Enhancing the Performance of Reinforcement Learning in Building Energy Management through Advanced Exploration Strategies

Background and Objective

Efficient energy management in buildings plays a key role in reducing operational costs, optimizing energy consumption, and improving the integration of renewable energy sources into the power grid. The research group Combined Smart Energy Systems (CoSES) provides a state-of-the-art research environment for developing advanced reinforcement learning (RL) methods for building energy management systems (BEMS).

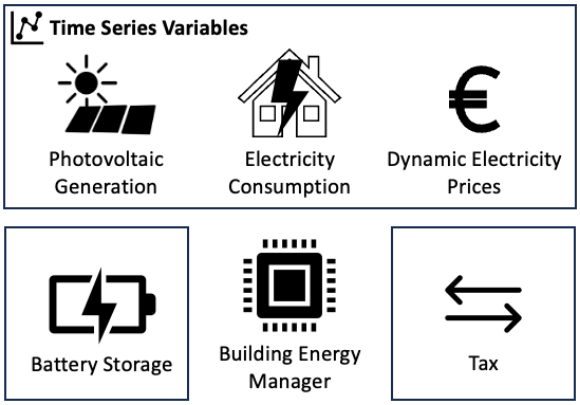

Figure 1: Basic example of a BEMS.

Current approaches such as Model Predictive Control (MPC) yield mathematically optimal results under idealized conditions with perfect forecasts. However, in practice, these methods often lack the flexibility required for widespread deployment in real-world buildings. RL methods, which are inherently more adaptive, have the theoretical potential to reach near-optimal solutions through exploration. Yet, their applicability in BEMS is significantly hindered by highly volatile energy prices, which directly impact the reward signal and introduce substantial noise. This issue is well known in RL research as the “sparse reward” problem and poses a major challenge for standard algorithms such as Proximal Policy Optimization (PPO), which tend to get stuck in local optima.

A promising and generally applicable solution to this challenge lies in advanced exploration strategies. These methods directly address the sparse reward problem and specifically enhance the ability of RL algorithms to learn meaningful policies in complex, noisy environments. As a result, they offer great potential not only for BEMS but also for a wide range of real-world RL applications.

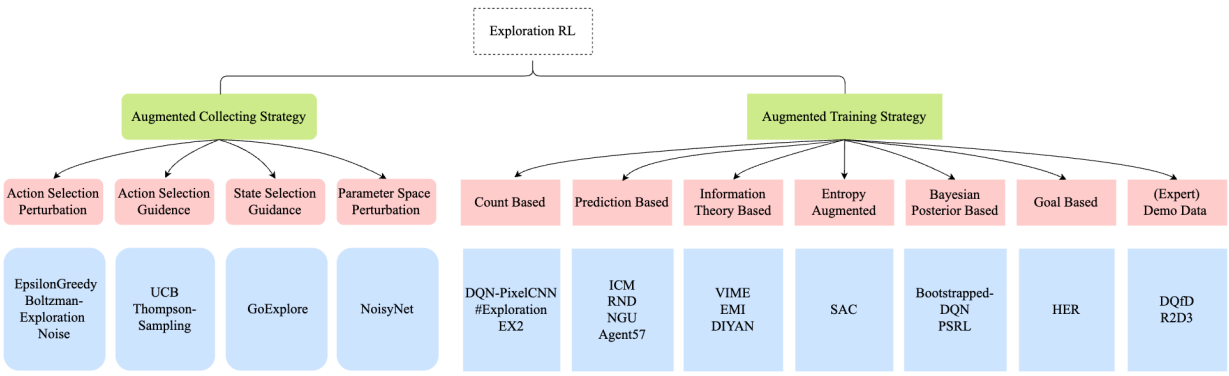

Figure 2: Taxonomy of exploration methods in RL.

Project Goal

The goal of this project is to investigate, develop, and implement advanced exploration strategies to improve the performance of RL algorithms in the context of BEMS under dynamic energy pricing.

1. **Literature Review and Domain Understanding**: Conduct a comprehensive literature review on recent exploration strategies for sparse reward problems in RL, with a focus on applications in building energy management.
2. **Development**: Identify and develop suitable exploration strategies to address the sparse reward challenge.

3. **Evaluation**: Systematically evaluate the developed strategies in comparison to established RL approaches (e.g., PPO) and MPC methods. Optimize the exploration strategies to enhance learning performance.
4. **Validation (Optional)**: Apply and validate the optimized exploration strategies in realistic scenarios using the CoSES lab environment, assessing their practical feasibility, robustness, and adaptability.

The results are expected to make a significant contribution to bringing RL methods into real-world BEMS. The overarching goal is to simplify energy management in buildings and thereby actively support the renewable energy transition.

Requirements

- Strong programming skills, preferably in Python.
- Solid background in machine learning, preferably with experience in RL.
- An understanding of BEMS is desirable but not mandatory.
- Excellent problem-solving abilities and a willingness to learn and adapt to new challenges.

Learning Outcomes

- Gain hands-on experience with advanced RL techniques, especially exploration strategies.
- Acquire a deep understanding of the challenges and solutions associated with building energy management.
- Strengthen skills in literature review, modeling, implementation, testing, and validation.
- Have the Opportunity to contribute to cutting-edge research with high practical relevance in BEMS and beyond.

Contact

Ulrich Ludolfinger, M.Sc.

Chair of Renewable and Sustainable Energy Systems (Prof. Dr. rer. nat. T. Hamacher)

Tel: +49 (0)871 - 506 336, Email: ulrich.ludolfinger@tum.de