Driving Energy Efficiency at Scale by Mass Deployment of AI-based Chiller Energy Optimization

Exploring a Scalable Framework for AI-driven Chiller Energy Optimization: A Case Study on Mass Deployment in Hong Kong

K.T. Lau*
ktlau@emsd.gov.hk
Engineer, the Electrical and Mechanical Services Department, HKSAR
CHINA

K.F. LEUNG
leungkf@emsd.gov.hk
Engineer, the Electrical and Mechanical Services Department, HKSAR
CHINA

P.M. SO
pmso@emsd.gov.hk
Senior Engineer, the Electrical and Mechanical Services Department, HKSAR
CHINA

ABSTRACT

This paper presents a comprehensive case study on the mass deployment of an AI-driven platform for chiller energy optimization in chiller plants. It discusses the development and implementation of a scalable approach, emphasizing the significance of utilizing a semantic format for data representation and storage. The paper also addresses the process of selecting an appropriate AI model for chiller energy optimization and presents the results and performance metrics achieved through the platform's implementation. This case study serves as a valuable reference for organizations seeking to deploy AI-based energy optimization solutions in chiller plants at scale.

CCS CONCEPTS

• Computing methodologies → Artificial intelligence; **• Theory of computation → Semantics and reasoning**; **Machine learning theory**; **Data integration**.

KEYWORDS

Building energy efficiency, Energy modeling and simulation, AI-driven energy optimization, Smart buildings

ACM Reference Format:
K.T. Lau, K.F. LEUNG, and P.M. SO. 2024. Driving Energy Efficiency at Scale by Mass Deployment of AI-based Chiller Energy Optimization: Exploring a Scalable Framework for AI-driven Chiller Energy Optimization: A Case Study on Mass Deployment in Hong Kong. In *2024 the 8th International Conference on Innovation in Artificial Intelligence (ICIAI 2024), March 16–18, 2024, Tokyo, Japan.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3655497.3655503

*Corresponding Author

ICIAI 2024, March 16–18, 2024, Tokyo, Japan

ACM ISBN 979-8-4007-0930-2/24/03
https://doi.org/10.1145/3655497.3655503

1 INTRODUCTION

Climate change poses a global challenge, with human activities contributing to increased global warming and more frequent extreme weather events. Hong Kong is not immune to these effects, experiencing rising temperatures, a decrease in cold days, and more frequent heavy rainfall. To address these issues, the Government of the Hong Kong Special Administrative Region (HKSAR) of the People's Republic of China has established the Hong Kong Climate Action Plan 2050, which sets ambitious targets and strategies for decarbonization. One of the key targets is to reduce electricity consumption in commercial buildings by 30% to 40% and in residential buildings by 20% to 30% by 2050 [7]. Based on the Hong Kong Energy End-use Data for 2022, air conditioning emerged as a significant consumer of electricity, accounting for approximately 48,000 Terajoules [6]. The building industry is actively researching and implementing energy-saving measures, including chiller optimization, to align with the goals of reducing energy consumption.

The Electrical and Mechanical Services Department (EMSD) plays a crucial role in continuously improving energy efficiency standards and exploring innovative technologies. Leveraging artificial intelligence (AI) technology, the industry is exploring smart energy management solutions to optimize chiller plant operations and enhance energy efficiency in air conditioning design. Successful studies have demonstrated the potential and feasibility of this AI-driven application. Some of them gives more than 10% energy saving even for real data testing [2, 5, 8, 20].

However, little research has been done on the overall framework for digitizing an existing building for online data collection and AI model deployment in a chiller system. Additionally, existing studies mainly focus on developing specific models for individual buildings only [17]. The redeployment of one model from one building to another may not be directly transferable and compatible.

This paper presents the proposed methodologies for effective data acquisition, transmission, and storage, ensuring the availability of high-quality data for analysis and optimization purposes to facilitate the mass deployment of AI-driven Chiller Energy Optimization in Hong Kong. The paper emphasizes the significance of utilizing a semantic format for data representation and storage, enabling streamlined analysis and integration of AI algorithms and machine learning models. Furthermore, the paper addresses the crucial process of selecting an appropriate AI model for chiller energy optimization. It explores the tools and techniques that assist in identifying the most suitable model for different building types and applications.

2 DATA COLLECTION AND PREPARATION

Hong Kong has over 42,000 existing buildings, including industrial, domestic, commercial, private, and public buildings. The government owns over 8,000 buildings and facilities and is currently collecting over 600,000 data points per day per building to understand the performance of building services systems. To handle the large volume of data, the EMSD has established a framework for data collection.

2.1 Defining the standard format

To enable the mass deployment of chiller optimization, a standardized data type and format (Table 1)are required for all venues. The EMSD has set up an Integrated Building Managed System (iBMS) guideline that stipulates the minimal data points required for the Electrical, Mechanical, and Building Service (EMABS) System, including the Chiller system. The guideline outlines the data sampling frequency of every 15 minutes and the use of trend log function for data recording.

2.2 External Weather Data

The Hong Kong Observatory (HKO) operates over 50 Automatic Weather Stations across all districts in Hong Kong, providing meteorological data. This instant and historical weather data is available to the public and can be accessed via an Application Programming Interface (API). For chiller load prediction purposes, the following weather data (Table 2) is acquired from HKO for further study [8].

Table 2: Weather Data from HKO

Weather Data	Example	Data Format
Temperature (Celsius)	10.0	Real
Humidity (%)	64.8	Real
UV Index	0.5	Real
Rainfall (mm)	1.0	Real

2.3 Data collection infrastructure

Sensor data is collected locally on-site and stored in the Integrated Building Management System (iBMS) for local control and monitoring. The Regional Digital Control Centre (RDCC) consolidates the data from all buildings for central review, benchmarking, and data analysis. The RDCC utilizes AI to analyze the collected data for predictive maintenance, plant optimization, and energy efficiency. The dataflow is illustrated in the flow chart below (Figure 1).

2.4 Data Processing

Data processing is a crucial step that precedes data exploration and prediction in any data-driven analysis or modeling task. Raw data often requires preprocessing and transformation to ensure its quality, consistency, and compatibility with the analysis techniques employed. With domain knowledge, data cleansing and feature extraction are carried out. The major data processing steps are detailed in the Table 3 below.

Table 1: Standard Parameters for Chillers

System Parameter	Data format	Data Objectives
Off/On Status	Boolean	O&M
Auto/Manual Status	Enumerated	O&M
Trip/Fault Alarm	Boolean	O&M
Running Hour	Real	O&M, M&V
Running Amps	Real	O&M, M&V
Running Line Voltage	Real	O&M, M&V
Refrigerant Discharge Pressure	Real	O&M, M&V
Refrigerant Suction Pressure	Real	O&M, M&V
Compressor Off/On Status	Boolean	O&M, M&V
Chilled Water Supply Water Temperature	Real	O&M, M&V
Chilled Water Return Water Temperature	Real	O&M, M&V
Chilled Water Flow Rate	Real	O&M, M&V
Flow Switch Status	Boolean	O&M
Motorized value status	Real	O&M
Low Evaporating Pressure Cut-out	Boolean	O&M
High Condensing Pressure Cut-out	Boolean	O&M
Suction Air Temperature (Condenser)	Real	O&M, M&V
Discharge Air Temperature (Condenser)	Real	O&M, M&V
COP for Chiller & Chiller Plant	Real	O&M, M&V
COP for Plant	Real	O&M, M&V

Remarks: O&M – Operation and Maintenance;
M&V – Measurement and Verification

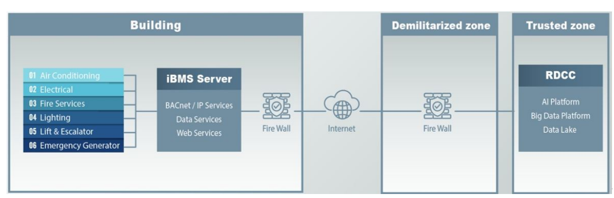

Figure 1: Data Collection Flow Chart

2.5 Semantic Artificial Intelligent

The adoption of artificial intelligence techniques in EMABS system and data analytics presents significant challenges due to the distinct knowledge and experience required in each domain. Optimizing control of complex EMABS systems necessitates aligning and standardizing knowledge across domains. However, a fundamental hurdle is the lack of well-structured and consistent identification

Table 3: Purpose and Details of the Data Processing

Process	Purpose and Details
Resampling	Resample data into hourly timeframe. (E.g., Cooling Load, Temperature, etc.)
Data Type Conversion	Convert datatype to float / integer for calculation and model prediction purpose.
Remove Negative Value	Remove Negative Value for engineering parameters (E.g., Energy Consumption).
Interpolation	Unify timestamp info for all setpoints, avoid missing timestamp.
Feature Combination	Combining two or more existing features to create a new one (E.g., Calculate Wet bulb by Temperature & Humidity).
Feature Aggregation	Aggregating features to create a new one (E.g., Chiller Sequence Column by counting On Off Status of chillers).
Feature Extraction	Extract different datetime information (E.g., Month, Year and Holiday from timestamp).

and classification of equipment. Building systems often have disorganized naming structures, leading to confusion and integration issues. Standardizing data naming and representations is crucial to address these concerns. Capturing entity properties and relationships throughout the asset lifecycle is also essential. Without unified semantic representation, understanding interrelated data requires extensive expertise. This poses difficulties for widespread AI deployment across buildings.

To unlock the potential of big data and AI in the E&M industry for mass deployment, the Semantic AI approach is recommended. Semantic model is an ontology-based framework following the Resource Description Framework (RDF) and Web Ontology Language (OWL) from World Wide Web Consortium (W3C) [17]. This ontology combines methodologies from Machine Learning, Knowledge Graph Modeling, Natural Language Processing, and Text Mining. Unlike traditional approaches, Semantic AI introduces a separate layer that enables the creation of AI models based on the semantic relationships among different equipment. By utilizing a unified semantic model that encapsulates various subsystems within a building, it becomes possible to programmatically explore the operational, structural, and functional aspects of the building. This approach can be extended for application in other buildings, facilitating the widespread adoption of AI-driven solutions. A semantic data platform, consisting of a graph database for storing the semantic model and a time-series database for building data, is utilized (Figure 2). The building data associated with the semantic model instance is ingested into a time series database. Entities in the instance can be reached through an application programming interface (API) for semantic path queries. Data queries can be formed by the semantic data platform to retrieve the time series data linked to the entity. This enables different domain experts to perform analytics and diagnostics separately, making the programs of the AI services portable across buildings in a shorter timeframe.

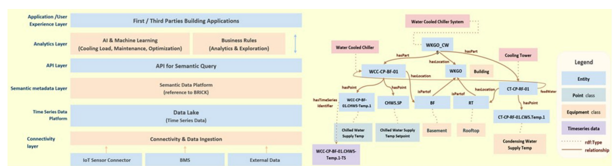

Figure 2: Architecture Overview of Semantic Data Platform and Partial View of a Water-Cooled Chiller System at A Government Building

Other building applications, like Building Information Model (BIM), can access the open and portable building data using our AI semantic model approach. This allows application developers to directly deploy their existing solutions, based on AI models developed for a similar building, in a new building. This approach significantly reduces time and cost for model development and data analytics. See the illustration below (Figure 3) for the Semantic Model integrated with the BIM model for daily Operation and Maintenance.

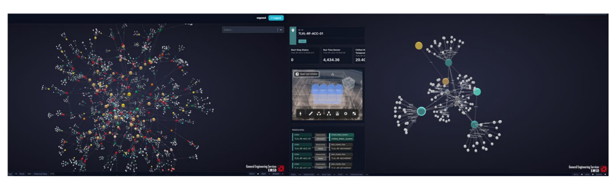

Figure 3: Comprehensive Semantic Model and Partial View of the Chiller System with Integration with BIM

The overall flow chart for creating the semantic file proposed is shown in the following graph (Figure4).

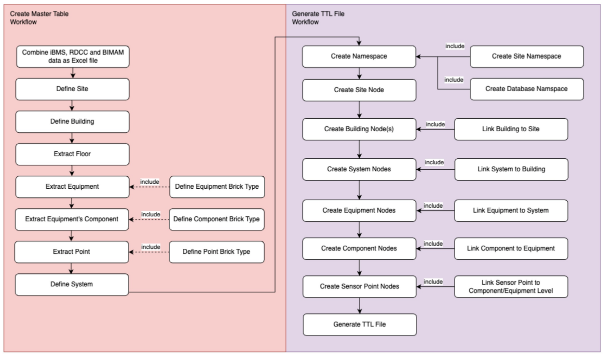

Figure 4: Flow Chart of Creating Master Table and the Semantic File

3 DATA EXPLORATION

Data exploration plays a pivotal role in comprehending and analyzing intricate systems. The primary objective is to visually represent and analyze current plant information in order to identify abnormalities prior to the application of AI algorithms for load prediction. By utilizing a diverse range of visualization tools and techniques, engineers can acquire valuable insights into the behavior and performance of the system, thereby facilitating proactive measures and enhancing overall system efficiency.

One of the key visualizations used in this framework is the heat map, which illustrates the cooling load of individual chillers across different hours of the day and months throughout the year (Figure 5). This visualization enables the identification of patterns and variations in cooling load, aiding in the detection of abnormal behavior or inefficiencies during specific time periods. Additionally, by comparing the heat maps of all the chillers within the same site, engineers can observe if there is any overloading or imbalanced load distribution among the chillers.

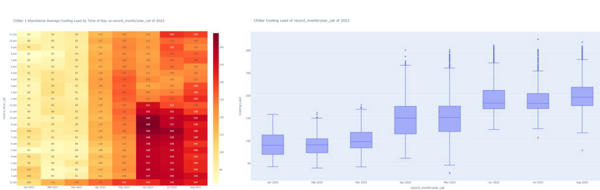

Figure 5: Individual Cooling Load by Data and Time and Total Chiller Plant Cooling Load

A box graph is employed to represent the monthly total system cooling load across the year. This visualization helps identify seasonal trends and variations in the cooling load, providing insights into the system's performance under different climatic conditions.

Analyzing the performance curves of chillers is crucial for evaluating efficiency across temperature ranges and part load ratios (Figure 6). Research has shown that over half of the energy is consumed at the low Part Load Ratio (PLR) range [13]. This analysis helps identify deviations and optimize energy consumption. By considering each chiller's performance characteristics and selecting the most efficient ones, engineers can enhance overall chiller system efficiency and achieve optimal energy utilization.

Figure 6: Chiller Performance Curve against Part Load Ratio under Different Wet Bulb Temperature Range

4 CHILLER LOAD PREDICTION AND SEQUENCE OPTIMIZATION

4.1 Model Selection

In chiller load prediction, different algorithms have been tested and found applicable. However, prediction accuracy can vary depending on factors like location, building type, and chiller system characteristics. It can be observed that there are several main directions for model selection. Some of them adopt Artificial Neural Network (ANN) [9–12, 14], while other techniques include Support Vector Machine (SVM) [3] and Extreme Gradient Boosting (XGBoost) [18]. When deploying at scale, with diverse building variations, it is recommended to test multiple algorithms and select the best model before fine-tuning. The traditional approach to chiller load prediction involves manual testing of each model or adopting a previous one. However, this approach is time-consuming and inefficient for mass deployment in diverse buildings. To address this challenge, a new approach has been adopted, leveraging online AI tools that enable a quick and efficient evaluation of multiple algorithms (Figure 7). These tools provide a platform for a brief trial of different algorithms, allowing for a comparative analysis of their performance. By utilizing these online AI tools, the process of algorithm selection and tuning can be streamlined, facilitating the mass deployment of chiller load prediction models.

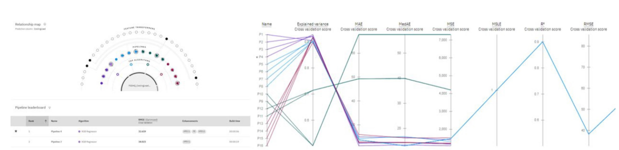

Figure 7: Relationship Map and Metric Chart of the Online AI Tool Running the Chiller Load and COP Prediction

The analysis of the results reveals that the top three performing algorithms for chiller load prediction are all variants of XGBoost with delete enhancement [18]. And the RMSE values are 32.619, 34.023, and 34.023, respectively. XGBoost is a popular Machine Learning (ML) algorithm with relative low programming complexity when comparing to Deep Learning algorithm. It provides faster training speed and more robust to noise and missing data and hence more suitable for mass deployment case. As a result, the XGBoost models are exported as plain code to undergo further fine-tuning and development for real-world application scenarios. This process allows for a more detailed optimization of the models to ensure their effectiveness and accuracy in practical use cases.

4.2 Model Training

The tool aids in selecting the most suitable algorithm for two key tasks: (i) cooling load prediction and (ii) Coefficient of Performance (COP) prediction. The next step involves training and fine-tuning hyperparameters to improve prediction accuracy. Different approaches like manual tuning, grid search, random search, and Bayesian optimization can be used. For the XGBoost algorithm with 12 hyperparameters, an automated framework performs an exhaustive search for optimal settings in each specific case. The Table 4 below provides a comprehensive overview and detailed explanations of hyperparameter settings.

The Tree-Structured Parzen Estimator method is used in hyperparameter tuning. It selects the parameter with the highest expected improvement, sorts them by scores, and divides them into two Parzen Estimators for further enhancements. Evaluating the performance of hyperparameter tuning involves training the loss function and striving for minimal of RMSE and MAE.

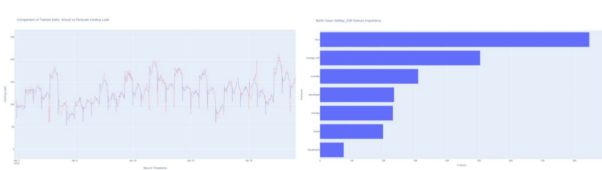

Figure 8: Predicted Cooling Load against Actual Cooling Load and the Feature Importance

Table 4: Hyperparameter Setting

Parameter	Range	Explanation
n_estimators	100, 2500, 100	Number of trees in the forest.
min_child_weight	1, 600, 1	Minimum sum of instance weight (hessian) needed in a child.
learning_rate	0.1, 1.0	Step size shrinkage used in update to prevents overfitting.
subsample	0.1, 1.0	Subsample ratio of the training instances.
eta	0.1, 1.0	Learning rate.
reg_alpha	0.1, 10.0	L1 regularization term on weights.
reg_lambda	0.1, 10.0	L2 regularization term on weights.
colsample_bytree	0.1, 1.0	Subsample ratio of columns when constructing each tree.
max_depth	1, 600	Maximum depth of a tree.
gamma	0.1, 10.0	Minimum loss reduction required to make a further partition on a leaf node of the tree.
num_parallel_tree	1, 50	Number of parallel trees constructed during each iteration.
nthread	1, 50	Number of parallel threads used to run XGBoost.
n_estimators	100, 2500, 100	Number of trees in the forest.

The chiller cooling load is accurately predicted using the trained and tuned XGBoost model (Figure 8). Additionally, feature importance analysis reveals that the most influential feature is the hour of the day, followed by average outdoor temperature, humidity, and the day of the week. These findings align with the domain understanding of chiller load variation, highlighting the significance of temporal factors and environmental conditions in determining the cooling load requirements. After prediction, the suggested chiller sequence is calculated, considering higher plant COP and the lowest total energy consumption. The statistical data are listed in the Table 5.

Based on the available information, the system will provide recommendations for the optimal high-level chiller sequence and control settings. The red area depicted in Figure 9 illustrates the maximum energy savings that can be attained by adopting these suggested settings. However, it is important to note that in real-life implementation, the actual energy savings may be slightly lower. This discrepancy arises from the presence of control constraints that are put in place to ensure the protection and longevity of the chiller system. These control constraints play a vital role in safeguarding the reliability and performance of the chiller system in the long run.

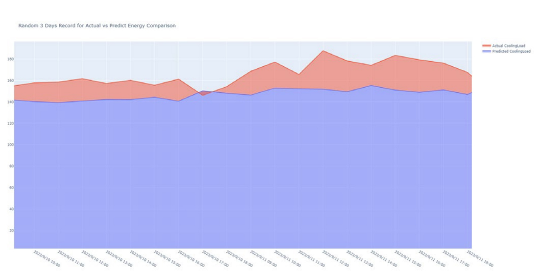

Figure 9: Energy Saving from Suggested Chiller Sequence and Settings

Table 5: Statistical Data of the Proposed Chiller Sequence

Chiller Sequence	Average COP	Standard Deviation	Min Average COP	25% Average COP	50% Average COP	75% Average COP	Max Average COP
No.1 and No.3	4.63	0.31	3.82	4.50	4.63	4.83	5.52
No.2 and No.3	4.92	0.45	3.82	4.60	4.90	5.30	5.93

4.3 Model Validation

The model's training phase involves splitting historical data into a 70% training set and a 30% evaluation set. After evaluation, a validation stage is conducted using three days of unseen actual data to test the model's performance in real-world scenarios. The Figure illustrates the hourly profiles of the actual cooling load and the predicted cooling load during this validation stage. The results, shown in the following Figure 10, demonstrate similar shapes between the actual and predicted cooling load profiles.

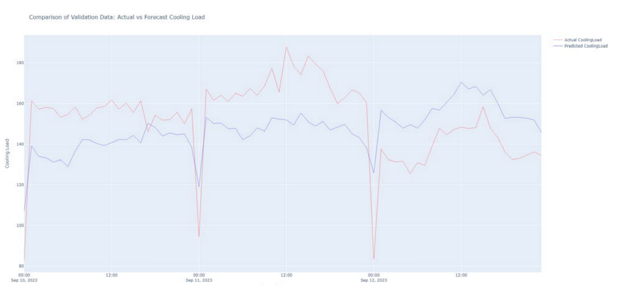

Figure 10: Validation of the Prediction Model on 3-Day Unseen Data

After analyzing the results, it is evident that the Root Mean Square Error (RMSE) and Mean Absolute Error (MAE) in the validation set are slightly higher compared to the training set (Table 6). However, they still remain within an acceptable range. This signifies that the model performs well and shows promise for real-world deployment.

Table 6: Comparison of the Accuracy of Evaluation and Validation Data

Stage	RMSE	MAE
Evaluation	11.079	6.87
Validation	19.84	18.59

5 IMPLEMENTATION AND RESULTS

From the study, the approach was first implemented in Site A and then expanded to Site B to test the expandability of the framework for mass deployment of AI-based Chiller Energy Optimization. The basic information of the sites is summarized in Table 7, with Site A being a large Office Complex and Site B being a small Clinic and Laboratory. The results have shown that the redeployment time in Site B has significantly reduced with the standardized chiller plant data format and infrastructure, consolidated database of local weather data, cloud data cleansing, and analytics platform.

Table 7: Characteristic and Result of Two Tested Sites

	Site A	**Site B**
Usage	Office Complex	Clinic and Laboratory
Construction Floor Area (sqm)	$\sim 90,000$	$\sim 1,100$
Type of chiller	Water-cooled	Air-cooled
Primary system	Variable flow	Variable flow
No. of Chillers	5	3
Total Cooling Load (kW)	$\sim 17,500$	~ 450
Possible Load Saving (%)	~ 13	$\sim$ 5.5 for 3-day validation period
Deployment time (Months)	11	5

6 CONCLUSION

In summary, this paper presents a comprehensive framework for the mass deployment of AI algorithms in chiller energy optimization. The framework encompasses standardized data formats, semantic AI, data exploration, data cleansing, and efficient model selection and tuning using market-available tools. By establishing standardized data formats and collection infrastructure, consistency and compatibility in data acquisition are ensured. The incorporation of semantic AI enhances data understanding and utilization by leveraging domain knowledge and equipment relationships. Thorough data exploration and cleansing techniques provide valuable insights and maintain data accuracy. The streamlined approach to model selection and tuning reduces redeployment time, enabling faster implementation and energy-saving benefits.

However, there are limitations and future considerations to address. For existing buildings, the absence of sensors and data collection infrastructure may impede deployment and necessitate additional time. Moreover, the cost implications of installing new sensors and infrastructure must be taken into account. Hence, it is advisable to prioritize deploying this framework in buildings already equipped with sensors and infrastructure. Furthermore, while this study encompassed two building types, specific usage buildings like sports stadiums or museums may present unique challenges warranting further investigation. Future work should concentrate on addressing these specific building types and exploring potential solutions for unforeseen issues.

Lastly, the use of a semantic model has shown promise in aiding model training through domain knowledge and allowing better transferability to different building types. Further study should focus on the development of graphical-embedded techniques to fully harness the additional information provided by the semantic model. Meanwhile, another direction is to generalize and summarize the types of AI models that should be adopted for different types of buildings and chiller system configurations, based on obtaining more cases.

In conclusion, this paper's framework establishes the groundwork for widespread adoption of AI-driven Chiller Energy Optimization in buildings. By addressing limitations and conducting additional research, we can continue to improve energy efficiency and sustainability, contributing to a greener future.

REFERENCES

[1] D. Azuatalam et al. 2020. Reinforcement learning for whole-building HVAC control and demand response. *EnergyAI* 2 (2020), 100020. https://doi.org/10.1016/j.egyai.2020.100020

[2] Y. Du et al. 2021. Intelligent multi-zone residential HVAC control strategy based on Deep Reinforcement Learning. *ApplEnergy* 281 (2021), 116117. https://doi.org/10.1016/j.apenergy.2020.116117

[3] C. Fan et al. 2020. Improving cooling load prediction reliability for HVAC system using Monte-Carlo simulation to deal with uncertainties in input variables. *Energy and Buildings* 226 (2020), 110372. https://doi.org/10.1016/j.enbuild.2020.110372

[4] C. Fan, F. Xiao, and Y. Zhao. 2017. A short-term building cooling load prediction method using deep learning algorithms. *Applied Energy* 195 (2017), 222–233. https://doi.org/10.1016/j.apenergy.2017.03.064

[5] A. Gupta et al. 2021. Energy-efficient heating control for smart buildings with deep reinforcement learning. *JBuildEng* 34 (2021), 101739. https://doi.org/10.1016/j.jobe.2020.101739

[6] Hong Kong Electrical and Mechanical Services Department. 2023. Hong Kong Energy End-use Data 2022.

[7] Hong Kong Environment and Ecology Bureau. 2021. Hong Kong's Climate Action Plan 2050.

[8] Hong Kong Observatory. 2023. Hong Kong Observatory Open Data. https://www.hko.gov.hk/en/abouthko/opendata_intro.htm Retrieved September 27, 2023.

[9] J.-H. Kim, N.-C. Seong, and W. Choi. 2019. Modeling and optimizing a chiller system using a machine learning algorithm. *Energies* 1215 (2019), 2860. https://doi.org/10.3390/en12152860

[10] J.-H. Kim, N.-C. Seong, and W. Choi. 2020. Forecasting the energy consumption of an actual air handling unit and absorption chiller using Ann Models. *Energies* 1317 (2020), 4361. https://doi.org/10.3390/en13174361

[11] N. Moghaddas-Zadeh et al. 2023. Ann-based procedure to obtain the optimal design and operation of the compression Chiller Network – energy, economic and environmental analysis. *JBuildEng* 72 (2023), 106711. https://doi.org/10.1016/j.jobe.2023.106711

[12] Nasruddin et al. 2019. Optimization of HVAC system energy consumption in a building using artificial neural network and multi-objective genetic algorithm. *Sustainable Energy Technologies and Assessments* 35 (2019), 48–57. https://doi.org/10.1016/j.seta.2019.06.002

[13] B.M. Seo and K.H. Lee. 2016. Detailed analysis on part load ratio characteristics and cooling energy saving of Chiller staging in an Office Building. *EnergyBuild* 119 (2016), 309–322. https://doi.org/10.1016/j.enbuild.2016.03.067

[14] L. Wang et al. 2019. Cooling load forecasting-based predictive optimisation for Chiller Plants. *Energy and Buildings* 198 (2019), 261–274. https://doi.org/10.1016/j.enbuild.2019.06.016

[15] T. Wei, S. Ren, and Q. Zhu. 2021. Deep reinforcement learning for joint datacenter and HVAC Load Control in distributed mixed-use buildings. *IEEETransSustainComp* 63 (2021), 370–384. https://doi.org/10.1109/tsusc.2019.2910533

[16] T. Wei, Y. Wang, and Q. Zhu. 2017. Deep reinforcement learning for building HVAC control. In *ProcDesignAutoConf.* https://doi.org/10.1145/3061639.3062224

[17] World Wide Web Consortium. 2014. Resource Description Framework (RDF). https://www.w3.org/RDF/ Retrieved November 1, 2023.

[18] XGBoost 2023. XGBoost Documentation - xgboost 2.0.3 documentation. Available at https://xgboost.readthedocs.io/en/stable/. Retrieved 1 November 2023.

[19] H. Zhao et al. 2021. Hybrid-model-based deep reinforcement learning for heating, ventilation, and air-conditioning control. *FrontEnergyRes* 8 (2021). https://doi.org/10.3389/fenrg.2020.610518

[20] S.L. Zhou et al. 2023. A comprehensive review of the applications of machine learning for HVAC. *Decarbon* 2 (2023), 100023. https://doi.org/10.1016/j.decarb.2023.100023