Doctoral Dissertation
Doctoral Program in Energetics (35^{th}cycle)

Scaling energy management in buildings with artificial intelligence

By

Giuseppe Pinto

******

Supervisor(s):
Prof. A. Capozzoli, Supervisor

Doctoral Examination Committee:
Prof. Enrico Fabrizio, Politecnico di Torino
Prof. Cheng Fan, Shenzhen University
Prof. Rongling Li, Technical University of Denmark - DTU
Prof. Adolfo Palombo, Università degli studi di Napoli Federico II
Dr. Marco Pritoni, Lawrence Berkeley National Laboratory - LBNL

Politecnico di Torino
2023

Declaration

I hereby declare that, the contents and organization of this dissertation constitute my own original work and does not compromise in any way the rights of third parties, including those relating to the security of personal data.

Giuseppe Pinto
2023

* This dissertation is presented in partial fulfillment of the requirements for **Ph.D. degree** in the Graduate School of Politecnico di Torino (ScuDo).

Abstract

The growing adoption of automation and control systems, and internet of things sensors in smart buildings has contributed to the unprecedented availability of monitoring data of the built environment, that could enable the deployment of Energy Management and Information Systems (EMIS) at scale. This dissertation aims at analyzing the potentialities provided by the exploitation of Artificial Intelligence (AI) techniques to scale EMIS, identifying promising directions and potential barriers for its real-world application. In this context, Grid-Interactive Efficient Buildings (GEB) are ideal candidates for the application of advanced energy management strategies. GEBs are energy-efficient buildings that uses smart technologies to provide demand flexibility while co-optimizing for energy cost, grid services, and occupant needs. However, when energy management is faced with shifting from a single building to multiple buildings, uncoordinated strategies for exploiting energy flexibility may have negative effects on the grid reliability, causing undesirable new peaks. The recent development of AI supported the creation of advanced data-driven control strategies, such as Deep Reinforcement Learning (DRL), however implementation focused on single buildings, neglecting the potentialities of applying this control strategy in multiple buildings. In this dissertation, three different applications that leveraged DRL at scale are conceived and tested. DRL is a control method based on the paradigm of learning from interaction, encoding the environment using deep neural networks. The developed applications used CityLearn, a simulation environment for the implementation of DRL in multiple buildings, focusing on 4 buildings equipped with thermal energy storage and renewables, benchmarking the DRL controller against a rule-based controller. In the first application, a centralized DRL controller was implemented to optimize the electrical demand profiles, reducing costs and peaks, understanding the effects of advanced control strategies at different scales (single building, district, grid). This application showed the potential of applying DRL in multiple buildings, achieving a 4% cost and 12% peak reduction. Then, a

second application analyzed the role of different reinforcement learning architectures, comparing a centralised (coordinated) controller and a decentralised (cooperative) controller to also consider different renewable energy systems. The two controllers reduced the costs by 3% and 7% respectively, and 10% and 14% respectively for peak demand. The study showed that the multi-agent cooperative approach may be more suitable for districts with heterogeneous objectives within the individual buildings. In a third application, the role of HVAC flexibility was investigated, exploiting deep neural networks to simulate the building thermal dynamics of the buildings, adding to the previously introduced framework the possibility to control the HVAC. In this case, the DRL controller was conceived to optimise the electrical demand profiles and provide services to the grid without penalising indoor comfort conditions. The developed DRL controller reduced the overall district electricity costs, while decreasing the peak energy demand by 23% and the Peak to Average Ratio by 20%, without penalizing indoor temperature control. The third application showed how deep neural networks are effective as a lightweight data-driven model to predict building thermal responses, highlighting their reliance on a large amount of data, that clashes with the potential limited data availability in most existing buildings. Therefore, the last application focused on data-driven models, identifying in transfer learning a way to overcome data reliance, describing its role in supporting building energy management. The application conducted a suite of experiments that leveraged 250 data-driven models based on a synthetic dataset of a building to study the influence of several features, isolating their contribution. The performance of the transfer learning process was compared against a classical machine learning approach, identifying guidelines for its application in buildings. Lastly, findings and outcomes of the present research study were discussed, providing a robust reasoning on the application of DRL controllers at large scale and how data-driven models can boost their adoption. Eventually, a wide overview on the lessons learned is proposed, outlining the future opportunities and barriers of scaling energy management in buildings using artificial intelligence.