A Reinforcement Learning Approach to Home Energy Management for Modulating Heat Pumps and Photovoltaic Systems

Lissy Langer, Technische Universität Berlin, Working Group Production and Operations Management, langer@pom.tu-berlin.de
Thomas Volling, Technische Universität Berlin, Working Group Production and Operations Management, volling@pom.tu-berlin.de

Overview

Buildings are one of the main drivers of global energy consumption and CO_2 emissions. Efficient energy management systems will have to integrate renewable energy sources with heat supply to mitigate climate change. In this study, we analyze the potential of deep reinforcement learning (DRL) to control a smart home with a modulating air-sourced heat pump, a photovoltaic system, a battery, as well as a thermal storage system for floor heating and hot-water supply. We describe a structured solution path to transform a mixed-integer linear program (MILP) into a DRL implementation. In our numerical analysis, we benchmark our results based on the deep deterministic policy gradient (DDPG) algorithm with the optimal model predictive control (MPC) result using a MILP under full information, as well as a practice-oriented rule-based approach.

Methods

We use the deep reinforcement learning (DRL) algorithm called deep deterministic policy gradient (DDPG) to construct an autonomous agent for a smart home energy management system (SHEMS). We benchmark our results to a model predictive control (MPC) implementation and a rule-based approach.

Results

This paper describes how to transform a SHEMS MPC implementation (MILP formulation) into an RL environment in order to solve the problem using the DDPG algorithm. We formulate the problem as an MDP with an action representation and RL environment informed by domain-knowledge. We propose a structured solution path to achieve stable policy results.We show that our proposed DRL implementation outperforms the rule-based approach and achieves an autarky level of 86% with only limited comfort violations. Analyzing different DRL formulations, we conclude that domain-knowledge is key to formalizing an efficient decision problem with stable results. Our input data and models developed using the Julia programming language, are available open-source.

Conclusions

We find significant improvements compared to the rule-based approach used in practice. We introduce sensible reductions of our action space and thereby significantly unburden the learning process. A structured and transparent solution approach is necessary to ensure assessability of the model results. Our findings are in line with data-centered AI, a recent movement, which stresses that the keys to successful real-world application are domain-knowledge and higher quality of (data) preparation—and not necessarily incremental algorithm improvement. Overall, we find great potential to controlling a SHEMS using DRL, as this combines the data-centric strength of machine learning with the modeling capabilities of OR.