

ABSTRACT

Advanced building control algorithms, such as Model Predictive Control (MPC) and reinforcement learning (RL), hold significant promise for improving energy efficiency without sacrificing occupant comfort. Yet, real-world implementation remains limited. MPC faces challenges stemming from scarce high-quality training data and the need for highly customized models of building system dynamics. Meanwhile, RL struggles with prolonged training times, sensitivity to hyper-parameter tuning, and difficulties in transferring learned policies across buildings with diverse characteristics. Transfer learning (TL) in RL can potentially improve sample efficiency by reusing knowledge from related domains; however, it often suffers from negative transfer and model mismatches.

To address these challenges, this Thesis first introduces an AI-enabled framework that automates the creation of domain-randomized environments to enhance policy robustness. By intentionally varying model parameters, domain randomization generates a diverse set of building models, moving beyond single deterministic representations. This strategy exposes RL agents to a wider range of simulated conditions during training, improving generalization. Specifically, this work constructs a model universe (MU), a probability distribution over plausible building models inferred from incomplete input data, thereby automating model generation without requiring additional expert input. Two case studies on office buildings demonstrate the approach. The first case illustrates how the MU captures variability across regulatory standards, while the second applies RL to a real-world building control task using environments sampled from the MU. Results show that policies trained on MUs improve the 95th percentile of RL reward by 17% compared to policies trained on a single static model, confirming enhanced robustness and performance.

Second, this Thesis presents an efficient TL framework that mitigates negative transfer and model mismatches by leveraging the MU as the source domain and applying an online heuristic policy selection strategy. The framework comprises three main steps: (1) generating a set of plausible source buildings from the MU; (2) pretraining candidate RL policies using a domain-adapted deep Q-learning algorithm for building control (DQN-BC); and (3) selecting the most effective candidate policies through brief deployment in the target building.

Experimental results validate the proposed DQN-BC algorithm's superior performance over standard DQN. Applying the TL framework to both single-zone and multi-zone office building

scenarios, with Department of Energy reference models as ground truth, shows that selected policies achieve up to 11.46% and 50.58% improvements in jumpstart rewards and asymptotic gains of approximately 49.98% and 37.42%, respectively. The approach closely matches the performance of policies trained directly on true models while requiring significantly less deployment time, achieving strong performance within two days and outperforming baseline methods. These results demonstrate that the proposed approach, which integrates domain randomization with transfer learning, enables an adaptable and data-efficient reinforcement learning framework for building control, thereby advancing practical deployment in complex real-world settings.