Building energy efficiency recommendations with reinforcement learning

Odyssefs Diamantopoulos*
diamandopoulos@aueb.gr

Spiros Chadoulos*†
spiroscha@aueb.gr
sc@pleg.ma

Iordanis Koutsopoulos*
jordan@aueb.gr

George C. Polyzos*
polyzos@aueb.gr

Nikolaos Ipiotis†
ni@pleg.ma

Abstract

With buildings accounting for a significant portion of the grid's total energy demand, it is evident that consumers should be engaged in energy efficiency. In this work, we propose a reinforcement learning approach that conducts energy efficiency recommendations for buildings, in the form of load-shift suggestions for different devices/assets. The adopted methodology can continuously learn consumer energy behavior and preferences to minimize energy costs, while preserving comfort by jointly training a single agent for all the building assets. The agent utilizes user feedback on the recommendations and integrates it in the reward function. Preliminary experiment results with simulated data show that the agent's reward is increasing throughout time.

CCS Concepts: • **Computing methodologies** → **Machine learning**; • **Hardware** → **Smart grid**.

Keywords: Smart buildings, reinforcement learning

ACM Reference Format:
Odyssefs Diamantopoulos, Spiros Chadoulos, Iordanis Koutsopoulos, George C. Polyzos, and Nikolaos Ipiotis. 2023. Building energy efficiency recommendations with reinforcement learning. In *Proceedings of 5th Summit on Gender Equality in Computing (GEC '23)*. ACM, New York, NY, USA, 2 pages. https://doi.org/XXXXXXX.XXXXXXX

*Mobile Multimedia Laboratory, Department of Informatics, School of Information Sciences and Technology, Athens University of Economics and Business, Patission 76, 104 34, Athens, Greece

†Plegma Labs S.A., Neratziotissis 115, 151 24, Marousi, Greece

GEC '23, June 26–27, 2023, Athens, Greece

ACM ISBN 978-1-4503-XXXX-X/18/06...$15.00
https://doi.org/XXXXXXX.XXXXXXX

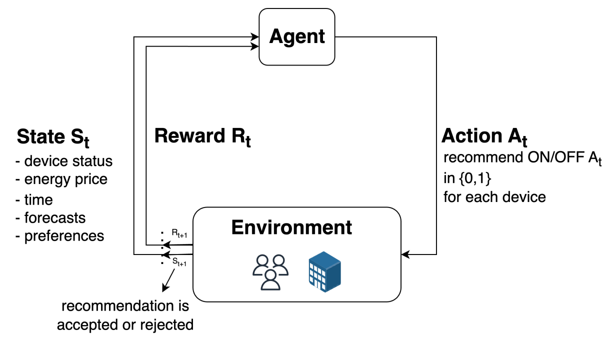

Figure 1. RL model formulation

1 Introduction and related work

Consumers play a critical role in the electricity grid, assuming a direct involvement, while buildings bear the responsibility for roughly one-third of the world's energy consumption, with residential buildings contributing to 22% of the global energy demand [1]. Engaging consumers toward energy efficiency through appropriate recommendations is crucial for the EU's energy transition [2], with the recent advancements in AI, such as Reinforcement Learning (RL), playing an important role. O'Neill et al. [5] proposed a Consumer Automated Energy Management System that takes into account both electricity prices and consumer energy demand. Mocanu et al. [4] implemented an On-Line Building Energy Management System (EMS) using Deep Reinforcement Learning. The RL agent makes multiple actions at each time-step related to switching ON/OFF events of different types of electrical devices. In [6], the authors design and validate an EMS that utilizes deep reinforcement learning along with occupant feedback and activities. The system controls different types of smart devices with the goal of minimizing the home's total energy cost plus the residents' discomfort.

2 Model and preliminary results

We propose a novel formulation for the problem of energy efficiency recommendations in buildings using RL, depicted in Figure 1. The RL agent continuously learns the energy behavior of end-users through various data sources, including

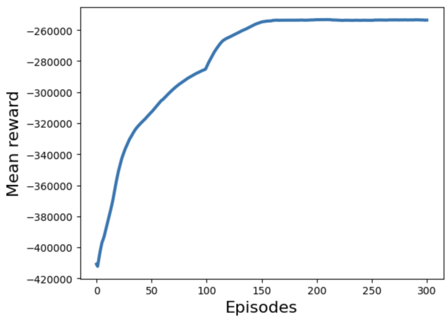

Figure 2. Mean Reward

historical and real-time data, energy prices, and comfort parameters. The agent then offers load shift recommendations to increase energy efficiency while preserving user comfort. The state S_t at each time step t incorporates multiple parameters related to the state of the building. Specifically, it includes the current state of all building assets (e.g. if they are ON/OFF), the energy price, energy consumption forecasts which can help the agent understand the future behavior of the consumers [3], and specific user preferences regarding the usage of the devices. The agent's action A_t at each step is the energy efficiency recommendation for the assets/devices of the building in the form of ON/OFF or specific power consumption values depending on the type of each asset/device. The reward R_t that the agent receives from the system (environment) at each timestep includes parameters such as the total device energy costs and the user's feedback (i.e. if the recommendation was followed from the asset/device energy demand). While existing works utilize a separate agent for each device, we adopt a building-level approach that jointly takes into account all the available devices/assets at the same time since multiple correlations between devices can appear, with the usage of one device affecting another. The envisioned system can swiftly adapt to the user's behavior and preferences while at the same time reducing the electrical bill of the household without compromising comfort.

In Figure 2, preliminary experiment results for a simplified version of the described model are presented. Specifically, the Advantage Actor-Critic (A2C) algorithm is utilized to train an agent for 300 episodes, with each episode having 8400 hourly time slots. A simulated user is integrated into the environment, with separate transition probabilities between two states (ON/OFF) for each device (5 devices in total). In addition, real energy price signals from NYISO are used for the reward function, which is calculated as follows: For each of the values of the agent's action A_t (representing a recommendation for one of the devices) that match the actual device state, meaning that the recommendation is accepted, the device reward is $-p_t P_t^d$, where p_t is the energy price and P_t^d the energy consumption of device d for time slot t. On the other hand, if the action's value does not match the device state (the recommendation was rejected or ignored) the device reward is set to -100 which is a hyperparameter. All the device rewards are added to form the total reward R_t for time slot t. It is evident from Figure 2 that the mean reward increases as the episodes pass, meaning that the agent learns a policy that minimizes the energy cost while following user preferences.

3 Conclusion and future work

In this work, an energy efficiency recommendations approach is presented, utilizing an RL model to train an agent that makes device load-shift suggestions to consumers. The RL agent is capable of continuously learning the energy behavior of the building occupants regarding device usage while taking into account their preferences in order to minimize energy costs and preserve comfort. In terms of future work, we plan to release a GYM environment that implements the model and formulation presented, in order to provide a fully customizable testbed to researchers who want to test their own RL approaches in a simulated building environment. In addition, experiments will be conducted to determine the specific environment parameters, states, rewards, and action types, while a demonstration with real buildings will be carried out.

Acknowledgments

This work was supported by the EMERGENT project which has received funding from the 2nd I-NERGY open call under Grant 101016508.

References

[1] Thibaut Abergel, Brian Dean, and John Dulac. 2017. UN Environment and International Energy Agency (2017): Towards a zero-emission, efficient, and resilient buildings and construction sector. *Global Status Report* (2017).

[2] Spiros Chadoulos, Iordanis Koutsopoulos, and George C Polyzos. 2020. Mobile Apps Meet the Smart Energy Grid: A Survey on Consumer Engagement and Machine Learning Applications. *IEEE Access* 8 (2020), 219632–219655.

[3] Spiros Chadoulos, Iordanis Koutsopoulos, and George C Polyzos. 2021. One model fits all: Individualized household energy demand forecasting with a single deep learning model. In *Proceedings of the Twelfth ACM International Conference on Future Energy Systems*. 466–474.

[4] Elena Mocanu, Decebal Constantin Mocanu, Phuong H Nguyen, Antonio Liotta, Michael E Webber, Madeleine Gibescu, and Johannes G Slootweg. 2018. On-line building energy optimization using deep reinforcement learning. *IEEE Transactions on Smart Grid* 10, 4 (2018), 3698–3708.

[5] Daniel O'Neill, Marco Levorato, Andrea Goldsmith, and Urbashi Mitra. 2010. Residential demand response using reinforcement learning. In *2010 First IEEE international conference on smart grid communications*. IEEE, 409–414.

[6] Salman Sadiq Shuvo and Yasin Yilmaz. 2022. Home energy recommendation system (hers): A deep reinforcement learning method based on residents' feedback and activity. *IEEE Transactions on Smart Grid* 13, 4 (2022), 2812–2821.