41 kW)	SOTA	2.29	3.97	1.94
100 % (2.68 kW)	SOTA	2.52	3.97	1.94
10 % (0.27 kW)	Ours	0.27	4.66	2.43
20 % (0.54 kW)	Ours	0.54	4.31	2.00
30 % (0.80 kW)	Ours	0.80	4.08	2.00
40 % (1.07 kW)	Ours	1.07	4.01	1.99
50 % (1.34 kW)	Ours	1.34	3.98	1.99
60 % (1.61 kW)	Ours	1.61	3.97	2.00
70 % (1.88 kW)	Ours	1.88	3.96	2.00
80 % (2.14 kW)	Ours	2.14	3.96	2.00
90 % (2.41 kW)	Ours	2.41	3.96	2.00
100 % (2.68 kW)	Ours	2.52	3.96	2.00
Avg SOTA		1.42	4.46	2.26
Avg Ours		1.46	4.09	2.04

energy consumption while maintaining acceptable thermal comfort. At a 20 % power constraint, the proposed method (*Ours*) consumes 4.31 kWh, which is 47.4 % lower than the baseline method (*SOTA*) at 8.19 kWh. As the power constraint increases, energy consumption for both methods decreases; however, *Ours* consistently achieves lower consumption. At 100 % power constraint, both methods converge, consuming approximately the same energy: 3.96 kWh for *Ours* and 3.97 kWh for *SOTA*.

While *SOTA* achieves slightly lower ATD values (from 3.65 °C to 1.94 °C), the proposed method maintains deviations within 2 °C, which is acceptable in cost-prioritized scenarios. Additionally, both methods respect the maximum peak power limit of 2.68 kW. On average, the proposed method achieves 8.3 % lower energy consumption (4.09 kWh

vs. 4.46 kWh) and a slightly better ATD (2.04 °C vs. 2.26 °C) than the baseline, demonstrating a better trade-off between energy efficiency and thermal comfort in cost-sensitive settings.

To further illustrate performance, we highlight the case under a 30 % power constraint. Fig. 11 displays hourly power usage for Rooms 1–4. In *SOTA* (Fig. 11(c)), early pre-cooling starts at 06:45 for Rooms 1 and 2, 08:30 for Room 3, and 07:15 for Room 4. These long pre-cooling periods lead to higher energy use. In contrast, *Ours* (Fig. 11(d)) schedules energy usage, initiating pre-cooling closer to occupancy times: Room 1 at 07:30, Room 2 at 08:15, Room 3 at 09:15, and Room 4 at 07:45. This strategic timing avoids unnecessary energy usage while ensuring thermal comfort.




Fig. 11. Comparison of baseline and proposed methods at weight 0.9 (energy cost priority): (a) baseline temperature, (b) proposed temperature, (c) baseline power, and (d) proposed power.

Although *Ours* results in a slightly higher ATD (2.00°C) compared to *SOTA* (1.98°C) in the 30% constraint case, it achieves 4% greater energy savings. This small trade-off in thermal comfort is acceptable in cost-focused applications, confirming the effectiveness of the proposed method.

4.4. Comparison with state-of-the-art methods

To evaluate the performance of our proposed Symbolic Regression-based Model Predictive Control (SR-MPC) approach, we compared it with several established HVAC control strategies: rule-based control, model-based MPC, data-driven MPC, and reinforcement learning (RL)-based control. The comparison is summarized in Table 9, which outlines key characteristics such as energy savings, computational demands, interpretability, and real-time feasibility.

Rule-based control is one of the simplest control strategies and has been widely used in practice. It relies on predefined rules, making it highly interpretable and computationally efficient. However, the energy savings achieved by this method are limited due to its inability to adapt to varying building conditions or optimize performance based on real-time data. Rule-based control is effective in stable environments but struggles to maintain energy efficiency under dynamic conditions, which are common in modern HVAC systems [16,17,27].

Model-based MPC, on the other hand, uses mathematical models derived from first principles to optimize control decisions. While it provides moderate to high energy savings and high interpretability due to its reliance on physical models, it has notable limitations. The need for manual model calibration and system identification can be time-consuming and computationally intensive. Moreover, model-based MPC methods may not effectively capture the complex, nonlinear dynamics of real-world HVAC systems, reducing their effectiveness in some situations. Although it achieves reasonable real-time feasibility, the method typically requires adjustments for each specific system, which can hinder its adaptability [19,21,22].

Data-driven MPC, such as those using neural networks, is able to capture complex system dynamics and achieve high energy savings by learning from large datasets. These models, however, come at the cost

of high computational demands and low interpretability, as neural networks are often viewed as “black boxes.” The requirement for large amounts of training data and extensive tuning further increases computational costs. Additionally, while data-driven MPC can offer high energy savings, its real-time feasibility is limited by the time required for data processing and model updating, which is especially problematic in fast-changing environments [33,34].

Reinforcement learning (RL)-based control represents a more flexible approach that can learn and adapt to changing environments. It has the potential to deliver significant energy savings due to its ability to optimize actions over time. However, RL-based methods are computationally expensive and require extensive training, which can make them impractical for real-time control, especially in buildings with varying occupancy and environmental conditions. Moreover, RL approaches often lack interpretability, as the decision-making process is not easily understood by humans. These factors lead to lower real-time feasibility and generalizability compared to other methods [18,32].

In contrast, our SR-based MPC offers several advantages over the aforementioned methods. It combines the best features of traditional model-based MPC with the efficiency of data-driven methods, but with significantly lower computational demands. The SR-based approach generates closed-form expressions that are highly interpretable, which is a key benefit over black-box methods like neural networks. This interpretability makes it easier to diagnose issues and adjust the system if needed, offering transparency in decision-making. Additionally, SR-MPC is computationally efficient, making it ideal for real-time control applications where quick decision-making is crucial. Its ability to balance high energy savings with low computational requirements and real-time feasibility positions SR-MPC as a strong candidate for practical implementation in HVAC systems.

By focusing on the key attributes outlined in Table 9, we demonstrate that our SR-based MPC approach outperforms many traditional and modern methods in terms of interpretability, computational efficiency, and real-time feasibility while maintaining high energy savings. The comparison highlights how SR-MPC provides a practical, scalable solution that combines the benefits of both traditional and data-driven approaches while minimizing their drawbacks.

Table 9
Comparison of control methods in terms of key characteristics.

Method	Energy savings	Computational demand	Interpretability	Real-time feasibility
Rule-Based Control	Low	Low	High	High
Model-Based MPC	Medium to High	Medium	High	Medium
Data-Driven MPC (Neural Networks)	High	High	Low	Medium
Reinforcement Learning (RL)	High	Very High	Low	Low to Medium
Proposed SR-Based MPC	High	Low	High	High

5. Discussion

Different simulation results demonstrate that the proposed coordinated method effectively optimizes energy consumption while maintaining thermal comfort across various operating priorities. Our method consistently outperforms the baseline, achieving significant energy savings with minimal average temperature deviation (ATD). When $\omega = 0.1$ (comfort-oriented), *Ours* achieves a 10 % reduction in energy usage and improves ATD by 0.05°C under a 30 % power constraint. The first two scenarios show slightly higher ATD values with our proposed method under tight peak power limitations. This occurs because the decentralized *SOTA* method independently tracks each room's setpoint more aggressively, without coordination between rooms. As the power limitation is relaxed (e.g., 40 %–100 %), *SOTA* units can follow their temperature setpoints more closely, resulting in marginally lower ATD values.

However, this comes at the expense of higher energy consumption and less efficient power distribution, as the decentralized approach does not consider interactions between rooms. In contrast, our centralized control method balances comfort across all rooms while adhering to the overall power constraint. Though this can lead to slightly higher ATD values in certain cases (e.g., 0.13 vs. 0.16°C), the deviations are minimal and outweighed by improved coordination and energy distribution. Our method consistently results in lower energy consumption across most scenarios, demonstrating its advantage in multi-zone VRF scheduling with better system-wide control.

When $\omega = 0.5$ (balancing energy cost and comfort), the results show that while the ATD values for *Ours* are slightly higher than the baseline after the 20 % power constraint, the difference is minimal (e.g., 1.19°C for *Ours* vs. 1.36°C for *SOTA* on average). This reflects the trade-off between maintaining comfort and optimizing energy use. *Ours* consistently achieves lower energy consumption across all scenarios, saving up to 16 % in energy use at a 30 % power constraint, thus proving its efficiency in managing energy while still ensuring acceptable comfort. In contrast, the decentralized *SOTA* method, while maintaining lower ATD values in less constrained scenarios, consumes more energy due to its lack of coordination between rooms, resulting in inefficient power distribution and higher energy costs.

Thus, *Ours* offers a more energy-efficient solution without significant loss of comfort, especially in scenarios where balancing energy costs and thermal comfort is critical. The dynamic pre-cooling schedule in *Ours* optimizes power usage, making it a more cost-effective choice for multi-zone VRF systems. In the energy cost priority scenario ($\omega = 0.9$), *Ours* experiences a slight increase in ATD values compared to *SOTA* after the 20 % power constraint, with the ATD rising from 2.00°C for *Ours* to 1.94°C for *SOTA*. However, this increase is minimal and is outweighed by significantly lower energy consumption. At a 30 % power constraint, *Ours* reduces energy use by approximately 4 %, saving energy while maintaining acceptable comfort levels.

The key advantage of *Ours* is its ability to maintain stable thermal comfort while optimizing energy usage. As power limitations become stricter, *Ours* efficiently allocates energy across the system without prolonged pre-cooling, resulting in better overall energy efficiency. In contrast, the decentralized *SOTA* method tends to over-consume energy due to its lack of coordination, leading to higher energy usage and less efficient power distribution. Overall, our method provides an efficient solution for applications that prioritize energy savings, particularly when

facing tight power constraints. It strikes a balance between maintaining thermal comfort and reducing energy consumption, outperforming the baseline method in terms of energy savings without significant comfort compromise.

In comparison, rule-based control methods rely on fixed operational logic and lack responsiveness to changing environmental or system conditions. As a result, they tend to cause either discomfort or unnecessary energy consumption when operating under constrained scenarios. While model-based MPC approaches offer improved flexibility, they are built on predefined mathematical models that require accurate system identification. This dependence limits their robustness and makes them less suitable for scenarios where energy availability is highly restricted or when building conditions vary.

Data-driven methods based on neural networks can learn complex patterns but often lack transparency and are sensitive to noise, which makes their performance unreliable when used in situations that differ from the training data. Controllers based on learning methods, such as reinforcement learning, require large amounts of training data, and their black-box nature makes them hard to validate or adjust for real-time control in HVAC systems. In contrast, our approach combines the clarity of symbolic models with the flexibility of MPC, providing dependable performance even under strict power constraints and changing conditions.

Another advantage of the proposed method lies in its generalizability. Unlike rule-based or model-based strategies that are designed for specific building setups, our framework can be applied to various building types including residential, commercial, and industrial buildings, without needing major changes to the system structure. Since the framework is based on historical data and does not rely heavily on detailed building-specific information, it can be easily adapted to different building types with only minor adjustments. These results suggest that our method not only performs well under strict energy limitations but also provides a practical solution for future energy-efficient building control.

6. Conclusion and future work

The study employs a data-driven method to propose a coordinated operation scheduling for multiple-room VRF systems. The primary goal is to maximize energy savings while ensuring thermal comfort in all rooms. The primary contribution of this research is the development and evaluation of the effects of peak power limitations on overall energy consumption, peak power demand, and thermal comfort. Using actual historical data from an educational facility at the University of Osaka in Japan, we conducted simulations in cooling mode during the summer. Under three priority settings, comfort ($\omega = 0.1$), trade-off ($\omega = 0.5$), and cost ($\omega = 0.9$) at different power restrictions, the simulation results are assessed by contrasting the proposed coordinated approach with the baseline uncoordinated method. The proposed method continuously outperformed the baseline approach, resulting in significant improvements in thermal comfort and energy efficiency.

Under a 30 % constraint, the proposed method achieves energy savings of up to 16 %, and even under stricter constraints, it maintains acceptable thermal comfort, with only minor increases in ATD compared to the baseline. In most cases, the average temperature deviations (ATD) are continuously below 2°C, ensuring optimum thermal comfort in every room. Through coordinated control and dynamic adjustments to

pre-cooling schedules and power allocation, the proposed method improves power distribution efficiency and reduces energy waste during occupied periods.

While the proposed method offers significant advantages, some trade-offs were observed, particularly in cost-prioritized scenarios ($\omega = 0.9$), where thermal comfort was somewhat decreased to save more energy. However, these ATD increases are minimal and acceptable when cost reduction is the primary objective, showing the method's flexibility across operational priorities. The system's ability to balance energy efficiency and comfort across all scenarios highlights its potential for modern building management systems.

To further enhance the proposed method's performance and applicability, future research should explore the incorporation of renewable energy sources such as solar or wind power, which could further improve the system's sustainability and reduce reliance on the grid. Additionally, investigating the impact of real-time weather data and occupancy predictions on system performance could yield further improvements in energy efficiency and comfort. Expanding the scope of simulations to include diverse building types (e.g., residential, commercial) and climates would also help generalize the findings and validate the system's robustness in different settings.

CRedit authorship contribution statement

Theint Theint Thu: Writing – original draft, Visualization, Validation, Software, Methodology, Formal analysis, Data curation, Conceptualization; **Kenshiro Kato:** Visualization, Validation, Methodology, Investigation, Data curation, Conceptualization; **Dafang Zhao:** Writing – review & editing, Project administration, Validation, Supervision, Software, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization; **Hiroki Nishikawa:** Writing – review & editing, Supervision, Formal analysis; **Ittetsu Taniguchi:** Validation, Supervision, Project administration, Investigation, Funding acquisition, Formal analysis, Conceptualization; **Takao Onoye:** Supervision, Project administration, Conceptualization.

Data availability

Data will be made available on request.

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.

Acknowledgment

This work was supported by JSPS, Japan KAKENHI Grant Number 22H03697, 24K20901, and DAIKIN Industries, Ltd.

References

- [1] I.R.E. Agency, World energy transitions outlook 2024, 2024. Accessed: 2024-12-18, https://www.irena.org/-/media/Files/IRENA/Agency/Publication/2024/Nov/IRENA_World_energy_transitions_outlook_2024.pdf.
- [2] 2023 Global Status Report for Buildings and Construction: Beyond Foundations - Mainstreaming Sustainable Solutions to Cut Emissions from the Buildings Sector, United Nations Environment Programme, 2024. <https://doi.org/10.59117/20.500.11822/45095>
- [3] M. González-Torres, L. Pérez-Lombard, J.F. Coronel, I.R. Maestre, D. Yan, A review on buildings energy information: trends, end-uses, fuels and drivers, *Energy Rep.* 8 (2022) 626–637.
- [4] I. IEA, World Energy Outlook 2022, IEA, Paris, France, 2022, <https://www.iea.org/reports/world-energy-outlook-2022>.
- [5] X. Yu, D. Yan, K. Sun, T. Hong, D. Zhu, Comparative study of the cooling energy performance of variable refrigerant flow systems and variable air volume systems in office buildings, *Appl. Energy* 183 (2016) 725–736.
- [6] M. Qian, D. Yan, T. Hong, H. Liu, Operation and performance of VRF systems: mining a large-scale dataset, *Energy Build.* 230 (2021) 110519.
- [7] B. Seo, Y.B. Yoon, B.H. Yu, S. Cho, K.H. Lee, Comparative analysis of cooling energy performance between water-cooled VRF and conventional AHU systems in a commercial building, *Applied Thermal Engineering*, 170, 114992, 2020, Elsevier.
- [8] P. Kumar, S. Akhai, Effective energy management in smart buildings using VRF/VRF systems, In: Additive Manufacturing in Industry 4.0, 27–35, 2022, CRC Press.
- [9] Y.H. Yau, U.A. Rajput, A. Badarudin, A comprehensive review of variable refrigerant flow (VRF) and ventilation designs for thermal comfort in commercial buildings, *Journal of Thermal Analysis and Calorimetry*, 149, 1935–1961, 2024, Springer.
- [10] Y. Arroyo Gómez, J.F. San José-Alonso, L.J. San José-Gallego, J.M. Rey-Hernández, A. Sanz-Tejedor, F.J. Rey-Martínez, Optimizing Energy Efficiency and Sustainability in Winter Climate Control: Innovative Use of Variable Refrigerant Flow (VRF) Systems in University Buildings, *Applied Sciences*, 15, 2374, 2025, MDPI.
- [11] M.M. Olama, T. Kuruganti, J. Nutaro, J. Dong, Coordination and control of building HVAC systems to provide frequency regulation to the electric grid, *Energies* 11 (7) (2018) 1852.
- [12] C.R. Laughman, H. Qiao, S.A. Bortoff, Integrated control of multi-zone buildings with ventilation and VRF systems in cooling mode (2018).
- [13] C. Winstead, M. Bhandari, J. Nutaro, T. Kuruganti, Peak load reduction and load shaping in HVAC and refrigeration systems in commercial buildings by using a novel lightweight dynamic priority-based control strategy, *Appl. Energy* 277 (2020) 115543.
- [14] J.W. Moon, Y.K. Yang, E.J. Choi, Y.J. Choi, K.H. Lee, Y.S. Kim, B.R. Park, Development of a control algorithm aiming at cost-effective operation of a VRF heating system, *Applied Thermal Engineering*, 149, 1522–1531, 2019, Elsevier.
- [15] Y. Lee, W. Kim, Development of an optimal start control strategy for a variable refrigerant flow (VRF) system, *Energies*, 14, 2, 271, 2021, MDPI.
- [16] R. Karunakaran, S. Iniyar, R. Goic, Energy efficient fuzzy based combined variable refrigerant volume and variable air volume air conditioning system for buildings, *Applied Energy*, 87, 4, 1158–1175, 2010, Elsevier.
- [17] A. Jindal, S. Deb, A. Roy, B. Khan, S. Misra, A heuristic-based smart HVAC energy management scheme for university buildings, *IEEE Trans. Ind. Electron.* 65 (5) (2018) 3559–3569.
- [18] G. Gao, J. Li, Y. Wen, Energy-efficient thermal comfort control in smart buildings via deep reinforcement learning, *arXiv:1901.04693* (2019).
- [19] X. Li, S. Chen, H. Li, Y. Lou, J. Li, Multi-dimensional analysis of air-conditioning energy use for energy-saving management in university teaching buildings, *Build. Environ.* 185 (2020) 107246.
- [20] D. Kim, J. Lee, S. Do, P.J. Mago, K.H. Lee, H. Cho, Energy modeling and model predictive control for HVAC in buildings: A review of current research trends, *Energies*, 15, 19, 7231, 2022, MDPI.
- [21] G. Serale, M. Fiorentini, A. Capozzoli, D. Bernardini, A. Bemporad, Model Predictive Control (MPC) for enhancing building and HVAC system energy efficiency: problem formulation, applications and opportunities, *Energies* 11 (2018). <https://doi.org/10.3390/en11030631>
- [22] N.S. Raman, B. Chen, P. Barooah, On energy-efficient HVAC operation with Model Predictive Control: A multiple climate zone study, *Applied Energy*, 324, 119752, 2022, Elsevier.
- [23] K. Kato, D. Watari, D. Zhao, H. Nishikawa, I. Taniguchi, T. Onoye, Scheduling for multiple HVAC systems with electrical power allocation, in: 2022 IEEE 11th Global Conference on Consumer Electronics (GCCE), IEEE, 2022, pp. 226–227.
- [24] E. Biyik, A. Kahraman, A predictive control strategy for optimal management of peak load, thermal comfort, energy storage and renewables in multi-zone buildings, *J. Build. Eng.* 25 (2019). <https://doi.org/10.1016/j.jobbe.2019.100826>
- [25] Y. Ozawa, D. Zhao, D. Watari, I. Taniguchi, T. Suzuki, Y. Shimoda, T. Onoye, Data-driven HVAC control using symbolic regression: design and implementation, in: 2023 IEEE Power & Energy Society General Meeting (PESGM), IEEE, 2023, pp. 1–5.
- [26] Y. Wang, Z. Zhang, B. Dong, Model-free HVAC control in buildings: a review, *Energies* 16 (20) (2023) 7124. <https://doi.org/10.3390/en16207124>
- [27] P.C. Tabares-Velasco, A. Speake, M. Harris, A. Newman, T. Vincent, M. Lanahan, A modeling framework for optimization-based control of a residential building thermostat for time-of-use pricing, *Appl. Energy* 242 (2019) 1346–1357. <https://doi.org/10.1016/j.apenergy.2019.01.241>
- [28] J. Drgoña, J. Arroyo, I. Cupeiro Figueroa, D. Blum, K. Arendt, D. Kim, E.P. Ollé, J. Oravec, M. Wetter, D.L. Vrabie, L. Helsén, All you need to know about model predictive control for buildings, *Ann. Rev. Control* 50 (2020) 190–232. <https://doi.org/10.1016/j.arcontrol.2020.09.001>
- [29] J.S. Li, C.A. Alonso, J.C. Doyle, Frontiers in scalable distributed control: SLS, MPC, and beyond, 2021 American Control Conference (ACC), 2720–2725, 2021, IEEE.
- [30] B. Jeddi, Y. Mishra, G. Ledwich, Distributed load scheduling in residential neighborhoods for coordinated operation of multiple home energy management systems, *Appl. Energy* 300 (2021). <https://doi.org/10.1016/j.apenergy.2021.117353>
- [31] D. Xie, L. Yu, T. Jiang, Y. Zou, Distributed energy optimization for HVAC systems in university campus buildings, *IEEE Access* 6 (2018) 59141–59151. <https://doi.org/10.1109/ACCESS.2018.2872589>
- [32] H. Zhang, S. Seal, D. Wu, B. Boulet, F. Bouffard, G. Joos, Data-driven model predictive and reinforcement learning based control for building energy management: a survey, *arXiv:2106.14450* (2021).
- [33] 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). <https://doi.org/10.1016/j.apenergy.2020.115147>
- [34] A. Kathirgamanathan, M. De Rosa, E. Mangina, D.P. Finn, Data-driven predictive control for unlocking building energy flexibility: a review, *Renew. Sustain. Energy Rev.* 135 (2021). <https://doi.org/10.1016/j.rser.2020.110120>

- [35] D. Zhao, D. Watari, Y. Ozawa, I. Taniguchi, T. Suzuki, Y. Shimoda, T. Onoye, Data-driven online energy management framework for HVAC systems: An experimental study, *Appl. Energy* 352 (2023) 121921.
- [36] J. Leprince, C. Miller, M. Frei, H. Madsen, W. Zeiler, Fifty shades of black: uncovering physical models from symbolic regressions for scalable building heat dynamics identification, in: *Proceedings of the 8th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, 2021, pp. 345–348.
- [37] W. La Cava, B. Burlacu, M. Virgolin, M. Kommenda, P. Orzechowski, F.O. de França, Y. Jin, J.H. Moore, Contemporary symbolic regression methods and their relative performance, *Adv. Neural Inform. Process. Syst.* 2021 (DB1) (2021) 1.
- [38] C.C. Ukachukwu, Symbolic regression in energy engineering (2024).
- [39] M. Jung, P.R. da Costa Mendes, M. Önnheim, E. Gustavsson, Model predictive control when utilizing LSTM as dynamic models, *Eng. Appl. Artif. Intell.* 123 (2023) 106226.
- [40] M. Aria, C. Cuccurullo, A. Gnasso, A comparison among interpretative proposals for Random Forests, *Machine Learning with Applications*, 6, 100094, 2021, Elsevier.
- [41] G. Menghani, Efficient deep learning: A survey on making deep learning models smaller, faster, and better, *ACM Computing Surveys*, 55, 12, 1–37, 2023, ACM New York, NY.
- [42] D. Zhao, Z. Chen, Z. Li, X. Yuan, I. Taniguchi, Improving building temperature forecasting: a data-driven approach with system scenario clustering, in: *2024 IEEE Power & Energy Society General Meeting (PESGM)*, IEEE, 2024, pp. 1–5.
- [43] D. Zhao, D. Watari, Y. Ozawa, I. Taniguchi, T. Suzuki, Y. Shimoda, T. Onoye, A thermal comfort and peak power demand aware VRF heating/cooling management framework: simulation and on-site experiment, *J. Inform. Process.* 30 (2022) 476–485. <https://doi.org/10.2197/ipsjip.30.476>
- [44] S. Shekhar, A. Bansode, A. Salim, A comparative study of hyper-parameter optimization tools, *2021 IEEE Asia-Pacific Conference on Computer Science and Data Engineering (CSDE)*, 1–6, 2021, IEEE.
- [45] T. Koch, T. Berthold, J. Pedersen, C. Vanaret, Progress in mathematical programming solvers from 2001 to 2020, *EURO Journal on Computational Optimization*, 10, 100031, 2022, Elsevier.
- [46] A. Tonda, Review of PySR: high-performance symbolic regression in Python and Julia, 2025, Springer.