ABSTRACT

In recent years, energy management in buildings has gained attention, as buildings account for a significant share of global energy consumption. Heating, Ventilation, and Air Conditioning (HVAC) systems are among the largest energy consumers and are characterized by non-linear and stochastic dynamics, making robust and adaptive control strategies essential. Traditional Rule-Based Controllers (RBCs) highlighting the need for advanced strategies since they are reactive, unable to balance multiple objectives and to adapt to dynamic changes in real-time. Model Predictive Control (MPC) emerged as a solution for its ability to optimize energy system operations over a shifting time horizon, accounting for both current and future building dynamics. However, deploying MPC in real buildings is constrained by its model-based nature, relying on accurate building systems characterization to achieve optimal performance. Deep Reinforcement Learning (DRL) offers a promising solution by autonomously learning control policies through a trial-and-error approach. Ideally, DRL controllers would be implemented directly in real buildings, but require an initial training phase that could increase energy consumption and discomfort, making direct deployment impractical. Therefore, DRL controllers are typically pre-trained offline using surrogate models emulating building dynamics. However, this approach is not scalable as buildings vary in geometry, envelope properties, and operating conditions, making model development complex and time-consuming. This dissertation proposes three methodologies, applied in five applications, to enhance DRL scalability for real-world deployment. Three applications were evaluated in simulations using EnergyPlus and Python, while the last two were tested in a real office building.

The first methodology employs Long Short-Term Memory networks to approximate building dynamics, allowing DRL pre-training to manage the supply water temperature in a heating system. This approach reduced indoor temperature violations by 80% compared to RBC, with similar energy consumption. The second methodology introduces Online Transfer Learning (OTL) to share DRL policies between buildings with similar (homogeneous TL) or different (heterogeneous TL) energy systems. In heterogeneous TL, energy systems include photovoltaic panels and a battery. The second and third applications evaluate OTL for managing a cooling system with a chiller and thermal storage. In homogeneous TL, OTL improved temperature conditions by 50-80% and reduced electricity cost by up to 40% compared to RBC and online DRL. In heterogeneous TL, OTL lowered electricity cost by 10% and improved temperature control by 10-40%, maximizing self-sufficiency and self-consumption by

9-11% compared to RBC and online DRL. Although OTL performed slightly worse than offline DRL, it required no modeling effort, making it a more practical and scalable solution. The fourth application tests behavioral cloning (i.e., third methodology), in a real office building in Switzerland. Here, a DRL controller first mimics the RBC operation and then manages the energy systems while actively improving its policy. This approach reduced energy consumption by 40% and improved temperature control by 43% when implemented to manage the valve opening for a Thermally Activated Building System. The fifth application demonstrates the real-world implementation of OTL, transferring a DRL controller between two similar office spaces with different boundary conditions, considering the same action space as the fourth application. The OTL reduced energy consumption by 20-40% and improved temperature control by 30-60% compared to RBCs, PI, and a DRL controller without pre-training. In conclusion, this dissertation highlights key findings and suggests future directions for TL in DRL control applications, including testing over longer periods and developing methods to assess domain similarity, paving the way for more efficient energy management in buildings.