

Heating, Ventilation, and Air Conditioning System

Subjects: Automation & Control Systems

Contributor: Zhiang Zhang, Xuyang Zhong

The heating, ventilation, and air conditioning (HVAC) system is a major energy consumer in office buildings, and its operation is critical for indoor thermal comfort.

Keywords: HVAC control ; deep reinforcement learning

1. Model Predictive Control

Model predictive control (MPC) has become popular over the past few years due to its potential for significant HVAC energy savings. MPC uses a building model to predict the future building performance, in which case the optimal control decisions for the current time step can be made. There have been a number of studies using MPC to control HVAC systems, such as controlling the supply air temperature setpoint for air handling units (AHUs) ^[1], controlling the on/off status of the HVAC system ^[2], controlling the ventilation airflow rates ^[3], and controlling zone air temperature setpoints ^[4], and most of them show significant energy savings.

While promising, MPC is still hard to implement in the real world because of the difficulties of HVAC modeling. The classic MPC requires low-dimensional and differentiable models; for example, the linear quadratic regulator needs a linear dynamics and quadratic cost function ^[5]. This is difficult for HVAC systems, especially for the supervisory control of centralized HVAC systems, not only because it has nonlinear dynamics but also because it involves a number of control logics that make it non-continuous. For example, the control logic for a single-speed direct-expansion (DX) coil may be “turn on the DX-coil if there is indoor air temperature setpoint not-met in more than five rooms”. Such logic is hard to represent with a continuous mathematical model because of the if-then-else condition. Therefore, in most previous MPC studies, either the building had an independent air conditioner for each room rather than a centralized system (such as ^[6] ^[7] ^[8]), or the MPC was used to directly control the local actuators rather than to set supervisory-level commands (such as ^[9] ^[10] ^[11]). Neither way generalizes well for typical multizone office buildings, which usually have centralized HVAC systems and non-uniform HVAC design.

To address the modeling difficulties of MPC for HVAC systems, white-box building model (physical-based model) based predictive control was proposed in ^[1] ^[6] ^[12]. This method may significantly reduce the modeling difficulties of MPC, because the white-box building model generalizes well for different buildings, and there are a number of software tools available for modeling. However, white-box building models, such as EnergyPlus models, are usually high-dimensional and non-differentiable. Heuristic search must be implemented for MPC. Given the fact that the white-box building model can be slow in computation, the scalability and feasibility of this type of MPC in the real world are questionable.

2. Model-Free Reinforcement Learning HVAC Control

Since model-based optimal control, such as MPC, is hard to use for HVAC systems, model-free reinforcement learning control becomes a possible alternative. To the authors' knowledge, reinforcement learning control for HVAC systems has not yet been well studied. Either the reinforcement learning methods used are too simple to reveal their full potential, or the test buildings are too unrealistic. For example, Liu and Henze ^[13] applied very simple discrete tabular-setting Q-learning to a small multizone test building facility to control its global thermostat setpoint and thermal energy storage discharge rate for cost savings. Regardless of the limited real-life experiment showing 8.3% cost savings compared with rule-based control, the authors admitted that the “curse of dimensionality” of such a simple reinforcement learning method limited its scalability. In the following research by the same authors ^[14], a more advanced artificial neural network (ANN) was used to replace simple tabular-setting Q-learning; however, the results indicate that the use of ANN did not show clear advantages, probably due to the limited computation resources at that time.

The deep neural network (DNN) has become enormously popular lately in the machine learning community due to its strong representation capacity, automatic feature extraction, and automatic regularization ^[15] ^[16] ^[17]. Deep reinforcement

learning methods take advantage of DNN to facilitate end-to-end control, which aims to use raw sensory data without complex feature engineering to generate optimal control signals that can be directly used to control a system. For example, Mnih et al. [18] proposed a deep Q-network that could directly take raw pixels from Atari game frames as inputs and play the game at a human level.

Deep reinforcement learning methods have been widely studied not only by machine learning and robotics communities but also by the HVAC control community. **Table 1** summarizes the HVAC control studies performed in recent years using deep reinforcement learning. Researchers have demonstrated via simulations and practical experiments that deep reinforcement learning can improve the energy efficiency for various types of HVAC systems. However, there are sparse data describing the implementation of end-to-end control for multizone buildings. On the one hand, the test buildings in several studies, including [19][20][21][22][23], were single zones with independent air conditioners. On the other hand, conventional deep reinforcement learning methods cannot effectively solve multizone control problems. Yuan et al. [24] showed that the direct application of deep Q-learning to a multizone control problem would make the training period too long. Ding et al. [25] proposed a multi-branching reinforcement learning method to solve this problem, but the method required a fairly complicated deep neural network architecture and therefore could not be scaled up for large multizone buildings. Based on deep reinforcement learning, Zhang et al. [26] proposed a control framework for a multizone office building with radiant heating systems. However, “reward engineering” (i.e., a complicated reward function of reinforcement learning) needed to be designed to help ensure that the reinforcement learning agent could learn efficiently, in which case end-to-end control could not be achieved.

Table 1. An overview of studies focusing on deep reinforcement learning methods for HVAC systems.

Reference	Reinforcement Learning Method	Test Building	HVAC System
[19]	Linear state-value function approximation	Two buildings, each having a single zone	Independent air conditioners
[20]	Continuous-action Q-learning	Multiple apartments, each having a single zone	Independent air conditioners
[21]	Continuous-action Q-learning	A hall with multiple single rooms	Independent air conditioners
[22]	Model-assisted fitted Q-learning	A lab room	An independent HVAC unit
[23]	Model-assisted fitted Q-learning	A single chamber	A heat pump system
[24]	Deep Q-learning	A multizone office building	A variable air volume system
[25]	Multi-branching reinforcement learning	A multizone office building	A variable air volume system
[26]	Policy gradient	A multizone office building	Radiant heating systems

References

1. Zhao, J.; Lam, K.P.; Ydstie, B.E.; Loftness, V. Occupant-oriented mixed-mode EnergyPlus predictive control simulation. *Energy Build.* 2016, 117, 362–371.
2. Zhao, T.; Wang, J.; Xu, M.; Li, K. An online predictive control method with the temperature based multivariable linear regression model for a typical chiller plant system. *Build. Simul.* 2020, 13, 335–348.
3. Talib, R.; Nassif, N. “Demand Control” an Innovative Way of Reducing the HVAC System's Energy Consumption. *Buildings* 2021, 11, 488.
4. Dong, B.; Lam, K.P. A real-time model predictive control for building heating and cooling systems based on the occupancy behavior pattern detection and local weather forecasting. *Build. Simul.* 2014, 7, 89–106.
5. Ma, X.; Bao, H.; Zhang, N. A New Approach to Off-Line Robust Model Predictive Control for Polytopic Uncertain Models. *Designs* 2018, 2, 31.
6. Ascione, F.; Bianco, N.; De Stasio, C.; Mauro, G.M.; Vanoli, G.P. Simulation-based model predictive control by the multi-objective optimization of building energy performance and thermal comfort. *Energy Build.* 2016, 111, 131–144.
7. Garnier, A.; Eynard, J.; Caussanel, M.; Grieu, S. Predictive control of multizone heating, ventilation and air-conditioning systems in non-residential buildings. *Appl. Soft Comput.* 2015, 37, 847–862.

8. Wang, C.; Wang, B.; Cui, M.; Wei, F. Cooling seasonal performance of inverter air conditioner using model prediction control for demand response. *Energy Build.* 2022, 256, 111708.
9. Váňa, Z.; Cigler, J.; Šíroký, J.; Žáčková, E.; Ferkl, L. Model-based energy efficient control applied to an office building. *J. Process Control* 2014, 24, 790–797.
10. Kumar, R.; Wenzel, M.J.; ElBsat, M.N.; Risbeck, M.J.; Drees, K.H.; Zavala, V.M. Stochastic model predictive control for central HVAC plants. *J. Process Control* 2020, 90, 1–17.
11. Toub, M.; Reddy, C.R.; Razmara, M.; Shahbakhti, M.; Robinett, R.D.; Aniba, G. Model-based predictive control for optimal MicroCSP operation integrated with building HVAC systems. *Energy Convers. Manag.* 2019, 199, 111924.
12. Kwak, Y.; Huh, J.-H.; Jang, C. Development of a model predictive control framework through real-time building energy management system data. *Appl. Energy* 2015, 155, 1–13.
13. Liu, S.; Henze, G.P. Experimental analysis of simulated reinforcement learning control for active and passive building thermal storage inventory: Part 2: Results and analysis. *Energy Build.* 2006, 38, 148–161.
14. Liu, S.; Henze, G.P. Evaluation of Reinforcement Learning for Optimal Control of Building Active and Passive Thermal Storage Inventory. *J. Sol. Energy Eng.* 2006, 129, 215–225.
15. Jayalaxmi, P.L.S.; Saha, R.; Kumar, G.; Kim, T.-H. Machine and deep learning amalgamation for feature extraction in Industrial Internet-of-Things. *Comput. Electr. Eng.* 2022, 97, 107610.
16. Chen, Y.; Tong, Z.; Zheng, Y.; Samuelson, H.; Norford, L. Transfer learning with deep neural networks for model predictive control of HVAC and natural ventilation in smart buildings. *J. Clean. Prod.* 2020, 254, 119866.
17. Othman, K. Deep Neural Network Models for the Prediction of the Aggregate Base Course Compaction Parameters. *Designs* 2021, 5, 78.
18. Mnih, V.; Kavukcuoglu, K.; Silver, D.; Graves, A.; Antonoglou, I.; Wierstra, D.; Riedmiller, M. Playing Atari with Deep Reinforcement Learning. *arXiv* 2013, arXiv:1312.5602.
19. Dalamagkidis, K.; Kolokotsa, D.; Kalaitzakis, K.; Stavrakakis, G.S. Reinforcement learning for energy conservation and comfort in buildings. *Build. Environ.* 2007, 42, 2686–2698.
20. Fazenda, P.; Veeramachaneni, K.; Lima, P.; O'Reilly, U.-M. Using reinforcement learning to optimize occupant comfort and energy usage in HVAC systems. *J. Ambient. Intell. Smart Environ.* 2014, 6, 675–690.
21. Capozzoli, A.; Piscitelli, M.S.; Gorrino, A.; Ballarini, I.; Corrado, V. Data analytics for occupancy pattern learning to reduce the energy consumption of HVAC systems in office buildings. *Sustain. Cities Soc.* 2017, 35, 191–208.
22. Costanzo, G.T.; Iacovella, S.; Ruelens, F.; Leurs, T.; Claessens, B.J. Experimental analysis of data-driven control for a building heating system. *Sustain. Energy Grids Netw.* 2016, 6, 81–90.
23. Fang, J.; Feng, Z.; Cao, S.-J.; Deng, Y. The impact of ventilation parameters on thermal comfort and energy-efficient control of the ground-source heat pump system. *Energy Build.* 2018, 179, 324–332.
24. Yuan, X.; Pan, Y.; Yang, J.; Wang, W.; Huang, Z. Study on the application of reinforcement learning in the operation optimization of HVAC system. *Build. Simul.* 2021, 14, 75–87.
25. Ding, X.; Du, W.; Cerpa, A. OCTOPUS: Deep Reinforcement Learning for Holistic Smart Building Control. In *Proceedings of the 6th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (BuildSys '19)*, New York, NY, USA, 13–14 November 2019.
26. Zhang, Z.; Chong, A.; Pan, Y.; Zhang, C.; Lam, K.P. Whole building energy model for HVAC optimal control: A practical framework based on deep reinforcement learning. *Energy Build.* 2019, 199, 472–490.