

Chiller Plant Replacement by MiMEP and Optimization with Reinforcement Learning Algorithm

Chelsea H.C. LI, Chris T.K. WONG, Tim C.Y. LAI and Sammy S.K. YEUNG
Electrical and Mechanical Services Department,
The Government of the Hong Kong Special Administrative Region,
Hong Kong, China

W.K. YEUNG, Paul Y.H. TSOI and Andy C.Y. TSANG
Innovative Application and Development Department, Capax Technology Limited,
Hong Kong, China

Abstract

In pursuit of carbon neutrality, the Electrical and Mechanical Services Department (EMSD) of the Government of the Hong Kong Special Administrative Region implemented Multi-trade Integrated Mechanical, Electrical and Plumbing (MiMEP) for the first time during the Tai Lung Veterinary Laboratory's chiller plant replacement. To further optimize the chiller plant efficiency and reduce energy consumption, chiller optimization with partially observable reinforcement learning (RL) algorithm was proposed, where three machine learning models have been developed to forecast cooling demand, predict cooling load, and predict energy consumption. By leveraging zero-inflated regression technique, these models establish an environment configuration for the RL algorithm. When compared with the default setting, the optimization approach can enhance the overall chiller plant efficiency by approximately 20% based on simulation. These findings highlight the potential of combining MiMEP with artificial intelligence for sustainable energy management, emphasizing the importance of technological integration in achieving carbon reduction objectives.

Keywords: Carbon neutrality, Energy saving, MiMEP, Artificial intelligence, Reinforcement learning, Chiller plant optimization

1 INTRODUCTION

Central air-conditioning system, particularly the chiller plant, is recognized as the largest energy consumer in modern buildings. In Hong Kong, air conditioning has been reported to account for 30% of total electricity consumption [1]. This energy demand is even more pronounced in office buildings, where air conditioning can comprise over 45% of the total energy usage [1], and this percentage is projected to continue rising. The escalating energy consumption can have a significant impact on the environment, resulting in increased pollution levels [2], and exacerbation of climate change [3]. Therefore, an effective energy management for chiller plants is vital for minimizing energy usage, reducing operational costs and mitigating environmental footprint.

In line with the goal of carbon neutrality, the Electrical and Mechanical Services Department (EMSD) of Hong Kong has successfully completed a pilot project that adopted Multi-trade Integrated Mechanical, Electrical and Plumbing (MiMEP) [4] in the replacement of chiller plant at Tai Lung Veterinary Laboratory (Figure 1).




Figure 1. MiMEP modules of the chiller plant at Tai Lung Veterinary Laboratory

The uninterrupted operation of the chiller plant is essential for the laboratory's critical role in conducting diverse veterinary experiments. The objectives of the replacement were to enhance the chiller plant efficiency and minimize carbon emission by implementing MiMEP technology during the replacement process. The adoption of MiMEP has demonstrated significant successes in accomplishing carbon neutrality targets, and the potential for wider use of MiMEP in Repair, Maintenance,

Alteration, and Addition (RMAA) projects. These enhancements consist of a 70% reduction in material wastage, as well as reductions of over 50% in construction time and manpower, achieved by off-site prefabrication (Figure 2). These changes contribute to improved environmental sustainability and are in line with the objective of achieving carbon neutrality.



Figure 2. Off-site prefabrication of chilled water pump set

In order to further enhance the efficiency of the chiller plant, the use of Artificial Intelligence (AI) models was proposed to optimize its control logic. The building management system (BMS) plays a crucial role in monitoring and regulating the control of the heating, ventilation and air conditioning (HVAC) systems. While the BMS has made notable contributions in coordinating the chiller plant operation, the existing control strategy of the system relies on rule-based approaches which could lead to ineffective operation and potential energy wastage. Although BMS control settings can be adjusted manually during the daily operation, this approach is deemed impractical and ineffectual when it applies to optimizing the chiller plant in the long term. Extensive manpower effort is necessary to modify the control parameters for optimal performance under various loads and weather conditions. Therefore, the development of a dynamic optimization strategy for achieving automatic optimization of chiller plant operations under various circumstances is essential.

In pursuit of dynamic optimization, the potential implementation of AI technology was being explored. In recent years, various AI applications have been developed to optimize the operation of chillers. Instead of depending on mathematical models derived from physical laws, the optimization approach is built with data-driven models that are generated through

historical operational data obtained from the chiller plant. There are numerous data-driven optimization approaches available for chiller optimization, including but not limited to particle swarm optimization (PSO) [5], genetic algorithms [6], machine learning algorithms [7], neural networks [6], and reinforcement learning algorithms [8]. However, these existing optimization approaches often encounter issues like low convergence rate, incapability to handle high-order optimization, or optimization of expensive black-box functions. In contrast, the combination of reinforcement learning with Bayesian optimization for chiller plant optimization remains largely unexplored. In this paper, a partially observable reinforcement learning (RL) algorithm integrated with Bayesian optimization [9] was proposed to achieve dynamic optimization of the chiller plant. This approach can effectively tackle high-order optimization challenges and manage costly black-box functions. The algorithm incorporates three machine learning algorithms to construct an environment configuration for cooling load prediction, cooling demand forecast, and energy consumption prediction. Considering the limited dataset size and the prevalence of missing data and zero counts in our dataset, the zero-inflated regression technique, combining a categorical boosting (CatBoost) regressor and an extremely randomized trees (Extra-Trees) classifier was employed to train the three models. The resultant cooling load prediction model and energy consumption prediction model establish a simulation environment for RL agents to explore, enabling them to learn and adapt dynamic control policies that aim at minimizing the chiller plant's energy consumption while fulfilling the cooling load requirements.

The optimization algorithm's performance was evaluated using data from 2023, where statistical analysis was employed to evaluate the improvement of the chiller plant in terms of the coefficient of performance (COP) and energy consumption.

2 PROBLEM STATEMENT

The chiller system to be optimized comprises three chillers and four chilled water pumps. The calculation of the cooling load supplied by the chiller plant was based on the following formula, taking into account the water volume flow rate, supply water temperature, and return water temperature.

$$Q = \sum_{i=1}^3 V_i * p * C_p * \Delta T_i \quad (1)$$

where

- Q : Cooling load (kW)
- V : Volume flow rate (m³/s)
- p : Water density (1000 kg/m³)
- C_p : Water specific heat capacity (4.19 kJ/(kg*K))
- ΔT : Temperature difference between chilled water supply and return

The energy consumption of the chiller plant was determined by summation of the active power of each chiller and chilled water pump, as indicated in the following formula.

$$E = \sum_{i=1}^3 E_{chiller_i} + \sum_{j=1}^4 E_{pump_j} \quad (2)$$

where

- E : Total energy consumption of chiller plant (kW)
- $E_{chiller}$: Energy consumption of chiller (kW)
- E_{pump} : Energy consumption of chilled water pump (kW)

The objective of the optimization process is to minimize the overall energy consumption of the chiller plant, while ensuring the cooling load requirement of the Laboratory is met, as represented in the following objective function:

$$\min |C_{forecast} - C_{predict}|, E \quad (3)$$

subject to

$$C_{forecast} * (1 - C_t) < C_{predict} < C_{forecast} * (1 + C_t) \quad (4)$$

$$E > 0 \quad (5)$$

where

- C_t : Tolerance for cooling load
- $C_{forecast}$: Forecast cooling load
- $C_{predict}$: Predicted cooling load
- E : Energy Consumption

3 METHODOLOGY

3.1 Dataset

The chiller plant operating data used for the model development were collected from the BMS between August 2022 and February 2024, at 15-minute intervals. The data was subsequently divided into training, validation, and testing sets in a ratio of 7:2:1 for the purpose of model development. Throughout the data collection process, occasional issues such as data corruption, incompleteness, missing values, and out-of-range readings, were encountered. Using the irrelevant or incorrect data when developing AI models may compromise the model's effectiveness and result in false optimization outcomes. Therefore, multiple data pre-processing procedures were implemented to tackle these problems and filter out unreliable data. For instance, zero or negative values of the chilled water flow rate and COP were excluded due to the routine maintenance. In addition, data pre-processing involved the assignment of zero values to data points that corresponded to periods when the chiller or chilled water pump was not running, the rectification of data through gain adjustments, the removal of outliers, the calculation of cooling load and energy consumption, and the conversion of data types. These steps collectively contributed to the development of models for analysis.

3.2 Chiller Plant Optimization Algorithm

The proposed chiller plant optimization approach employed a partially observable RL algorithm, coupled with multi-objective Bayesian optimization, to dynamically optimize water-side chiller plant operation. This methodology leveraged the capabilities of adapting to changing preferences or priorities in multi-objective optimization problems, allowing it to find trade-off solutions between different objectives. Furthermore, it benefited from the efficient search space exploration and exploitation offered by Bayesian optimization. Through iteratively learning and optimization, the RL agent continuously refined its control strategies by interacting with a simulated environment. This enabled effective response to diverse cooling demands and environmental factors, ultimately resulting in significant energy savings, cost reduction, and improved sustainability.

Three machine learning regression models were developed: a cooling load prediction model, an energy consumption prediction model, and a cooling demand prediction model. These models collectively established an environment configuration for the RL process. During the RL process, the cooling load prediction model and energy consumption prediction model were utilized to create the simulation environment for RL agent to interact with. Figure 3 shows an overview of the design.

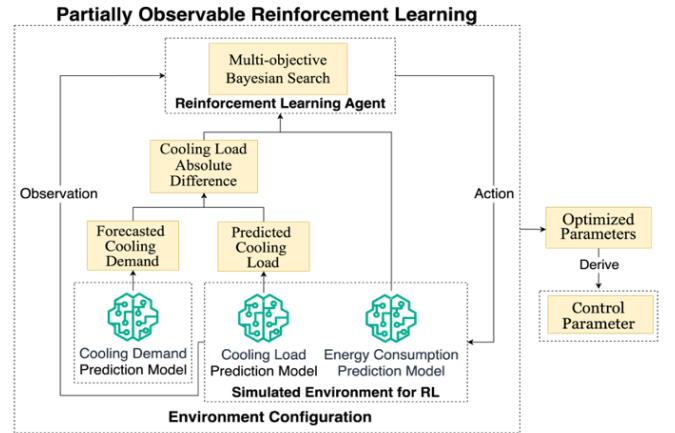


Figure 3. Overview design of the partially observable RL algorithm for chiller plant optimization

The RL process consisted of several steps. First, the cooling demand prediction model forecasted the required cooling load for the chiller plant. Then, the RL agent was employed to initialize and adjust parameters through predefined iterations, treating them as “actions” that were inputted into machine learning models for cooling load prediction and energy consumption prediction. The predicted cooling load and energy consumption served as “observations”, while the “rewards”

were determined by calculating the negative absolute difference between the forecasted and predicted cooling load, as well as the negative predicted energy consumption. These rewards guided the agent in re-initializing the parameters.

The re-initialization process utilized Bayesian optimization to iteratively update the parameters based on the rewards obtained. Following each iteration, the agent aimed at maximizing the reward by adjusting the chiller plant's operating parameters based on the observed rewards. Penalties were imposed in cases where the predicted energy consumption dropped below zero or if the predicted cooling load deviated from the range defined by the forecasted cooling load and the tolerance value specified in the objective function. Through successive iterations, the parameters that achieved the lowest energy consumption while meeting the cooling load target were identified. These optimal parameters were then derived into control settings to optimize the operation of the chiller plant.

3.3 Machine Learning Regression Models

To facilitate the optimization process and address the challenges posed by limited dataset size, as well as the prevalence of missing data and zero counts when chiller is not in operation, the zero-inflated regression technique, as illustrated in Figure 4, was employed to construct the three machine learning regression models. This technique combined a classifier and a regressor to differentiate between periods when the chillers and chilled water pumps were active or inactive. The classifiers were employed to identify the on/off status of each chiller and chilled water pump, providing valuable information for the regression model to accurately predict the cooling load, as well as the energy consumption.

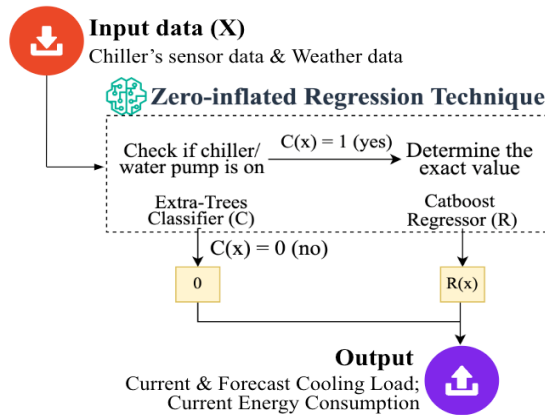


Figure 4. Schematic diagram illustrating the mechanism of zero-inflated regression model

The classifier was formulated using Extra-Trees [10], a variant of the random forest algorithm known for its ability to reduce data bias and training time. For the regressor, CatBoost was

adopted, followed by the Extra-Trees classifiers. The CatBoost regressor was well suited for handling categorical features, preventing model overfitting, and achieving superior performance through a combination of random permutations and ordered boosting [11]. They were designed to predict precise cooling load and energy consumption values. The combination of Extra-Trees classifier and CatBoost regressor enabled the capturing of intrinsic relationship between the operational states of the chillers and chilled water pumps, and their corresponding cooling load and energy consumption.

Table 1. Input attributes used for cooling load prediction and energy consumption prediction models

Attributes	Type	Description
<i>MCHWSWT</i>	Float	Main Chilled Water Supply Water Temperature
<i>MCHWSFWR</i>	Float	Main Chilled Water Supply Flow Rate
<i>CHR-PLANT-SWTSP</i>	Float	Chiller - Plant Room - Supply Water Temperature Setpoint
<i>VSD-DPRES</i>	Float	Variable Speed Drive - Differential Pressure Setpoint
<i>CHR_i-S</i>	Bool	On/Off Status of Chiller <i>i</i>
<i>CHR_i-PAMP1</i>	Float	Current for Each Phase 1 of Chiller <i>i</i>
<i>CHR_i-PAMP2</i>	Float	Current for Each Phase 2 of Chiller <i>i</i>
<i>CHR_i-PAMP3</i>	Float	Current for Each Phase 3 of Chiller <i>i</i>
<i>CHR_i-DAT</i>	Float	Discharge Air Temperature (Condenser) of Chiller <i>i</i>
<i>CHR_i-SUAT</i>	Float	Suction Air Temperature (Condenser) of Chiller <i>i</i>
<i>CHR_i - CHWSWT</i>	Float	Chilled Water Supply Water Temperature of Chiller <i>i</i>
<i>CHR_i - CHWFWR</i>	Float	Chilled Water Flow Rate of Chiller <i>i</i>
<i>CHR_i - RUNTIME</i>	Float	Running Time of Chiller <i>i</i>
<i>CHWP_i-S</i>	Bool	On/Off status for Each of Chilled Water Pump <i>i</i>
<i>CHWP_i-PAMP1</i>	Float	Current for Each Phase 1 of Chilled Water Pump <i>i</i>
<i>CHWP_i-PAMP2</i>	Float	Current for Each Phase 2 of Chilled Water Pump <i>i</i>
<i>CHWP_i-PAMP3</i>	Float	Current for Each Phase 3 of Chilled Water Pump <i>i</i>
<i>CHWP_i - VSDSPD</i>	Float	VSD Speed of Chilled Water Pump <i>i</i>
<i>CHWP_i - VSDSPDC</i>	Float	VSD Speed Setpoint of Chilled Water Pump <i>i</i>
<i>CHWP_i - WDPRE</i>	Float	Water Discharge Pressure of Chilled Water Pump <i>i</i>
<i>CHWP_i - WSUPRE</i>	Float	Water Suction Pressure of chilled water pump <i>i</i>
<i>PLANT-RMRH</i>	Float	Plant Room - Room Relative Humidity
<i>PLANT-RMT</i>	Float	Plant Room - Room Temperature
<i>T_{outdoor}</i>	Float	Outdoor Air Temperature
<i>H_{outdoor}</i>	Float	Outdoor Air Humidity

A Pearson correlation analysis was conducted to investigate the relationship between the sensor data from the chiller plant and the cooling load as well as energy consumption. The analysis pinpointed essential variables with a correlation coefficient exceeding 0.1, which were chosen for the input attributes of the

cooling load prediction model and the energy consumption prediction model. These selected attributes are listed in Table 1.

Rather than predicting the current cooling load and energy consumption, the cooling demand prediction model was specifically designed to forecast the cooling load of the chiller plant at a 30-minute ahead interval. To achieve accurate prediction of the cooling demand, additional input attributes were incorporated into the model. These attributes, as outlined in Table 2, encompassed features associated with the chiller's return water, as well as forecasted outdoor temperature and humidity. They are used to ensure the model would consider pertinent factors for precise forecast of cooling demand.

The performance of these three machine learning models was assessed using Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) as performance metrics. The use of zero-inflated regression technique in combination with the CatBoost regressor and the Extra-Trees classifier yielded the lowest MAE and RMSE score compared to other tree-based algorithms (Table 3). A lower MAE and RMSE score indicate higher model accuracy, demonstrating the algorithm's capability in simulating the chiller plant environment.

Table 2. Additional input attributes used for cooling demand prediction model

Attributes	Type	Description
<i>MCHWRWT</i>	Float	Main Chilled Water Return Water Temperature
<i>MCHWRFWR</i>	Float	Main Chilled Water Return Flow Rate
<i>CHR_i - CHWRWT</i>	Float	Chilled Water Return Water Temperature of Chiller <i>i</i>
<i>T_{outdoor-forecast}</i>	Float	Forecast Outdoor Air Temperature
<i>H_{outdoor-forecast}</i>	Float	Forecast Outdoor Air Humidity

Table 3. Performance comparison of cooling load prediction model | energy consumption prediction model | cooling demand prediction model

Model	MAE	RMSE
Random Forest	7.40 2.19 11.46	13.10 4.94 16.95
XGBoost	7.32 2.28 9.93	12.56 3.45 14.99
CatBoost	5.90 1.07 8.69	8.32 1.56 12.63
Zero-inflated technique (CatBoost regressor + Extra-trees classifier)	5.74 1.04 8.56	8.30 1.51 12.53

4 RESULTS AND DISCUSSION

The algorithm was tested using data collected in 2023. Figures 5 and 6 show the chiller plant's COP before and after optimization respectively in relation to the cooling load under different outdoor temperature conditions. Figure 5 represents the logged data obtained under the default chiller plant control setting, whereas Figure 6 illustrates the simulated COP of the chiller

plant following optimization using the partially observable RL algorithm.

When comparing the results with the default control setting, the COP of the chiller plant after optimization are consistently higher, with an average theoretical improvement of approximately 20%. It is noteworthy to highlight that as the outdoor air temperature increases, the enhancement in COP becomes more prominent. Conversely, the degree of improvement diminishes when the outdoor temperature is low.

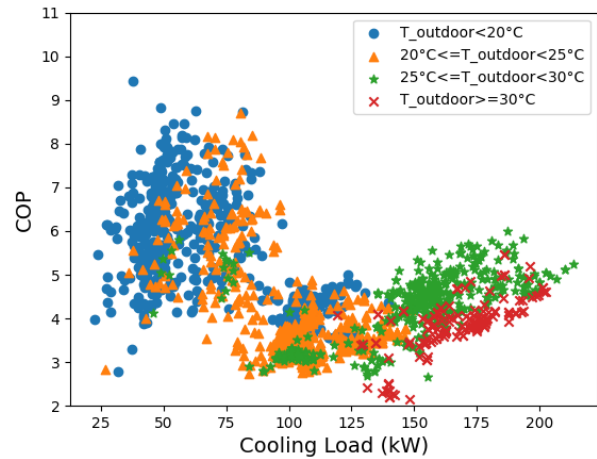


Figure 5. Chiller plant COP vs cooling load before optimization




Figure 6. Chiller plant COP vs cooling load after optimization

A Wilcoxon signed-rank test was performed to evaluate the statistical significance of the average improvement, yielding a p-value smaller than 0.0001 as depicted in Figure 8a, showing a high statistical significance in the observed improvement. This suggests the potential of the partially observable RL algorithm in optimizing the performance of the chiller plant,

leading to a more efficient operation.

Similar findings can be observed in terms of energy consumption during the chiller plant operation. The energy consumption of the chiller plant under the optimization decreases significantly by approximately 16% theoretically, with a p-value smaller than 0.0001 (Figure 8b) compared to the chiller plant under the default control setting (Figure 7). Despite the significant decrease in energy consumption, the cooling load with less than a 7.5% difference compared to the default chiller plant control setting. This optimization result highlights the 20% increase in COP, which led to a 16% reduction in energy consumption while maintaining the cooling capacity. This reduction in energy usage is essential for the sake of sustainability and cost-effectiveness as it leads to reduced electricity costs and a diminished environmental impact.

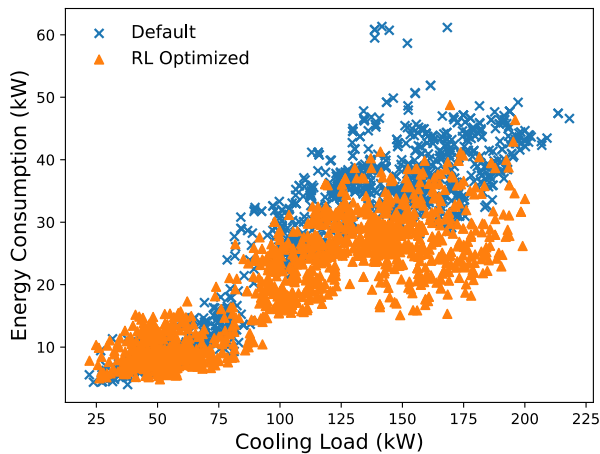


Figure 7. The comparison of energy consumption vs cooling load before and after optimization

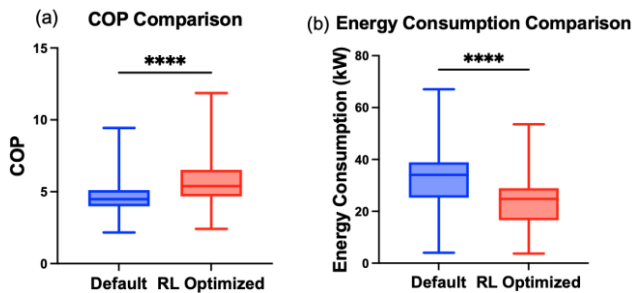


Figure 8. The comparison of magnitude of (a) COP and (b) energy consumption, before and after optimization (**** $p < 0.0001$, Wilcoxon signed-rank test)

5 CONCLUSIONS

The aim of this study is to replace chiller plant through MiMEP and optimize their operation through reinforcement learning algorithm. Chiller plant optimization using partially observable

reinforcement learning algorithm along with multi-objective Bayesian optimization was proposed due to their effectiveness in handling expensive black-box functions and capability in catering high-order optimization issues. Three machine learning models were developed to configure the environment, utilizing zero-inflated regression techniques to address issues arising from insufficient data, the prevalence of missing data and zero counts in the dataset. By interacting with the simulated environment, the reinforcement learning agent can constantly improve its control strategy to the chiller plant through iterative learning and optimization. The proposed algorithm demonstrated a theoretical reduction in the energy consumption by 16% and an improvement in coefficient of performance by 20%. It showcases the potential of the algorithm in improving environmental sustainability and contributes towards the goal of carbon neutrality. Further investigation will be conducted to evaluate the whole year performance of the chiller plant when the proposed algorithm is being deployed in the real-world setting.

REFERENCES

- [1] The Electrical and Mechanical Services Department of HKSAR Government, Hong Kong Energy End-use Data 2023, 2023.
- [2] I. Dincer, "Environmental impacts of energy," *Energy Policy*, vol. 27, no. 14, pp. 845–854, 1999.
- [3] U. F. AKPAN and G. E. Akpan, "The Contribution of Energy Consumption to Climate Change: A Feasible Policy Direction," *IJEPP*, vol. 2, no. 1, pp. 21–33, 2012.
- [4] The Architectural Services Department and the Electrical and Mechanical Services Department of HKSAR Government, "A game-changing shift in building services - MiMEP in government buildings," *Hong Kong Engineer*, vol. 51, pp. 8-15, 2023.
- [5] J. Kennedy and R. Eberhart, "Particle swarm optimization," *Proceedings of ICNN'95 - International Conference on Neural Networks*, pp. 1942–1948, 1995.
- [6] T. T. Chow, G. Q. Zhang, Z. Lin, and C. L. Song, "Global optimization of absorption chiller system by genetic algorithm and Neural Network," *Energy and Buildings*, vol. 34, no. 1, pp. 103–109, Jan. 2002.
- [7] W. T. Ho and F. W. Yu, "Chiller system optimization using K nearest neighbour regression," *Journal of Cleaner Production*, vol. 303, p. 127050, Jun. 2021.
- [8] S. Qiu, Z. Li, Z. Li, and X. Zhang, "Model-free optimal chiller loading method based on Q-Learning," *Science and Technology for the Built Environment*, vol. 26, no. 8, pp. 1100–1116, May 2020.
- [9] M. Balandat, B. Karrer, D. R. Jiang, S. Daulton, B. Letham, A. G. Wilson, and E. Bakshy, "Botorch: A framework for efficient monte-carlo bayesian optimization," *Advances in Neural Information Processing Systems*, no. 1807, pp. 21524–21538, 2020.
- [10] P. Geurts, D. Ernst, and L. Wehenkel, "Extremely randomized trees," *Machine Learning*, vol. 63, no. 1, pp. 3–42, 2006.
- [11] L. Prokhorenkova, G. Gusev, A. Vorobev et al., "CatBoost: unbiased boosting with categorical features," *Advances in neural information processing systems*, pp. 6639–6649, 2018.

Contact E-mail Address: hcli@emsd.gov.hk