

# Constraint-aware generative operation strategy for multiple AHUs using graph recurrent forecasting

Seungkeun Yeom<sup>a</sup>, Juui Kim<sup>a</sup>, Jinwoo Choi<sup>a</sup>, Gaurav Kattel<sup>b</sup>, Anup Paudel<sup>b</sup>, Jun-Ki Choi<sup>b</sup>,  
Taehoon Hong<sup>a,\*</sup>

<sup>a</sup> Department of Architecture and Architectural Engineering, Yonsei University, Seoul, 03722, South Korea

<sup>b</sup> Department of Mechanical Engineering, University of Dayton, Dayton, OH, 45469-0238, USA

## ARTICLE INFO

### Keywords:

Air handling unit  
Conditional variational autoencoder  
Adaptive graph convolutional recurrent network  
Indoor environmental quality  
Energy consumption

## ABSTRACT

The operation of air handling units (AHUs) significantly affects indoor environmental quality (IEQ) and building energy consumption, especially in systems with multiple AHUs where control is complicated by spatial interactions and operational uncertainty. This study proposes a data-driven framework for optimal operation of multiple AHUs by integrating conditional scenario generation, spatiotemporal IEQ forecasting, energy consumption prediction, and constraint-based strategy selection. A conditional variational autoencoder (CVAE) generates AHU operating schedules conditioned on weather, temporal characteristics, and previous operating history, and the generated schedules are further expanded through rule-based augmentation. An adaptive graph convolutional recurrent network (AGCRN) is then employed to predict indoor air temperature and CO<sub>2</sub> concentration dynamics across multiple AHUs, while a gated recurrent unit (GRU)-based sequence-to-sequence model predicts building-level energy consumption. The generated scenarios are evaluated using these surrogate models, and feasible candidates are selected by enforcing indoor air temperature and CO<sub>2</sub> concentration constraints. Among the feasible candidates, the scenario with the minimum cumulative energy consumption is selected as the final operating strategy. A case study of a large multi-purpose building equipped with 10 AHUs shows that the proposed framework can generate realistic and diverse operating scenarios, maintain acceptable IEQ conditions, and achieve meaningful energy savings. The results also reveal distinct seasonal operating types and provide practical guidelines for operation of multiple AHUs. These findings suggest that the proposed framework can support data-driven multi-AHU operation by screening historically plausible scenarios under IEQ constraints.

## 1. Introduction

The building sector accounts for approximately 30% of total global energy consumption, with air handling units (AHUs) comprising more than 40% of building energy use [1,2]. As a core facility that simultaneously manages indoor thermal environment control and indoor air quality (IAQ) through outdoor air intake, optimizing AHU operation is a key challenge for achieving both improved building energy efficiency and the assurance of thermal comfort and IAQ [3,4]. Particularly in

large-scale multi-purpose facilities or multi-zone buildings, multiple AHUs operate simultaneously, exhibiting characteristics where supply and return air conditions, indoor CO<sub>2</sub> concentration, and energy consumption are interconnected according to changes in outdoor conditions and occupancy patterns. Therefore, in large-scale multi-purpose facilities, AHU operation should not be limited to independent control of individual devices but approached with an integrated operational strategy considering spatial interdependence among multiple AHUs [5, 6].

**Abbreviation:** AGCRN, Adaptive graph convolutional recurrent network; AHU, Air handling unit; BAS, Building automation system; CVAE, Conditional variational autoencoder; GRU, Gated recurrent unit; HVAC, Heating, Ventilation, And air conditioning; IAQ, Indoor air quality; IEQ, Indoor environmental quality; KL, Kullback-Leibler; LSTM, Long short-term memory; MAE, Mean absolute error; MAPE, Mean Absolute percentage error; ML, Machine learning; MLP, Multi-layer perceptron; RMSE, Root mean squared error; RNN, Recurrent neural network; Seq2Seq, Sequence-to-sequence.

\* Corresponding author at: Yonsei University, 50 Yonsei-ro, Seodaemun-gu, Seoul, 03722, South Korea.

**E-mail addresses:** [kevinysk@yonsei.ac.kr](mailto:kevinysk@yonsei.ac.kr) (S. Yeom), [caution0210@yonsei.ac.kr](mailto:caution0210@yonsei.ac.kr) (J. Kim), [jinwoo818@yonsei.ac.kr](mailto:jinwoo818@yonsei.ac.kr) (J. Choi), [kattelg1@udayton.edu](mailto:kattelg1@udayton.edu) (G. Kattel), [paudela5@udayton.edu](mailto:paudela5@udayton.edu) (A. Paudel), [jchoil@udayton.edu](mailto:jchoil@udayton.edu) (J.-K. Choi), [hong7@yonsei.ac.kr](mailto:hong7@yonsei.ac.kr) (T. Hong).

<https://doi.org/10.1016/j.buildenv.2026.115115>

Received 3 June 2026; Received in revised form 26 July 2026; Accepted 11 August 2026

Available online 14 August 2026

0360-1323/© 2026 Elsevier Ltd. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

Against this background, data-driven approaches such as machine learning (ML) and deep learning (DL) have been actively studied over the past decade for the purpose of establishing and optimizing AHU operational strategies, including applications to equipment-level fault detection and diagnosis [7–9] as well as operational strategy optimization. Recent studies can be broadly categorized into three research directions: (i) collaborative operation and optimization of multiple AHUs, (ii) generative scenario modeling for exploring diverse operation strategies, and (iii) graph-based spatiotemporal prediction of indoor environmental conditions.

The first research direction focuses on collaborative operation and optimization of multiple AHUs serving interconnected indoor spaces. Representative approaches include model predictive control (MPC), reinforcement learning (RL), and multi-objective optimization, which optimize HVAC operation while considering interactions among multiple zones. For example, Ra et al. [10] demonstrated the applicability of real-time MPC for HVAC operation in a large factory building, whereas Soleimani et al. [11] combined deep learning with Koopman predictive control for multi-objective optimization of HVAC systems. Zhang et al. [12] further integrated building automation systems (BAS), the Internet of Things (IoT), and building information modeling (BIM) into a graph representation to support coordinated operation of central air-conditioning systems. Although these approaches improve operational efficiency, they generally optimize operating conditions within predefined control spaces and provide limited capability for exploring previously unobserved operation strategies. Moreover, approaches based on multi-objective optimization typically require assigning explicit weights to competing objectives such as energy savings and occupant comfort; however, comfort is a human-centered outcome that is difficult to weight objectively, and the acceptable trade-off varies substantially across space types.

The second research direction aims to expand the range of candidate operation strategies through generative scenario modeling. Unlike deterministic prediction models that evaluate only historically observed operation patterns, generative approaches can synthesize statistically plausible but previously unobserved operating scenarios, thereby expanding the candidate operation space under limited training data. Such approaches are particularly beneficial for newly constructed buildings or environments with insufficient operational data. For example, Fu et al. [13] employed a conditional diffusion model to generate synthetic building energy meter data conditioned on building metadata, while Song et al. [14] applied a sequence generative adversarial network to forecast residential indoor temperature patterns. Parisio et al. [15] proposed a scenario-based predictive control approach for building HVAC management, generating multiple plausible operating scenarios to account for uncertainty in building operating conditions. However, these studies mainly focus on scenario generation or scenario-based control itself and do not provide a framework for estimating the indoor environmental quality (IEQ) and energy consequences of newly generated, previously unobserved operation schedules.

The third research direction focuses on graph-based spatiotemporal prediction of indoor environmental conditions. Recent studies have shown that graph neural network (GNN)-based models effectively capture spatial dependencies among multiple zones and HVAC systems. Huang et al. [16] proposed a graph-based spatiotemporal framework for predicting indoor CO<sub>2</sub> concentration considering dynamic occupancy and environmental conditions. Wen et al. [17] developed an adaptive spatiotemporal graph convolutional network to predict high-resolution indoor temperature fields, while Zhao et al. [18] presented a multi-task learning-based deep neural network for predicting temperature and relative humidity at multiple indoor locations. Zhao et al. [19] additionally applied long short-term memory (LSTM) to predict building energy consumption for HVAC operation. Although these studies significantly improve prediction accuracy by incorporating spatial relationships, they primarily focus on environmental forecasting rather

than supporting operational strategy evaluation.

Despite these advances, several important research gaps remain in integrating these research directions into a practical operational framework for multiple AHUs. First, although collaborative control approaches improve operational efficiency, they generally optimize operation within predefined control spaces and do not sufficiently explore previously unobserved operating scenarios. Second, while generative approaches can expand the operation scenario space, they do not estimate the resulting IEQ and energy performance of the generated schedules. Third, graph-based prediction models effectively capture spatial dependencies among multiple AHUs but are primarily designed to improve prediction accuracy rather than support operational strategy selection under practical IEQ constraints.

To overcome these limitations, this study proposes a multiple AHUs operational framework integrating operational scenario generation, spatiotemporal environment prediction, and constraint-based operational strategy selection. First, a conditional variational autoencoder (CVAE) is utilized to stochastically generate weekly AHU operational patterns conditioned on outdoor conditions, temporal features, and past operational history, thereby securing diverse candidate operational scenarios. Because the available operational data cannot represent all possible AHU operating schedules, the generated scenarios include historically plausible but previously unobserved operation patterns. Expanding the scenario space in this manner enables the framework to explore candidate strategies beyond those contained in the historical dataset.

Next, the performance of each scenario is evaluated by combining an adaptive graph convolutional recurrent network (AGCRN)-based indoor air temperature and CO<sub>2</sub> concentration prediction model and a gated recurrent unit (GRU)-based energy consumption prediction model. Since the indoor environmental and energy responses of these newly generated scenarios are unknown, the AGCRN and GRU models are employed to estimate their IEQ and energy consequences. Consequently, the three models operate as complementary components of a unified framework rather than as independent prediction modules.

Finally, scenarios satisfying IEQ constraints for indoor air temperature and CO<sub>2</sub> concentration are selected, and the operational strategy with the minimum energy consumption within the selected scenario set is chosen to derive a control strategy applicable to actual building operation (refer to Fig. 1). Unlike predictive control or optimization-based approaches that require an explicit system dynamics model and iterative online optimization over continuous control variables, this offline, scenario-screening design was adopted to reflect the discrete, schedule-based nature of AHU control in the case-study building.

## 2. Material and methods

### 2.1. Data description

#### 2.1.1. Target building and monitoring system

In this study, a multi-purpose building located on the central campus of a university in Seodaemun-gu, Seoul, South Korea, was selected as the case study. The building is a steel-reinforced concrete structure with a total floor area of 28,300 m<sup>2</sup> and includes halls and dining facilities for seminars, exhibitions, and large-scale events (refer to Fig. 2). The building is served by 10 AHUs installed across individual rooms and zones, each operating under a constant air volume scheme to regulate ventilation and thermal conditions within its designated area. The operation of each AHU is continuously monitored and controlled through the BAS, which also enables the systematic collection and storage of operational data.

The monitoring system consists of AHU-level operational sensors, zone-level indoor environmental sensors, and occupancy sensors. At the AHU level, the BAS records key operating variables, including supply, return, and mixed air temperatures, supply and return airflow rates, fan or inverter signals, supply air pressure, CO<sub>2</sub> concentration, relative




Fig. 1. Research framework.




Fig. 2. Information and layout of target building.

humidity, and AHU on/off status. These variables were used to characterize both the operating state of each AHU and the

corresponding environmental response of the served zones. At the zone level, indoor air temperature and occupancy status were collected at 10-minute intervals from sensors installed in the spaces served by each AHU. The sensors were positioned to represent the thermal and occupancy conditions of the corresponding zones, including halls, dining areas, and event-oriented spaces. The spatial relationship between AHUs and their served zones is illustrated in Fig. 2.

Energy consumption data were obtained from YEPS, the campus energy data platform, where hourly energy use of the target building was recorded over the same analysis period [20]. These data were used to train and validate the energy consumption prediction model, enabling a quantitative assessment of the effects of AHU operating conditions on overall building energy consumption. The indoor environmental data were integrated with AHU operational data to analyze indoor environmental responses to AHU operation and to develop the IEQ prediction model.

### 2.1.2. AHU data set

In this study, AHU operational data collected through the BAS were integrated with weather observations to account for outdoor conditions at each time step. The two datasets were temporally aligned to construct a unified dataset. The resulting AHU dataset consists of multivariate time series data containing sensor-measured variables that characterize both AHU operation and indoor environmental conditions, including supply, return, and mixed air temperatures, supply and return airflow rates, CO<sub>2</sub> concentration, and supply air pressure. All variables were collected at 10-minute intervals over a continuous one-year period from January 1 to December 31. A total of 10 AHUs were included in the analysis, with each AHU defined as an independent node in the prediction model. The input variables for each node were grouped into three categories: control variables, state variables, and exogenous variables. Detailed definitions of these variables are provided in Table 1.

The AHU supply air set point in the target building was fixed at 23 °C throughout the study period. On and off indicator representing system operating status was also included as a separate state variable. Exogenous variables were defined as factors that vary independently of AHU operating conditions but influence system behavior. These include weather variables such as outdoor air temperature, relative humidity, wind speed, and precipitation, as well as temporal features. To account for temporal periodicity, time-of-day and day-of-week information were encoded using sine–cosine cyclic representations and included in the input features. Zone-level indoor environmental measurements were obtained using wall-mounted Pt1000 resistance temperature detector (RTD)-based temperature sensors (model KRTH-21), relative humidity sensors, and non-dispersive infrared (NDIR)-based CO<sub>2</sub> sensors, installed at approximately 1.5 m height on interior walls, while occupancy was measured using ceiling-mounted passive infrared (PIR) occupancy sensors. The temperature and humidity sensors provide measurement accuracies of approximately ±0.5 °C and ±3%RH, respectively, and the CO<sub>2</sub> sensors have a measurement range of 0–2000 ppm with an accuracy of ±3% full scale.

**Table 1**  
AHU data features.

| Variable group    | Included measurements (unit)                                               |
|-------------------|----------------------------------------------------------------------------|
| Temperature       | Supply air temperature, Mixed air temperature, Return air temperature (°C) |
| Fan/drive signals | Supply/ventilation inverter outputs (%)                                    |
| IAQ               | CO <sub>2</sub> concentration (ppm)                                        |
| Flow/pressure     | Supply/ventilation airflow (m <sup>3</sup> /h), supply air pressure (mmAq) |
| Moisture          | Relative humidity (%)                                                      |
| Occupancy         | Room-level occupancy rate derived from occupancy sensors (%)               |

## 2.2. AHU operation pattern generation

### 2.2.1. AHU operation scenario generation

AHU operation can be formulated as a conditional generative problem influenced by outdoor conditions, temporal characteristics, recent operational history, and historical patterns from the same day of the week. Based on these characteristics, this study develops a CVAE-based scenario generation model as illustrated in Fig. 3. The CVAE is a conditional latent variable model that learns the posterior distribution of target variables given a condition vector, enabling the sampling of diverse operational scenarios [21]. Compared to generative adversarial network models, the CVAE provides more stable training. In contrast to deterministic regression models, it enables probabilistic sampling in the latent space, which naturally enhances generative diversity. This is particularly suitable for AHU operation data, where patterns vary significantly depending on weekday and weekend, seasonal changes, and weather variability.

The CVAE is trained by maximizing the evidence lower bound [22] (refer to Eq. (1)).

$$\log p(x|c) \geq \mathbb{E}_{q_\phi(z|x,c)}[\log p_\theta(x|z,c)] - \text{KL}(q_\phi(z|x,c) \| p(z|c)) \quad (1)$$

Here,  $q_\phi(z|x,c)$  denotes the posterior distribution of the latent variable estimated by the encoder. In this study, it is assumed to follow a Gaussian distribution (refer to Eq. (2)).

$$q_\phi(z|x,c) = \mathcal{N}(z; \mu_\phi(x,c), \text{diag}(\sigma_\phi^2(x,c))) \quad (2)$$

The latent variable is sampled using the reparameterization trick, and the decoder generates the estimated operational parameters  $\hat{x} = f_\theta(z,c)$  from the latent vector and condition vector.

The target variable  $x$  represents the daily operational pattern of each AHU. For AHU  $n$ , the generated operation scenario is defined as  $x = \{s_n, u_n, b_n\}$ , where  $s_n$ ,  $u_n$ ,  $b_n$  represent the start time, operation duration, and binary operation status. The start time and duration are represented at an hourly resolution, while  $b_{n,d} \in \{0, 1\}$  indicates whether the AHU operates on that day. Based on the observed data, each daily AHU schedule is simplified as a single contiguous operation block.

The condition vector  $c$  is constructed by combining exogenous variables and historical operational information that influence AHU operation. It includes (1) weather features, (2) temporal features, (3) previous operation history. Weather features consist of outdoor temperature, relative humidity, wind speed, wind direction, and precipitation, represented as statistical summaries of high-resolution time series data. Temporal features including day of week, month, and week of year are encoded using sinusoidal representations. In addition, the condition vector includes the previous day's operational parameters and historical operation patterns from the same day of the week over the past four weeks.

The network consists of MLP-based encoder and decoder modules, with a self-attention mechanism to capture interdependencies among multiple AHUs. Since operational patterns differ across AHUs even under identical external conditions, learnable embedding vectors are introduced for each AHU. Specifically, the condition vector is transformed into a shared representation through an MLP, which is then combined with AHU specific embeddings to form AHU level condition tokens. Additionally, day embeddings are incorporated to capture variations across different days, allowing the model to represent day specific operational behaviors for each AHU.

Although the CVAE is capable of generating diverse operational scenarios by learning conditional distribution, the generated samples tend to concentrate around dominant patterns in the training data. This characteristic may limit the exploration of AHU operation strategies with high energy saving potential. To address this limitation, rule-based augmentation was applied to the scenarios generated by the CVAE in order to expand scenario diversity.

Specifically, several transformation strategies were introduced,




Fig. 3. Framework of the CVAE-based AHU operational scenario generation and evaluation.

including temporal shifts of operation schedules, adjustment of operation duration, modification of on and off status for specific days, and incorporation of previous week operational patterns. In addition, a scenario shuffling strategy was applied, in which daily operation patterns corresponding to the same day of the week are exchanged across different generated scenarios. For this purpose, operational statistics over the past 12 weeks were precomputed for each AHU. These statistics include daily operation status, operation start time, operation duration, daily state transition probabilities, and historical operation patterns. Based on these statistics, augmented scenarios were designed to remain consistent with historical distributions while exploring a broader solution space than CVAE generated scenarios alone.

### 2.2.2. Scenario generation settings and evaluation

The CVAE was trained using one year of AHU operational data and weather data collected at a 10-minute resolution. The dataset was split into training, validation, and test sets with ratios of 70 percent, 15 percent, and 15 percent, respectively, while preserving temporal order. The AdamW optimizer and a cosine annealing learning rate scheduler were applied. The batch size and total number of training epochs were set to 32 and 100. To prevent posterior collapse during early training, a warm up strategy was adopted in which the weight of the KL term was linearly increased over the first 40 epochs. Gradient clipping with a maximum norm of 1.0 was applied to ensure training stability, and the model with the lowest validation loss was selected as the final model.

To quantitatively evaluate the quality of the generated AHU operation scenarios, a multi-dimensional evaluation framework was established. This framework assesses both how well the generated distribution reflects real operational patterns and how effectively it explores diverse operational strategies. This approach follows prior studies that jointly consider distribution similarity and scenario diversity in generative model evaluation [23,24].

- **Distribution similarity:** The Kullback-Leibler divergence (KL divergence) was calculated to evaluate the similarity between the generated distribution and the real distribution. A value close to zero indicates high similarity, suggesting that the generative model

accurately reproduces the probabilistic structure of actual operational patterns [13,25]. In general, values below 0.5 are considered to indicate stable reproduction of the real distribution in generative model evaluation studies. In this study, the KL divergence was used to verify that the distribution of generated operational scenarios does not significantly deviate from real AHU operation patterns.

- **Coverage:** To evaluate how well the generated scenarios span the valid range of the real operational distribution, the proportion of generated samples falling within the empirical 5th to 95th percentile range of the real data was calculated. This metric allows simultaneous identification of two types of failure, namely excessive generation of unrealistic values outside the real distribution and insufficient exploration of the extreme but valid range of the distribution [26]. A coverage value above 0.9 is generally interpreted as indicating that the generated distribution sufficiently covers the valid range of the real data.
- **Scenario diversity:** To evaluate the diversity within the generated scenario set, the average pairwise distance between scenarios was computed. Pairwise distance is commonly used to diagnose mode collapse, where generated samples become overly similar [27–29]. Larger pairwise distance indicates that diverse operational patterns are sufficiently generated. In this study, higher pairwise distance within the bounds of realistic distributions was interpreted as improved scenario diversity. In addition, histogram-based entropy was calculated to quantify the uncertainty and dispersion of generated operational parameters such as operation start time and duration. The entropy ratio, defined as the entropy of generated scenarios normalized by that of real operational data, was used as a key evaluation metric.

### 2.3. Spatial temporal indoor environmental quality prediction model

#### 2.3.1. AGCRN model architecture

AHUs within the same building interact with each other, and the effects of control inputs are reflected in indoor environmental conditions with temporal delays. Therefore, simple recurrent neural networks (RNN)-based or MLP-based time series models that treat each AHU

independently are insufficient to capture these spatial and temporal dependencies. To address this limitation, this study adopts the AGCRN, which is capable of jointly learning inter AHU interactions and temporal dynamics.

The AGCRN combines adaptive graph convolution for modeling spatial dependencies with an RNN for capturing temporal dependencies [30]. Rather than relying on a predefined physical adjacency matrix, AGCRN learns a node embedding matrix to construct an adaptive adjacency matrix during training. This allows the model to automatically discover latent spatial relationships among AHUs without requiring prior knowledge of their connectivity. Spatial features aggregated through graph convolution are then passed to a GRU layer, enabling the model to capture both cross-AHU interactions and dynamic temporal behavior simultaneously. In this study, each AHU is defined as a node in the graph. The input vector for node  $i$  at time  $t$ ,  $\mathbf{x}_{t,i} \in \mathbb{R}^{D_x}$ , includes control variables, state variables, weather conditions, and temporal features. The output vector  $\mathbf{y}_{t,i} \in \mathbb{R}^{D_y}$  consists of indoor air temperature, and CO<sub>2</sub> concentration. The prediction task is formulated as a sequence-to-sequence function approximation problem as Eq. (3), where the model predicts future indoor environmental states over a horizon of  $H$  time steps based on historical observations over  $L$  time steps.

$$\hat{\mathbf{Y}}_{t+1:t+H} = f_{\theta}(\mathbf{X}_{t-L+1:t}) \quad (3)$$

Indoor air temperature, and CO<sub>2</sub> concentration exhibit different physical characteristics and temporal dynamics. To capture these differences, a multi head structure is introduced after the shared encoder representation. In other words, each target variable is predicted through an independent projection head. A sequence-to-sequence architecture further mitigates error accumulation in multi-step prediction by generating future outputs autoregressively from an encoded hidden state.

In addition, CO<sub>2</sub> concentration exhibits cumulative dynamics, and directly predicting absolute values can lead to drift and error accumulation in long term predictions. To address this issue, the prediction target for CO<sub>2</sub> concentration is defined as the change relative to the value at the prediction start time (refer to Eq. (4)).

$$\Delta\text{CO}_{2,t+h,i} = \text{CO}_{2,t+h,i} - \text{CO}_{2,t,i} \quad (4)$$

By predicting relative changes instead of absolute values, the model is encouraged to focus on temporal variation patterns, which improves stability in long horizon prediction.

### 2.3.2. Model training and evaluation

Experiments were conducted on 10 AHUs using operational and environmental time series data resampled at a 10-minute interval. The input sequence length  $L$  was set to 72 time steps, corresponding to 12 h of historical data, and the prediction horizon  $H$  was set to 36 time steps, corresponding to 6 h ahead multi step forecasting. The input dimension at each time step was  $D_x = 19$ , and the output dimension was  $D_y = 2$ , consisting of indoor air temperature, and CO<sub>2</sub> concentration. The dataset was split in chronological order, with training, validation, and test ratios of 0.6, 0.2, and 0.2, respectively.

Model training was performed by minimizing the mean squared error over all prediction horizons, nodes, and output variables. As shown in Eq. (5), the loss function aggregates prediction errors across time steps and AHU nodes, enabling joint optimization of spatiotemporal predictions.

$$\mathcal{L} = \frac{1}{HND_y} \sum_{h=1}^H \sum_{i=1}^N \|\mathbf{y}_{t+h,i} - \hat{\mathbf{y}}_{t+h,i}\|_2^2 \quad (5)$$

For model configuration, the number of AHUs was set to  $N = 10$ , the hidden dimension was set to  $D_h = 64$ , and the number of recurrent layers was set to 2. The graph convolution employed a Chebyshev polynomial with order  $K = 2$ . The node embedding dimension for adaptive graph learning was set to  $d_e = 10$ . The model was trained using

the Adam optimizer with a learning rate of  $1 \times 10^{-3}$ . The batch size was set to 32, and gradient clipping was applied with a maximum norm of 5.0.

### 2.4. Energy consumption prediction model

In this study, a deep learning-based energy consumption prediction model was developed to estimate building-level energy consumption from multiple AHUs operational states using AHU sensor data. The model serves as a surrogate energy model for quantitatively evaluating the effects of AHU operating scenarios on building energy use and is repeatedly employed as a predictive module during the subsequent operation strategy exploration process.

The energy consumption prediction model was built on a GRU-based sequence-to-sequence (Seq2Seq) architecture [27,28]. Rather than directly predicting total energy consumption, the model is designed to predict the residual energy component  $\Delta E$ , defined as Eq. (6).

$$\Delta E(t) = E_{\text{Total}}(t) - E_{\text{Base}}(t) \quad (6)$$

Here,  $E_{\text{Base}}(t)$  denotes the baseline energy consumption attributable to outdoor conditions and temporal characteristics, whereas,  $\Delta E(t)$  represents the additional energy consumption beyond the baseline that results from changes in AHU operating conditions, such as on and off status, airflow rate, and set points. It is also consistent with time-series decomposition practices, in which trend or recurring components are separated to reduce the learning burden for nonstationary forecasting problems [31]. Based on this decomposition, the model is designed to first capture the baseline load driven by outdoor and temporal patterns and then separately learn the incremental contribution of AHU operational changes to total energy consumption.

The model consists of three encoders and one decoder. The weather encoder learns baseline energy patterns from outdoor weather variables and temporal features, while the energy encoder captures autoregressive energy patterns from historical energy time series. The AHU encoder, in turn, learns the effect of AHU operational variables on energy consumption. The decoder then estimates, at each time step in the prediction horizon, the baseline energy component and the AHU-driven residual component, which are combined to predict the final energy consumption.

The input data include AHU sensor data, outdoor weather data, temporal features, and historical energy consumption. AHU operational variables include supply airflow rate and operating status (on/off), collected at 10-minute intervals. To account for the combined influence of multiple AHUs, area-weighted aggregation was applied using the service area of each AHU as the weighting factor. Exogenous inputs include weather variables such as outdoor air temperature, relative humidity, wind speed, and precipitation, along with sinusoidal encodings of time-of-day and day-of-week.

To ensure consistency with the IEQ prediction model, the energy consumption model was trained and evaluated on the same dataset split and time period. The dataset was divided chronologically into training, validation, and test sets at ratios of 0.6, 0.2, and 0.2, respectively. Using 10-minute interval data, the model takes the previous 12 h ( $L = 72$ ) as input and predicts energy consumption over the subsequent 6 h ( $H = 36$ ), formulated as a multi-step sequence-to-sequence problem using a sliding-window approach. Model training was performed by minimizing Smooth L1 loss ( $\beta = 0.5$ ) over all prediction horizons. The AdamW optimizer was used with a learning rate of  $1 \times 10^{-4}$  and weight decay of  $1 \times 10^{-4}$ . The hidden dimensions were set to 64 for the weather and energy encoders and 32 for the AHU encoder.

### 2.5. Constraint-based analysis of AHU operation strategies

The final operation strategy was selected through a constraint-based sequential decision-making procedure. For each AHU operation scenario

generated by the CVAE, the IEQ prediction model and the energy consumption prediction model were applied to estimate indoor air temperature, CO<sub>2</sub> concentration, and energy consumption over the prediction horizon.

Feasibility filtering was then performed by enforcing indoor environmental quality constraints. Specifically, indoor air temperature was required to remain within a predefined acceptable range, while CO<sub>2</sub> concentration was required to remain below a specified threshold, both evaluated at the aggregate occupied-hour level. Scenarios satisfying these constraints were retained as feasible candidates. Among the feasible candidates, energy performance was evaluated, and the scenario with the minimum cumulative energy consumption over the prediction horizon was selected as the final operation strategy.

If no scenario satisfied all constraints simultaneously, a relaxed selection procedure was applied. Scenarios were first ranked based on the magnitude of constraint violations in temperature and CO<sub>2</sub> concentration, and those with the smallest violations were prioritized. Among these, the scenario with the lowest energy consumption was selected. This constraint-based selection framework reflects practical building operation strategies, where indoor environmental quality is prioritized before energy efficiency, thereby ensuring operational feasibility and applicability [15,32]. IEQ maintenance is treated as a priority screening criterion rather than a weighted objective in a multi-objective formulation, because occupant comfort is a human-centered outcome that is difficult to assign an objective weight to, and the acceptable trade-off between comfort and energy savings varies substantially across space types. Compared with iterative optimization methods, the proposed framework requires only forward inference of the trained surrogate prediction models for generated candidate schedules, thereby avoiding repeated optimization. This computational advantage extends to other search-based approaches: MPC requires repeated online optimization at each control interval, reinforcement learning requires extensive training and exploration to converge on a policy, and evolutionary algorithms require population-based iterative search over multiple generations. In contrast, the proposed framework performs a single forward inference pass per candidate schedule, remaining computationally lightweight and compatible with existing schedule-based AHU operation.

### 3. Results

#### 3.1. Evaluation results of CVAE-based operational scenario generation

The quality of the generated AHU operation scenarios was evaluated from three perspectives: distribution similarity, coverage, and scenario diversity (refer to Table 2).

First, distribution similarity between real operational patterns and generated scenarios was quantitatively assessed using KL divergence. The results show that the average KL divergence across all AHUs was 0.193, with most AHUs exhibiting values below 0.3. Although a relatively higher value (0.473) was observed for AHU08, which exhibits greater operational variability, the overall results indicate that the

generated scenarios closely follow the statistical characteristics of real AHU operation patterns. This suggests that the generative model effectively captures the underlying probabilistic structure of real operational data and produces realistic operation scenarios.

Next, coverage was evaluated to assess how well the generated scenarios span the valid range of real operational distributions. Specifically, the proportion of generated samples falling within the empirical  $[P_5, P_{95}]$  range of the real data was calculated, resulting in an average coverage of 0.915 (refer to Fig. 4). This indicates that most generated scenarios lie within the valid operating range of real data. At the same time, the model avoids excessive generation of unrealistic extreme values while adequately covering the effective range of actual operation patterns.

Finally, scenario diversity was evaluated by computing the average pairwise distance among generated scenarios. The results show that the average pairwise distance was 3.434 h for operation start time and 3.543 h for operation duration. This indicates that the generated scenarios do not converge to overly similar patterns but instead explore a sufficiently diverse range of operation start times and durations. In general, mode collapse in generative models leads to significantly reduced pairwise distances; however, such behavior was not observed in this study.

In addition, entropy analysis further confirms the diversity of the generated scenarios. The entropy of the generated distribution was higher than that of the real operational data, with entropy values of 2.53 for operation start time and 2.69 for operation duration, whereas the corresponding values for the real data were 1.70 and 2.11, respectively. These results indicate that the model preserves the key structural characteristics of real operation patterns while generating a broader range of plausible scenarios.

Overall, the generated AHU operation scenarios maintain strong distributional similarity to real operation patterns while achieving sufficient diversity. This balance provides a robust foundation for exploring a wide range of candidate operation strategies in subsequent optimization.

#### 3.2. Spatiotemporal prediction performance of indoor environmental quality

The IEQ prediction performance was evaluated under a temporally ordered data split into training, validation, and test sets. Fig. 5 presents a comparison between the actual and predicted time series over the test period, showing that the model accurately captures both short-term fluctuations and overall trends of temperature-related variables.

In terms of quantitative performance, AGCRN achieved the highest accuracy across all output variables compared to baseline models. Based on mean absolute percentage error (MAPE), AGCRN recorded 1.926% for indoor air temperature, and 5.729% for CO<sub>2</sub> concentration. These results indicate consistently lower prediction errors than those of other graph-based models, including attention based spatial temporal graph convolutional network (ASTGCN) [33], dynamic directed graph (DDGCRN) [34], multivariate time series graph neural network (MTGNN) [35], as well as conventional time series models such as GRU, LSTM, MLP, TCN, and Transformer (refer to Table 3).

To further contextualize the performance, the results were compared with existing studies on short-term IEQ prediction. Hamayat et al. [36] reported a MAPE range of 2.5–5.3% for indoor air temperature prediction in residential buildings using a Bi-LSTM model. Song et al. [14] achieved approximately 2.3% MAPE using a sequence generative adversarial network (SGAN) for short-term indoor air temperature prediction. For CO<sub>2</sub> concentration prediction, Yang et al. [37] reported approximately 15% MAPE for 30-minute ahead forecasting in school classrooms using a Bayesian-optimized LSTM model, while Taheri and Razban [38] reported MAPE values ranging from 16% to 21% using various machine learning models, including MLP. Compared with these studies, the proposed AGCRN model demonstrates competitive or superior performance, particularly for CO<sub>2</sub> concentration prediction.

**Table 2**

Evaluation metrics of generated AHU operation scenarios.

| AHU   | KL divergence | Pairwise distance (h) |          | Entropy    |          |
|-------|---------------|-----------------------|----------|------------|----------|
|       |               | Start time            | Duration | Start time | Duration |
| AHU01 | 0.163         | 3.70                  | 3.11     | 1.17       | 1.09     |
| AHU02 | 0.114         | 2.65                  | 4.09     | 2.75       | 1.60     |
| AHU03 | 0.311         | 3.78                  | 3.68     | 1.04       | 1.00     |
| AHU04 | 0.085         | 3.60                  | 3.10     | 1.94       | 1.55     |
| AHU05 | 0.068         | 3.37                  | 4.15     | 1.05       | 1.14     |
| AHU06 | 0.124         | 3.19                  | 3.63     | 3.40       | 1.28     |
| AHU07 | 0.212         | 3.89                  | 3.38     | 1.13       | 0.97     |
| AHU08 | 0.473         | 3.12                  | 3.55     | 1.68       | 1.17     |
| AHU09 | 0.246         | 3.97                  | 3.47     | 2.00       | 2.40     |
| AHU10 | 0.132         | 3.07                  | 3.28     | 1.41       | 1.41     |




Fig. 4. Distribution comparison between generated and real AHU operation scenarios.




Fig. 5. Comparison of predicted and actual indoor air temperature and CO<sub>2</sub> concentration on the test set.

The performance improvement of AGCRN can be attributed to its ability to jointly capture spatial interactions and temporal dependencies in multiple AHUs. In the target building, 10 AHUs operate in parallel, and their operational states interact through shared influences such as outdoor conditions, occupancy variations, and air mixing processes. AGCRN constructs an adaptive adjacency matrix from learnable node embeddings, enabling the model to capture these interdependencies in a data-driven manner. As a result, the model can effectively learn latent correlations among AHUs without relying on predefined physical adjacency or distance-based relationships. To further quantify the

contribution of the main architectural components, an ablation study was conducted by modifying the output-head structure, sequence-to-sequence decoder, future control inputs, and adaptive graph propagation. As summarized in Table 4, the full AGCRN achieved the most balanced overall performance. Removing control inputs, such as set-point temperature and supply flow, produced the largest degradation in temperature prediction, increasing the MAPE to 2.295%, which indicates that AHU operation information is important for capturing thermal responses. Removing the sequence-to-sequence structure also increased the indoor air temperature MAPE to 2.167%, although the

**Table 3**

Performance comparison of AGCRN and baseline models.

| Category          | Model       | Indoor air temperature (MAPE %) | Indoor CO <sub>2</sub> concentration (MAPE %) |
|-------------------|-------------|---------------------------------|-----------------------------------------------|
| Graph-based model | AGCRN       | 1.926                           | 5.729                                         |
|                   | ASTGCN      | 6.154                           | 11.658                                        |
|                   | DDGCRN      | 5.203                           | 9.086                                         |
|                   | MTGNN       | 5.351                           | 9.977                                         |
| Time-series model | GRU         | 7.133                           | 11.098                                        |
|                   | LSTM        | 7.025                           | 11.069                                        |
|                   | MLP         | 7.166                           | 11.292                                        |
|                   | TCN         | 6.763                           | 10.952                                        |
|                   | Transformer | 7.445                           | 10.973                                        |

**Table 4**

Performance comparison of AGCRN and baseline models.

| Study                     | Indoor air temperature (MAPE %) | Indoor CO <sub>2</sub> concentration (MAPE %) |
|---------------------------|---------------------------------|-----------------------------------------------|
| Full model                | 1.926                           | 5.729                                         |
| Single head               | 2.118                           | 6.077                                         |
| Without seq2seq structure | 2.167                           | 5.323                                         |
| Without control input     | 2.295                           | 5.635                                         |
| Identity graph            | 2.114                           | 5.552                                         |

indoor CO<sub>2</sub> concentration error decreased slightly, suggesting that the control-conditioned decoder contributes more strongly to thermal forecasting. Replacing the target-specific dual heads with a shared output head and removing adaptive graph propagation also reduced overall accuracy. These results indicate that future control conditioning, target-specific prediction heads, temporal decoding, and inter-AHU dependency modeling jointly contribute to the predictive performance of the full model.

In addition, the node-level performance comparison showed that prediction errors tended to increase in zones with higher operational variability, such as those characterized by large occupancy fluctuations or stronger outdoor influences (refer to Table 5).

Overall, the AGCRN-based IEQ prediction model successfully captured the spatiotemporal dependencies inherent in the multiple AHUs and provided reliable predictions of IEQ variations. By enabling the indoor environmental responses to generated operation scenarios to be evaluated in advance, the model can be used as a surrogate model to screen feasible candidate scenarios during the constraint-based operation strategy selection process.

### 3.3. Energy consumption forecasting performance

In this study, an AHU sensor-based time series regression model was developed to predict building-level energy consumption from multiple AHUs operational states, and its predictive performance was evaluated

**Table 5**AHU-wise prediction performance for indoor air temperature and CO<sub>2</sub> concentration.

| AHU   | Indoor air temperature (MAPE %) | Indoor CO <sub>2</sub> concentration (MAPE %) |
|-------|---------------------------------|-----------------------------------------------|
| AHU01 | 1.500                           | 9.017                                         |
| AHU02 | 1.747                           | 6.624                                         |
| AHU03 | 1.768                           | 6.937                                         |
| AHU04 | 2.041                           | 3.557                                         |
| AHU05 | 1.377                           | 8.399                                         |
| AHU06 | 1.866                           | 5.332                                         |
| AHU07 | 2.529                           | 6.116                                         |
| AHU08 | 1.354                           | 3.982                                         |
| AHU09 | 1.801                           | 2.920                                         |
| AHU10 | 3.281                           | 4.406                                         |

on the test set.

Fig. 6 presents a comparison between the actual energy consumption and the model predictions over the test period. Overall, the model reproduced both daily periodic variations and longer-term consumption trends with good agreement. Notably, the predicted results closely tracked the recurring peak patterns associated with AHU operation, indicating that the model effectively captured the temporal dynamics of building energy consumption.

Quantitative performance was assessed using mean absolute error (MAE), root mean squared error (RMSE), and mean absolute percentage error (MAPE). The model achieved an MAE of 0.476 Wh/m<sup>2</sup>, an RMSE of 0.681 Wh/m<sup>2</sup>, and a MAPE of 9.414%. These results demonstrate that the proposed time series-based prediction model, which jointly incorporates AHU operational variables and outdoor conditions, can provide reliable estimates of building energy consumption variations.

Previous studies on building energy consumption prediction have generally reported short-term forecasting accuracy of ML- and DL-based models within a MAPE range of approximately 10–20%. Table 6 compares the predictive performance of representative prior studies with that of the proposed model. Somu et al. [39] reported a MAPE of approximately 12–15% using a CNN-LSTM-based framework for building energy consumption prediction. Gao et al. [40] proposed an interpretable deep learning model with an attention mechanism and achieved prediction errors of approximately 10–13%. Jiao et al. [41] reported a MAPE range of approximately 10–18% using a hybrid machine learning approach based on modal decomposition and reconstruction. Zhang et al. [42] presented a Bi-LSTM-based short-term electricity consumption forecasting model incorporating building–occupant interaction variables and reported errors in the range of approximately 11–16%.

In comparison, the proposed GRU Seq2Seq model achieved a MAPE of 9.804%, outperforming the error ranges reported in these previous studies. This result suggests that the time series input configuration used in this study, which jointly captures multiple AHUs operational states and outdoor conditions, is effective for representing the temporal variability of building-level energy consumption.

To further examine the contribution of the residual learning structure and input feature groups, an ablation study was conducted for the building-level energy consumption forecasting model. Table 7 summarizes the prediction performance of the full model and ablated configurations. Compared with direct absolute energy prediction, residual learning improved the prediction accuracy, reducing MAE from 0.490 to 0.476 Wh/m<sup>2</sup> and RMSE from 0.691 to 0.681 Wh/m<sup>2</sup>. Among the input feature groups, removing time features caused the largest performance degradation, increasing MAE by 11.8%, followed by removing future AHU control variables, which increased MAE by 6.6%. This indicates that daily/weekly temporal patterns and AHU operation schedules are important factors in estimating building-level energy response. In contrast, removing AHU history or weather variables caused relatively smaller changes in MAE. Overall, the ablation results suggest that residual learning provides a consistent improvement, while the inclusion of future AHU control variables confirms that the energy surrogate reflects changes in candidate AHU operation scenarios.

### 3.4. Results of constraint-based AHU operational strategy selection

To evaluate the practical applicability of the proposed constraint-based operation strategy selection framework, the performance of candidate strategies satisfying indoor environmental constraints was analyzed among the generated AHU operation scenarios. Throughout this section, the comparison baseline refers to the facility's existing BAS operation implemented during the corresponding monitoring period. For each week, the strategy that simultaneously satisfied the indoor environmental constraints and the energy objective was selected as the final operation scenario.

The results indicate that indoor environmental conditions were




Fig. 6. Comparison of predicted and actual energy consumption on the test set.

Table 6

Comparison of MAPE values for building energy consumption prediction models.

| Study             | Model          | MAPE   |
|-------------------|----------------|--------|
| Somu et al. [39]  | CNN-LSTM       | 12–15% |
| Gao et al. [40]   | Attention LSTM | 10–13% |
| Jiao et al. [41]  | Hybrid ML      | 10–18% |
| Zhang et al. [42] | Bi-LSTM        | 11–16% |
| This study        | GRU Seq2Seq    | 9.414% |

Table 7

Ablation study for energy consumption prediction model.

| Mode                        | MAE (Wh/m <sup>2</sup> ) | RMSE (Wh/m <sup>2</sup> ) | MAPE (%) |
|-----------------------------|--------------------------|---------------------------|----------|
| Full model                  | 0.476                    | 0.681                     | 9.414    |
| Without residual learning   | 0.49                     | 0.691                     | 9.821    |
| Without ahu history pattern | 0.482                    | 0.687                     | 9.578    |
| Without weather input       | 0.484                    | 0.711                     | 9.243    |
| Without AHU control input   | 0.507                    | 0.731                     | 9.928    |
| Without time input          | 0.532                    | 0.778                     | 10.421   |

Table 8

Weekly comparison of indoor air temperature and CO<sub>2</sub> concentration between the selected optimal scenario and the actual operation.

| Week     | Indoor air temperature, selected (°C) | Indoor air temperature, actual (°C) | AHU-level mean range, selected (°C) | CO <sub>2</sub> concentration, selected (ppm) | CO <sub>2</sub> concentration, actual (ppm) | AHU-level mean range, selected (ppm) |
|----------|---------------------------------------|-------------------------------------|-------------------------------------|-----------------------------------------------|---------------------------------------------|--------------------------------------|
| Week #20 | 23.34 ± 0.70                          | 23.36 ± 0.64                        | 22.59–24.15                         | 888.7                                         | 897.2                                       | 680.95–1345.61                       |
| Week #21 | 23.48 ± 0.74                          | 23.60 ± 0.65                        | 22.46–24.61                         | 868.3                                         | 874.4                                       | 668.58–1308.96                       |
| Week #28 | 25.07 ± 0.86                          | 24.97 ± 0.85                        | 23.72–26.30                         | 854.5                                         | 867.6                                       | 689.18–1107.12                       |
| Week #31 | 25.89 ± 0.97                          | 25.79 ± 0.92                        | 24.09–26.84                         | 863.3                                         | 881.8                                       | 699.70–1100.27                       |
| Week #33 | 26.30 ± 1.11                          | 26.19 ± 1.01                        | 24.05–27.43                         | 836.4                                         | 845.1                                       | 604.74–1117.77                       |
| Week #34 | 26.44 ± 1.16                          | 26.33 ± 1.06                        | 24.14–27.82                         | 825.0                                         | 837.1                                       | 657.68–1061.70                       |
| Week #38 | 25.26 ± 1.23                          | 25.04 ± 1.16                        | 24.01–26.50                         | 810.7                                         | 834.8                                       | 639.90–1052.70                       |
| Week #40 | 25.48 ± 1.19                          | 25.52 ± 0.99                        | 24.54–26.46                         | 779.0                                         | 786.3                                       | 580.54–909.25                        |
| Week #43 | 24.15 ± 0.76                          | 24.12 ± 0.81                        | 23.11–25.01                         | 815.3                                         | 825.2                                       | 617.53–1028.83                       |
| Week #44 | 23.55 ± 1.24                          | 23.72 ± 1.28                        | 21.84–24.54                         | 829.0                                         | 840.9                                       | 637.58–1105.85                       |

maintained in a stable manner across all weeks. During occupied hours, indoor air temperature remained within a range nearly identical to that observed under actual operation, suggesting that the thermal environment was not materially degraded even though airflow rates and operating durations were adjusted during the optimization process. As shown in Table 8, the mean indoor air temperature ranged from approximately 23.3 to 26.4 °C, and the difference between the optimal operation scenario and actual operation was less than 0.2 °C in most weeks. The temporal variability of indoor air temperature also remained comparable to that of actual operation, with standard deviations ranging from approximately 0.7 to 1.2 °C. These findings indicate that the optimized AHU operation strategy effectively preserved indoor thermal stability.

Fig. 7 illustrates the distribution of weekly return-air temperature across individual AHUs for the selected scenarios compared with actual operation, together with the reference temperature threshold. Fig. 8 presents the corresponding distribution of CO<sub>2</sub> concentration. As shown, most AHUs maintained conditions comparable to actual operation and within the reference range; however, some AHUs (e.g., AHU02 for indoor air temperature, AHU01 for CO<sub>2</sub> concentration) exhibited occasional excursions beyond the threshold under both actual and selected




Fig. 7. AHU-level indoor air temperature, actual vs. selected (Threshold: 27 °C).




Fig. 8. AHU-level CO<sub>2</sub> concentration, actual vs. selected (Threshold: 1200 ppm).

operation, reflecting the spatial heterogeneity inherent in the building's zone-level occupancy and thermal loads.

To examine whether the proposed CVAE-based scenario selection could be replaced by simpler scheduling strategies, an additional comparison with rule-based baseline methods was conducted. The baseline strategies included fixed occupied-hour operation (weekdays, 08:00–18:00), rule-based schedule shifting ( $\pm 2$  h relative to the actual operation schedule), operation-duration reduction ( $-10\%$  and  $-20\%$ ), and a historical best schedule, defined as the actual schedule with the lowest energy consumption during the previous four weeks. All baseline strategies were evaluated using the same energy and IEQ surrogate models, occupied-hour definition, and constraint-based selection procedure as the proposed framework to ensure a fair comparison.

The mean indoor temperature was similar across all methods (approximately 24.9 °C), indicating that average indoor temperature alone was insufficient to distinguish the IEQ performance of different scheduling strategies. However, clear differences emerged when constraint satisfaction was evaluated, as summarized in Table 9. The proposed CVAE-selected scenarios met the adopted comparative IEQ screening criteria in all ten evaluation weeks, whereas none of the baseline strategies did so in more than two of the ten weeks including the historical-best schedule, despite its achieving the highest average estimated energy saving (5.13%) among all baselines.

Table 9

Comparison with baseline control scenarios.

| Method                             | Mean saving (%) | Mean indoor air temperature (°C) | Mean CO <sub>2</sub> concentration (ppm) | Satisfied weeks (n) |
|------------------------------------|-----------------|----------------------------------|------------------------------------------|---------------------|
| CVAE selected                      | 4.52            | 24.92                            | 837.00                                   | 10                  |
| Fixed occupied hours (08:00–18:00) | 3.72            | 24.91                            | 843.39                                   | 2                   |
| Schedule shift $-2$ h              | 3.33            | 24.98                            | 855.60                                   | 1                   |
| Schedule shift $+2$ h              | 2.94            | 24.92                            | 845.85                                   | 0                   |
| Duration $-20\%$                   | 4.86            | 24.98                            | 846.08                                   | 2                   |
| Duration $-10\%$                   | 4.03            | 24.98                            | 849.88                                   | 1                   |
| Previous-4-week best               | 5.13            | 24.96                            | 840.42                                   | 2                   |

These results suggest that selecting historically low-energy schedules or uniformly adjusting operating schedules can provide apparent energy-saving potential, but do not consistently meet the adopted comparative IEQ screening criteria at the AHU level. In contrast, the proposed CVAE-based framework first screens candidate schedules according to the adopted comparative IEQ criteria and then selects the

schedule with the lowest predicted energy consumption from the feasible candidate set. Consequently, the proposed framework demonstrated more consistent performance in meeting the adopted IEQ screening criteria while maintaining comparable or higher surrogate-estimated energy-saving potential (refer to Table 9).

A similar pattern was observed for IAQ. The mean CO<sub>2</sub> concentration remained comparable to or lower than that under actual operation across all weeks, with reductions typically ranging from several ppm to several tens of ppm. For example, the mean CO<sub>2</sub> concentration decreased from 897.2 ppm to 888.7 ppm in Week #20 and from 834.8 ppm to 810.7 ppm in Week #38. These results indicate that the optimized operation strategy was able to maintain IAQ within the range observed under actual operation, even while adjusting ventilation rates. Taken together, the proposed strategy represents a feasible operational solution that reduces energy consumption while maintaining both indoor air temperature and CO<sub>2</sub> concentration at levels comparable to those observed under actual operation.

Based on this result, the energy-saving performance of the operation strategy with the minimum energy consumption among the constraint-satisfying scenarios was further analyzed. Table 10 presents the weekly energy-saving performance obtained when selecting, among the scenarios that satisfied the allowable ranges of indoor air temperature and CO<sub>2</sub> concentration, the strategy with the lowest energy consumption. The results show that, in the weeks where energy savings were achieved, the reduction rate ranged from 0.9% to 15.4% relative to actual operation, with most weeks exhibiting stable savings of approximately 2–6%. In particular, Week #38 (08.25–08.31) showed the highest reduction rate of 15.4%, corresponding to a total energy saving of 1034 Wh/m<sup>2</sup> (60,737.6 kWh) and an estimated electricity cost saving of \$5819.2. In contrast, Week #44 (10.06–10.12) exhibited the lowest reduction rate of 0.9%, likely because relatively mild outdoor conditions limited the additional savings achievable through optimization.

Overall, when aggregated across the weeks in which reductions were observed, the proposed strategy yielded a total annual energy saving of 2422.5 Wh/m<sup>2</sup> (142,304.1 kWh), corresponding to an electricity cost saving of \$13,634.1. These findings demonstrate that the proposed constraint-based operation strategy can deliver meaningful energy and cost savings while preserving indoor environmental conditions.

#### 4. Discussion

To provide a more systematic interpretation of the optimal AHU operation scenarios derived from the constraint-based strategy, the weeks in which energy savings were observed were classified according to outdoor conditions and AHU load characteristics. Fig. 9 illustrates the spatial distribution of changes in supply airflow rates between the actual and optimal operation scenarios for each week.

The results show that, in most weeks, the observed patterns were not characterized by uniform reductions in airflow across all AHUs. Instead, a supply air redistribution pattern was identified, in which airflow rates increased in some AHUs while decreasing in others. This indicates that the constraint-based strategy operates by balancing airflow distribution at the building level, rather than simply reducing airflow in individual AHUs.

In particular, during cooling and transition periods (Weeks #20, #21, #28, #31, #33, and #34), a general reduction in supply airflow was observed across most AHUs (refer to Fig. 9), while some AHUs exhibited stable or increased airflow. Similar patterns were observed at the daily scale. During morning hours, when outdoor temperatures were relatively low, certain AHUs maintained or increased airflow to take advantage of favorable conditions. In contrast, during afternoon hours, rising outdoor temperatures increased cooling loads associated with outdoor air intake, leading to reduced airflow in most AHUs. In this process, ventilation rates were reduced in some zones while maintaining CO<sub>2</sub> concentration within acceptable limits, suggesting increased reliance on internal air recirculation.

**Table 10**

Weekly energy saving performance of constraint-based AHU operation strategies.

| Week                     | Estimated energy saving ratio (%) | 95% Confidence Interval (%) | Total energy saving potential (Wh/m <sup>2</sup> ) | Total energy saving potential (kWh) | Electricity cost saving potential (\$) |
|--------------------------|-----------------------------------|-----------------------------|----------------------------------------------------|-------------------------------------|----------------------------------------|
| Week #20 (04.21 – 04.27) | 3.3                               | [1.65, 5.23]                | 143.2                                              | 8412.4                              | 805.9                                  |
| Week #21 (04.28 – 05.04) | 3.5                               | [1.88, 5.53]                | 159.3                                              | 9359.2                              | 896.7                                  |
| Week #28 (06.16 – 06.22) | 3.6                               | [2.05, 5.27]                | 169.4                                              | 9953                                | 953.6                                  |
| Week #31 (07.07 – 07.13) | 5.9                               | [4.47, 7.48]                | 294.1                                              | 17,275                              | 1655.1                                 |
| Week #33 (07.21 – 07.27) | 3.4                               | [2.01, 5.16]                | 167.5                                              | 9840.4                              | 942.8                                  |
| Week #34 (07.28 – 08.03) | 4.6                               | [3.15, 6.24]                | 227.7                                              | 13,372.9                            | 1281.2                                 |
| Week #38 (08.25 – 08.31) | 15.4                              | [14.35, 16.71]              | 1034.0                                             | 60,737.6                            | 5819.2                                 |
| Week #40 (09.08 – 09.14) | 1.7                               | [0.11, 3.69]                | 70.8                                               | 4159.8                              | 398.5                                  |
| Week #43 (09.29 – 10.05) | 2.5                               | [0.91, 4.30]                | 114.8                                              | 6744                                | 646.1                                  |
| Week #44 (10.06 – 10.12) | 0.9                               | [−0.78, 2.84]               | 41.7                                               | 2449.8                              | 234.7                                  |
| Total (Year)             | 0.9 - 15.4                        |                             | 2422.5                                             | 142,304.1                           | 13,634.1                               |

A different pattern was observed during the late transition period (Weeks #43 and #44). In these weeks, lower outdoor temperatures during morning hours led to increased airflow in certain zones to meet heating demands. Conversely, in the afternoon, as outdoor temperatures increased, the demand for thermal conditioning decreased, resulting in reduced airflow or partial shutdown of AHU operation in many zones. This indicates that, under transition-season conditions, the operation strategy dynamically adjusts heating and cooling demands in response to time-varying outdoor temperatures.

Although the proposed framework is compatible with existing BAS, its primary role is not to replace the current control logic with a fully autonomous controller. Instead, it is designed to support facility managers by providing evidence-based operational recommendations derived from historical building operation and predictive scenario



Fig. 9. Weekly AHU supply air intake changes between optimal and actual scenario.

evaluation. In the target building, AHUs are already managed through the BAS by facility operators who routinely adjust schedules according to occupant requests, building events, and other operational requirements. Within this workflow, the proposed framework can periodically identify alternative AHU operation strategies that maintain indoor environmental quality while reducing energy consumption relative to the current operating practice. These recommendations can be incorporated into routine BAS operation as decision-support information, allowing facility managers to retain operational flexibility while systematically improving energy performance. Consequently, the proposed framework can be deployed without substantial modification of the existing BAS infrastructure, facilitating practical adoption in real building operation.

To provide a more systematic interpretation of the observed supply air redistribution patterns, the weeks in which energy savings were achieved were classified into five operational archetypes, denoted as Types A to E. The results indicate that weekly energy-saving patterns exhibited distinct characteristics depending on seasonal conditions and outdoor environments, suggesting that the underlying energy-saving mechanisms in the multiple AHUs vary with outdoor conditions, ventilation demands, space functions, and inter-AHU interactions. The multiple AHUs operational patterns and corresponding energy-saving mechanisms for each archetype are described below.

- **Type A (Early cooling transition):** Weeks #20 and #21 correspond to the early cooling transition period, before cooling loads begin to rise substantially. During this period, outdoor temperatures remain relatively mild, resulting in limited cooling demand. However, under actual operation, relatively high ventilation airflow tends to be maintained to ensure IAQ stability. The airflow analysis showed that, compared with actual operation, supply airflow was reduced in AHU03 (Conference Hall I), AHU05 (Conference Hall II), and AHU08 (Alumni Hall). These spaces are event-oriented zones characterized

by highly variable occupancy densities, and thus conservative ventilation strategies are typically applied in practice. In contrast, the proposed operation strategy re-evaluated ventilation requirements based on CO<sub>2</sub> concentration prediction and reduced unnecessary ventilation airflow. In addition, a balanced airflow adjustment pattern was observed in the central core zones served by AHU04–AHU06, where multiple AHUs simultaneously exhibited modest airflow reductions while accounting for thermal interactions among adjacent spaces. This redistribution of supply airflow reduced both fan power and cooling loads associated with outdoor air intake, resulting in an energy saving of approximately 6%.

- **Type B (Peak cooling season):** Weeks #28 and #31 correspond to the peak cooling season, during which cooling loads increase with rising outdoor temperatures. In this period, aggressive airflow reduction strategies may increase the risk of violating IEQ constraints due to high cooling demand across most AHUs. The airflow pattern analysis revealed that AHU07 (Grand Ballroom) and AHU10 (Lobby) have dominant influence on building-level energy consumption, supplying large airflow rates on the order of 30,000–40,000 m<sup>3</sup>/h. Accordingly, the optimal operation strategy maintained the baseline operation of these large-capacity AHUs, while mitigating peak airflow levels in medium-scale AHUs (e.g., AHU01, and AHU06) and adjusting their operating durations. This reflects a hierarchical airflow control strategy, in which the primary operation of large AHUs is preserved while fine-grained adjustments are applied to medium-scale units. Through this approach, energy savings of approximately 5–6% were achieved even during peak cooling conditions.
- **Type C (Post-peak cooling decline):** Weeks #33 and #34 correspond to the period following peak summer conditions, during which outdoor temperatures gradually decrease and cooling demand declines. During this period, the high ventilation rates maintained under actual operation tend to exceed the required ventilation demand,

leading to over-ventilation. The airflow analysis showed that supply airflow was reduced in AHU05 (Conference Hall II) and AHU08 (Alumni Hall) compared with actual operation. These zones, characterized by high variability in occupancy, are typically operated with conservative ventilation strategies. However, the optimized strategy re-evaluated ventilation requirements based on CO<sub>2</sub> concentration prediction and reduced unnecessary airflow. In addition, a coordinated airflow redistribution pattern was observed in the central core AHUs (AHU04–AHU06), where multiple units simultaneously exhibited modest airflow reductions. This cooperative multiple AHUs operation reduced both fan energy and cooling loads, resulting in an energy saving of approximately 6.8% in Week #34.

- **Type D (Transitional season):** Week #38 corresponds to a mild transitional season, during which outdoor temperatures approach indoor set point levels. This period exhibited the highest energy-saving potential in this study. As both cooling demand and ventilation requirements decrease, the high airflow maintained under actual operation tends to result in over-ventilation. The airflow analysis showed simultaneous reductions in multiple AHUs, including reduced airflow in AHU07, shortened operating durations in AHU03 and AHU05, and partial airflow reduction in AHU08. This indicates a system-wide airflow reduction pattern, involving both airflow magnitude and operating duration adjustments across multiple AHUs. In particular, reductions in large-capacity units such as AHU07 and AHU10, which serve extensive areas and deliver high airflow rates, had a substantial impact on overall energy consumption. As a result, a maximum energy saving of approximately 15.4% was achieved during this period.
- **Type E (Heating transition):** Weeks after #40 correspond to the heating transition or early heating season, during which outdoor temperatures gradually decrease. During this period, excessive outdoor air intake can increase heat loss, making the suppression of unnecessary ventilation and reduction of operating duration key strategies for energy savings. However, the extent to which outdoor air intake can be reduced is inherently constrained by the need to maintain indoor CO<sub>2</sub> concentration below the predefined threshold; unlike the cooling season, where surplus ventilation capacity is available for reduction without approaching the IAQ constraint, heating-season ventilation is already operated closer to the minimum level required for acceptable air quality, leaving comparatively little

room for further reduction without risking constraint violation. The airflow analysis showed that supply airflow in most AHUs remained similar to that under actual operation; however, some AHUs exhibited shortened operating durations, increased part-load operation, and selective operation focused on occupied zones. In particular, large-space AHUs such as AHU07 and AHU10, which handle significant outdoor air volumes, can reduce system energy consumption even through modest reductions in operating time. Nevertheless, due to the already low cooling demand, and the limited margin for further ventilation reduction under the IAQ constraint, the additional energy-saving potential was limited, resulting in reductions of approximately 2–3%. This seasonal constraint, rather than split membership in the surrogate models' training procedure, explains why the held-out test weeks (#40, #43, #44) show comparatively low estimated savings.

In this study, weekly energy-saving patterns were categorized into five operational archetypes, and practical guidelines for multiple AHUs operation strategies were derived based on the corresponding multiple AHUs operational mechanisms, as illustrated in Fig. 10.

Overall, the energy-saving potential of the multiple AHUs exhibited clear seasonal variations. In particular, during the early heating transition period (Type E), only limited operational strategies were identified that could simultaneously satisfy indoor environmental constraints while further reducing energy consumption. In fact, weeks corresponding to the full heating season, characterized by lower outdoor temperatures, were not selected as optimal scenarios. This suggests that, during this period, viable operational alternatives that simultaneously satisfy indoor air temperature and CO<sub>2</sub> concentration constraints while reducing energy consumption relative to actual operation are highly limited.

These findings are closely related to the inherent characteristics of AHU operation under heating conditions. During the heating season, increased outdoor air intake leads to higher heat losses and consequently greater heating demand. Conversely, excessive reduction of ventilation airflow results in elevated indoor CO<sub>2</sub> concentration, making it difficult to satisfy IAQ constraints. This reflects a strong trade-off between minimizing outdoor air intake for energy savings and maintaining sufficient ventilation to ensure acceptable IAQ.

Accordingly, the energy-saving potential of multiple AHUs was



Fig. 10. Seasonal operating types and practical guidelines for multiple AHUs operation.

found to be greater during cooling periods and cooling transition seasons than during heating periods. In particular, the highest savings potential was observed during the transitional period from summer to autumn, when both cooling demand and ventilation requirements decrease. The present framework should be regarded as a historically grounded schedule-screening approach rather than a complete AHU optimal control solution. The generated actions are limited to daily operating status, start time, and duration, with each daily schedule simplified as one contiguous operation block; continuous control actions such as airflow modulation, economizer operation, supply-air temperature reset, and demand-controlled ventilation are outside the scope of the present implementation. This scope was chosen because the case-study building operates CAV AHUs with a fixed supply-air temperature set point, making on/off scheduling the dominant controllable variable available for data-driven evaluation. Because the selected schedules were evaluated using offline surrogate models rather than through real-building implementation, the reported energy reductions should be interpreted as predicted saving potential rather than confirmed operational savings. To account for surrogate-model uncertainty, the residual distribution from the held-out energy-prediction test set was used to construct bootstrap confidence intervals for the predicted weekly savings; weeks for which the confidence interval includes zero are interpreted as directional rather than statistically distinguishable savings. Future work should further validate the proposed framework through real-building operation, virtual commissioning, or HVAC-level submetered evaluation.

## 5. Conclusions

In this study, a data-driven, scenario-based decision-support framework was proposed for evaluating the operation of multiple AHUs in a large-scale building. The framework integrates AHU operation scenario generation, indoor environmental prediction, energy consumption prediction, and constraint-aware scenario screening. Specifically, a CVAE was used to generate diverse and historically plausible AHU operation scenarios, while an AGCRN-based IEQ prediction model and a GRU-based energy consumption prediction model were employed to evaluate the potential impacts of each scenario. Among the generated scenarios, the candidate that satisfied the predefined indoor temperature and CO<sub>2</sub> constraints and showed the lowest predicted cumulative energy consumption was identified as the preferred operational candidate.

The offline case-study evaluation showed that the selected candidate scenarios had the potential to reduce energy consumption while maintaining indoor air temperature and CO<sub>2</sub> concentration at levels comparable to those observed under actual operation. In most evaluated periods, the predicted energy-saving potential was approximately 2–6%. The highest potential was observed during transitional seasons, when reduced cooling demand allowed system-wide airflow reduction and operating time adjustments across multiple AHUs. In contrast, during heating transition periods, the trade-off between ventilation requirements and heat loss limited the potential for additional energy reduction.

A key contribution of this study is the explicit treatment of IEQ requirements as screening criteria rather than as terms that are freely traded off against energy savings. By first identifying scenarios that satisfy indoor temperature and CO<sub>2</sub> constraints and then comparing their predicted energy performance, the proposed framework reflects practical building operation priorities, where occupant comfort and indoor air quality should be maintained before energy-saving alternatives are considered.

In addition, the CVAE-based scenario generation approach enabled the exploration of diverse combinations of AHU operation schedules within the historical operating domain represented in the available data. By integrating predictive models for both indoor environmental conditions and building-level energy consumption, the framework supports the offline evaluation of candidate AHU schedules before

implementation. In this sense, the proposed approach can complement conventional rule-based or experience-driven AHU operation by providing data-driven evidence for comparing alternative operation scenarios.

Overall, this study presents a case-study-based predictive screening framework for multiple AHU operation using real building operation data. The identified seasonal archetypes should be interpreted as case-specific operational insights rather than general guidelines for AHU operation. The energy-saving mechanisms identified across these archetypes were primarily achieved through supply air redistribution and operating-duration adjustment among multiple AHUs, consistent with the schedule-based nature of AHU operation and BAS data in the case-study building. The proposed framework relies on data commonly available in BAS-equipped buildings, including AHU operational data, zone-level IEQ and occupancy measurements, outdoor weather data, building energy consumption, and the mapping between AHUs and their served zones, thereby providing a practical basis for reproducing the proposed workflow in buildings with comparable monitoring infrastructure. Nevertheless, the present study was validated using a single multi-AHU building, and establishing the replicability of the framework requires long-term validation across additional buildings with diverse operational characteristics. These findings suggest that scenario-based generation and predictive screening can support more informed AHU operation planning. As richer BAS instrumentation and continuously controllable variables become available, further validation in additional buildings, extension toward MPC or reinforcement learning-based control approaches, and real-time deployment studies represent promising directions to confirm and extend the broader applicability of the proposed framework. In addition, given the reliance on building-level rather than AHU-submetered energy data, incorporating direct HVAC-level submetering would further strengthen the causal attribution of the reported energy-saving potential.

## CRedit authorship contribution statement

**Seungkeun Yeom:** Conceptualization, Formal analysis, Methodology, Validation, Visualization, Writing – original draft. **Juui Kim:** Formal analysis, Validation, Visualization, Writing – original draft. **Jinwoo Choi:** Resources, Visualization, Writing – original draft. **Gaurav Kattel:** Methodology, Validation. **Anup Paudel:** Software, Visualization. **Jun-Ki Choi:** Investigation, Project administration, Writing – review & editing. **Taecheon Hong:** Conceptualization, Funding acquisition, Investigation, Methodology, Project administration, Supervision, Writing – review & editing.

## 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.

## Acknowledgements

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MIST; Ministry of Science and ICT) (NRF-2021R1A3B1076769).

## Data availability

Data will be made available on request.

## References

- [1] L. Pérez-Lombard, J. Ortiz, C. Pout, A review on buildings energy consumption information, *Energy Build.* 40 (2008) 394–398, <https://doi.org/10.1016/j.enbuild.2007.03.007>.

- [2] D. Ürge-Vorsatz, L.F. Cabeza, S. Serrano, C. Barreneche, K. Petrichenko, Heating and cooling energy trends and drivers in buildings, *Renew. Sustain. Energy Rev.* 41 (2015) 85–98, <https://doi.org/10.1016/j.rser.2014.08.039>.
- [3] H. Huang, B.R. Hughes, Review of HVAC forecasting and control strategies for improved building performance, *Build. Environ.* 287 (2026) 113797, <https://doi.org/10.1016/j.buildenv.2025.113797>.
- [4] A. Ameen, F.L. Rashid, M.A. Al-Obaidi, N.M.L. Al Maimuri, M.K. Abdalrahman, E. B. Agyekum, A. Chibani, M. Kezzar, Ventilation strategies in net-zero energy buildings: balancing indoor air quality and energy efficiency, *Energy Build.* 351 (2026) 116724, <https://doi.org/10.1016/j.enbuild.2025.116724>.
- [5] A. Afram, F. Janabi-Sharifi, Theory and applications of HVAC control systems - A review of model predictive control (MPC), *Build. Environ.* 72 (2014) 343–355, <https://doi.org/10.1016/j.buildenv.2013.11.016>.
- [6] M. Killian, M. Kozek, Ten questions concerning model predictive control for energy efficient buildings, *Build. Environ.* 105 (2016) 403–412, <https://doi.org/10.1016/j.buildenv.2016.05.034>.
- [7] H. Zhang, C. Li, Q. Wei, Y. Zhang, Fault detection and diagnosis of the air handling unit via combining the feature sparse representation based dynamic SFA and the LSTM network, *Energy Build.* 269 (2022) 112241, <https://doi.org/10.1016/j.enbuild.2022.112241>.
- [8] H. Zhang, H. Zhao, Y. Fu, J. Ma, Y. Xiang, The class labels and spatial information based fault diagnosis of air handling unit via combining kernel Fischer discriminant analysis with an improved graph convolutional neural network, *Meas. J. Int. Meas. Confed.* 257 (2026) 118622, <https://doi.org/10.1016/j.measurement.2025.118622>.
- [9] H. Zhang, Y. Zhang, H. Meng, J.B. Lim, W. Yang, A novel global modelling strategy integrated dynamic kernel canonical variate analysis for the air handling unit fault detection via considering the two-directional dynamics, *J. Build. Eng.* 96 (2024) 110402, <https://doi.org/10.1016/j.jobe.2024.110402>.
- [10] S.J. Ra, J.H. Kim, C.S. Park, Real-time model predictive cooling control for an HVAC system in a factory building, *Energy Build.* 285 (2023) 112860, <https://doi.org/10.1016/j.enbuild.2023.112860>.
- [11] M. Soleimani, F.N. Irani, M. Yadegar, M. Davoodi, Multi-objective optimization of building HVAC operation: advanced strategy using Koopman predictive control and deep learning, *Build. Environ.* 248 (2024) 111073, <https://doi.org/10.1016/j.buildenv.2023.111073>.
- [12] J. Zhang, F. Xiao, A. Li, T. Ma, K. Xu, H. Zhang, R. Yan, X. Fang, Y. Li, D. Wang, Graph neural network-based spatio-temporal indoor environment prediction and optimal control for central air-conditioning systems, *Build. Environ.* 242 (2023) 110600, <https://doi.org/10.1016/j.buildenv.2023.110600>.
- [13] C. Fu, H. Kazmi, M. Quintana, C. Miller, Creating synthetic energy meter data using conditional diffusion and building metadata, *Energy Build.* 312 (2024) 114216, <https://doi.org/10.1016/j.enbuild.2024.114216>.
- [14] J. Song, T. Bian, G. Xue, H. Wang, X. Shen, X. Wu, Short-term forecasting model for residential indoor temperature in DHS based on sequence generative adversarial network, *Appl. Energy* 348 (2023) 121559, <https://doi.org/10.1016/j.apenergy.2023.121559>.
- [15] A. Parisio, M. Molinari, D. Varagnolo, K.H. Johansson, A scenario-based predictive control approach to building HVAC management systems, *IEEE Int. Conf. Autom. Sci. Eng.* (2013), <https://doi.org/10.1109/CoASE.2013.6654024>.
- [16] C. Huang, H.H.L. Kwok, K.H. Poon, Z. Wu, F. Hou, J. Ma, J.C.P. Cheng, Graph-based spatial-temporal prediction and feature interaction analysis of CO<sub>2</sub> and occupant in large indoor space, *Build. Environ.* 280 (2025) 112963, <https://doi.org/10.1016/j.buildenv.2025.112963>.
- [17] S. Wen, W. Zhang, N. Zhou, L. Yuan, Adaptive spatio-temporal graph convolutional network for rapid indoor temperature field prediction with limited sensors, *Build. Environ.* 283 (2025) 113346, <https://doi.org/10.1016/j.buildenv.2025.113346>.
- [18] T. Zhao, B. Jiang, Y. Li, Y. Rezgui, C. Zhang, P. Wang, Multi-point temperature or humidity prediction for office building indoor environment based on CGC-BiLSTM deep neural network, *Build. Environ.* 267 (2025) 112259, <https://doi.org/10.1016/j.buildenv.2024.112259>.
- [19] Y. Zhao, W. Li, J. Zhang, C. Jiang, S. Chen, Real-time energy consumption prediction method for air-conditioning system based on long short-term memory neural network, *Energy Build.* 298 (2023) 113527, <https://doi.org/10.1016/j.enbuild.2023.113527>.
- [20] Yonsei Energy Platform Service. <https://yeps.yonsei.ac.kr/>, 2025. <https://yeps.yonsei.ac.kr/>, accessed March 3, 2025.
- [21] D.P. Kingma, M. Welling, Auto-encoding variational bayes, in: 2nd int. Conf. Learn. Represent. ICLR 2014 - conf, Track Proc. 2 (2014), <https://doi.org/10.61603/ceas.v2i1.33>.
- [22] K. Sohn, X. Yan, H. Lee, Learning structured output representation using deep conditional generative models, *Adv. Neural Inf. Process. Syst.* (2015).
- [23] T. Gneiting, A.E. Raftery, Strictly proper scoring rules, prediction, and estimation, *J. Am. Stat. Assoc.* 102 (2007) 359–378, <https://doi.org/10.1198/016214506000001437>.
- [24] E. Betzalel, C. Penso, E. Fetaya, Evaluation metrics for generative models: an empirical study, *Mach. Learn. Knowl. Extr.* 6 (2024) 1531–1544, <https://doi.org/10.3390/make6030073>.
- [25] F. Guo, R. Xie, B. Huang, A deep learning just-in-time modeling approach for soft sensor based on variational autoencoder, *Chemom. Intell. Lab. Syst.* 197 (2020) 103922, <https://doi.org/10.1016/j.chemolab.2019.103922>.
- [26] M.C. Ruiz-Abellón, L.A. Fernández-Jiménez, A. Guilmón, A. Gabaldón, Applications of probabilistic forecasting in demand response, *Appl. Sci.* 14 (2024) 9716, <https://doi.org/10.3390/app14219716>.
- [27] S. Liu, X. Zhang, J. Wangni, J. Shi, Normalized diversification, *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern. Recognit.* (2019), <https://doi.org/10.1109/CVPR.2019.01055>.
- [28] D. Friedman, A.B. Dieng, The Vendi Score: a diversity evaluation metric for machine learning, *Trans. Mach. Learn. Res.* (2023).
- [29] C. Xiao, P. Zhong, C. Zheng, Bourgan: generative networks with metric embeddings, *Adv. Neural Inf. Process. Syst.* (2018).
- [30] L. Bai, L. Yao, C. Li, X. Wang, C. Wang, Adaptive graph convolutional recurrent network for traffic forecasting, *Adv. Neural Inf. Process. Syst.* (2020).
- [31] N. Mounir, H. Ouadi, I. Jhrilifa, Short-term electric load forecasting using an EMD-Bi-LSTM approach for smart grid energy management system, *Energy Build.* 288 (2023) 113022, <https://doi.org/10.1016/j.enbuild.2023.113022>.
- [32] S. Naqvi, K. Kar, S. Bhattacharya, V. Chandan, S. Mishra, T. Salisbury, Air quality and comfort constrained energy efficient operation of multi-zone buildings, *Build. Environ.* 244 (2023) 110716, <https://doi.org/10.1016/j.buildenv.2023.110716>.
- [33] S. Guo, Y. Lin, N. Feng, C. Song, H. Wan, Attention based spatial-temporal graph convolutional networks for traffic flow forecasting, in: 33rd AAAI Conference on Artificial Intelligence AAAI 2019, 31st Innovative Applications of Artificial Intelligence Conference IAAI 2019 9th AAAI Symposium on Educational Advances in Artificial Intelligence 33, AAAI, 2019, <https://doi.org/10.1609/aaai.v33i01.3301922>.
- [34] W. Weng, J. Fan, H. Wu, Y. Hu, H. Tian, F. Zhu, J. Wu, A decomposition dynamic graph convolutional recurrent network for traffic forecasting, *Pattern. Recognit.* 142 (2023) 109670, <https://doi.org/10.1016/j.patcog.2023.109670>.
- [35] Z. Wu, S. Pan, G. Long, J. Jiang, X. Chang, C. Zhang, Connecting the dots: multivariate time series forecasting with graph Neural networks, in: Proceeding ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2024, <https://doi.org/10.1145/3394486.3403118>.
- [36] F. Hamayat, M.O. Ullah, S. Zubair, S.M. Anwar, An indoor air temperature prediction framework for model predictive control in HVAC systems, in: 2022 27th International Conference on Industry of the Future and Smart Manufacturing ICAC 2022, 2022, <https://doi.org/10.1109/ICAC55051.2022.9911173>.
- [37] G. Yang, E. Yuan, Predicting the long-term CO<sub>2</sub> concentration in classrooms based on the BO-EMD-LSTM model, *Build. Environ.* 224 (2022) 109568, <https://doi.org/10.1016/j.buildenv.2022.109568>.
- [38] S. Taheri, A. Razban, Learning-based CO<sub>2</sub> concentration prediction: application to indoor air quality control using demand-controlled ventilation, *Build. Environ.* 205 (2021) 108164, <https://doi.org/10.1016/j.buildenv.2021.108164>.
- [39] N. Somu, G. Raman M R, K. Ramamritham, A deep learning framework for building energy consumption forecast, *Renew. Sustain. Energy Rev.* 137 (2021) 110591, <https://doi.org/10.1016/j.rser.2020.110591>.
- [40] Y. Gao, Y. Ruan, Interpretable deep learning model for building energy consumption prediction based on attention mechanism, *Energy Build.* 252 (2021) 111379, <https://doi.org/10.1016/j.enbuild.2021.111379>.
- [41] Y. Jiao, Z. Tan, D. Zhang, Q.P. Zheng, Short-term building energy consumption prediction strategy based on modal decomposition and reconstruction algorithm, *Energy Build.* 290 (2023) 113074, <https://doi.org/10.1016/j.enbuild.2023.113074>.
- [42] C. Zhang, L. Ma, Z. Luo, X. Han, T. Zhao, Forecasting building plug load electricity consumption employing occupant-building interaction input features and bidirectional LSTM with improved swarm intelligent algorithms, *Energy* 288 (2024) 129651, <https://doi.org/10.1016/j.energy.2023.129651>.