INCREASING BUILDING ENERGY EFFICIENCY THROUGH DEEP REINFORCEMENT LEARNING

Hans-Kristian Ovanger, NTNU, hanskrio@stud.ntnu.no
Anne Neumann, NTNU, anne.neumann@ntnu.no
Ruud Egging-Bratseth, SINTEF, ruud.egging@sintef.no

Overview and Method

This paper report explores how artificial intelligence (AI), specifically deep reinforcement learning (DRL) algorithms, can optimize the control of compressor frequency in heat pumps. Traditionally, proportional-integral (PI) control is employed to regulate compressor frequency, relying on complex, predefined mathematical models. In contrast, this project report leverages DRL for model-free learning, allowing the AI algorithm to autonomously learn optimal control strategies without prior knowledge of heat pump dynamics or the specific problem. This approach eliminates the dependency on computationally intensive mathematical frameworks, making it significantly more scalable and applicable across diverse energy systems. Deep Q-Networks (DQN) are trained to directly adjust the compressor frequency signal, aiming to minimize energy consumption and costs while maintaining thermal comfort. Through experiments conducted in BOPTEST, a standardized test environment for building energy systems and control strategies, DQN is compared with PI control using key performance indicators such as energy use, emissions, and thermal comfort.

Results

Preliminary reults demonstrate that artificial intelligence, through DRL, can effectively learn to control compressor frequency signals in heat pumps without any prior knowledge or domain-specific models. Moreover, the results highlight the potential for cost and energy savings that surpass the capabilities of traditional PI control.

Figure 1: Comparison of KPIs between baseline controller and the DRL agent

	Key Performance Indicators			
Controller	Discomfort (Kh/zone)	Energy Use (kWh/m^2)	Cost (€/m^2)	Emissions (CO_2 kg)
Baseline	8.382	3.478	0.8816	0.5808
DQN	**6.160**	**3.432**	**0.8701**	**0.5732**

(a) Peak heating period

	Key Performance Indicators			
Controller	Discomfort (Kh/zone)	Energy Use (kWh/m^2)	Cost (€/m^2)	Emissions (CO_2 kg)
Baseline	**9.446**	**1.773**	**0.4495**	**0.2961**
DQN	13.92	1.778	0.4506	0.2969

(b) Typical heating period

Figure 2: Control actions and performance over two test periods (weeks)

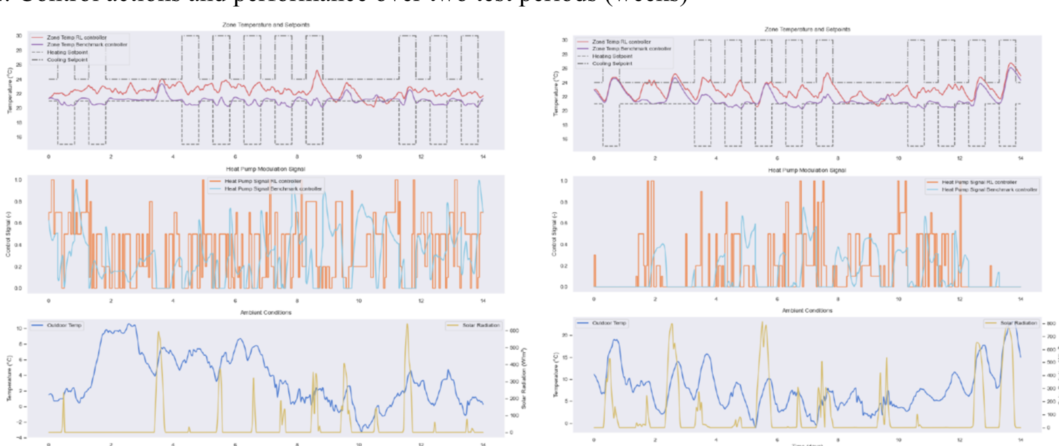

(a) January 17 - January 31: Peak heating period

(b) April 19 - May 3: Typical heating period

Conclusions

The results suggest that the utise of artificial intelligence can be a scalable and effective solution for sustainable, data-driven strategies in future energy management systems. The Deep Q-Network algorithm successfully learns control strategies for residential heat pumps, evidenced by the increasing trend in the mean reward graph in Figure 2. The agent's ability to adapt through learning provide evidence for the applicability of reinforcement learning to the heat pump control problem. The RL agent outperforms the baseline controller during colder periods as shown in Figure 1(a), indicating its ability to optimize heating strategies under high heating-demand conditions. This highlights the potential for improving energy efficiency and reducing operational costs during times of peak energy consumption. The agent's underperformance during warmer periods, as illustrated in Figure 2(b), is attributed to unanticipated high solar irradiation, showing a notable limitation in its current design. We are working on improving this and hope to be able to present final results during the conference.

References

T. Wei, Y. Wang, and Q. Zhu, "Deep reinforcement learning for building HVAC control," in Proc. 54th Annual Design Automation Conference (DAC), 2017, pp. 1–6.

M. Biemann, F. Scheller, X. Liu, and L. Huang, "Experimental evaluation of model-free reinforcement learning algorithms for continuous HVAC control," Applied Energy, vol. 298, p. 117164, 2021.

Z. Wang and T. Hong, "Reinforcement learning for building controls: The opportunities and challenges," Applied Energy, vol. 269, p. 115036, 2020.

L. Yu, S. Qin, M. Zhang, C. Shen, T. Jiang, and X. Guan, "A review of deep reinforcement learning for smart building energy management," IEEE Internet of Things Journal, vol. 8, no. 15, pp. 12046–12063, 2021.

D. Blum, J. Arroyo, S. Huang, et al., "Building optimization testing framework (BOPTEST) for simulation-based benchmarking of control strategies in buildings," Journal of Building Performance Simulation, vol. 14, no. 5, pp. 586–610, 2021.

Y. Yao and D. K. Shekhar, "State of the art review on model predictive control (MPC) in heating ventilation and air-conditioning (HVAC) field," Building and Environment, vol. 200, p. 107952, 2021.

A. Bloess, W.-P. Schill, and A. Zerrahn, "Power-to-heat for renewable energy integration: A review of technologies, modeling approaches, and flexibility potentials," *Applied Energy*, vol. 212, pp. 1611–1626, 2018.

V. Mnih et al., "Human-level control through deep reinforcement learning," Nature, vol. 518, pp. 529–533, 2015.