Current Trends in Computer Sciences & Applications
DOI: 10.32474/CTCSA.2023.02.000145
ISSN: 2643-6744
Research Article

Reinforcement Learning, A Promising Approach to The Management of Renewable Energy Communities

Giulia Palma*, Leonardo Guiducci and Antonio Rizzo

Department of Social, Political and Cognitive Sciences, University of Siena, Siena, Italy

***Corresponding author:** Giulia Palma, Department of Social, Political and Cognitive Sciences, University of Siena, Italy

Received: May 19, 2023 **Published:** June 01, 2023

Abstract

As the demand for sustainable energy solutions continues to rise, the effective management of energy flows in interconnected systems has become a critical challenge. Energy communities, comprising diverse energy sources and consumers, necessitate intelligent control strategies to optimize energy distribution and consumption. In recent years, Reinforcement Learning (RL) algorithms have emerged as a powerful tool for addressing complex decision-making problems in various domains. This paper presents an opinion on the relationship between Reinforcement Learning and Optimal Control theory for managing energy flows within energy communities, highlighting the potential of RL-based algorithms in this context. Moreover, it provides insights into the current trends in computer sciences and applications, emphasizing the significance of leveraging RL techniques for energy flows optimization.

Index Terms

Reinforcement learning; energy community; social welfare; energy management; online scheduling; optimal control

Introduction

Renewable energy communities RECs have gained significant attention as an innovative concept that promotes the integration of renewable energy sources, energy storage systems, and intelligent energy management strategies. As defined in [1], RECs are groups of individuals and organizations that work together to promote and use renewable energy sources, such as solar, wind, and hydroelectric power. A crucial issue within energy communities is the proper management of energy flows to ensure a balance between energy supply and demand. This problem has been addressed with various methods like classic optimization methods, heuristic algorithms and rule-based con- trollers [2,3]. Optimal control techniques play a pivotal role in ensuring efficient utilization of available resources and minimizing energy waste. However, such algorithms require perfect knowledge of the generation and demand profiles in order to provide the optimal solution to the problem. But a critical challenge in optimal scheduling of energy flows is the significant intermittency and stochasticity in renewable generation as well as the electricity demand. For this reason, much attention is shifting towards approaches based on Reinforcement Learning (RL), since this technique has demonstrated its effectiveness in addressing decision-making challenges in dynamic and uncertain environments. This opinion piece explores the relationship between RL algorhitm and optimal control theory in managing energy flows within RECs.

Efficient Energy Management: Reinforcement Learning Vs. Optical Control

Energy communities exhibit complex dynamics influenced by factors such as energy production, storage capacities, consumption patterns, and environmental conditions. The management of energy flows within RECs has seen significant advancements in recent years. Traditional approaches, such as rule-based controllers and optimal control techniques, have been commonly used but often struggle to adapt to dynamic and uncertain environments. The emergence of RL algorithms has revolutionized the field by offering adaptable and data- driven solutions for energy management. RL-based approaches have showcased their ability to optimize energy flows, improve efficiency, and integrate renewable energy sources effectively. Indeed, reinforcement learning algorithms can capture the system's state and define appropriate action spaces, enabling intelligent decision-making to optimize energy distribution and consumption. Designing suitable reward functions is crucial for RL algorithms to learn optimal control policies. In the context of energy communities, rewards can be based on objectives such as minimizing energy costs, maximizing renewable energy utilization, or reducing carbon emissions. The integration of domain-specific knowledge into reward design enhances the efficiency and effectiveness of RL-based control strategies. Reinforcement learning algorithms balance exploration and exploitation to find the optimal control policies. In energy communities, exploration enables the discovery of efficient energy flows, while exploitation leverages the learned policies to optimize energy distribution and consumption in real-time. RL algorithms provide an adaptive framework to dynamically adjust exploration-exploitation trade-offs based on changing system conditions and objectives. Thus, when comparing a RL approach to an optimal control one for managing energy flows in energy communities, RL offers several advantages that make it superior in addressing the complexities and uncertainties of these systems. Optimal control techniques traditionally rely on accurate system models and assumptions about the dynamics of the energy community. While effective in well-defined and predictable environments, optimal control approaches often struggle to adapt to dynamic and evolving conditions, such as changes in renewable energy generation, demand patterns, or unexpected disruptions. On the other hand, RL algorithms have the ability to learn directly from interaction with the environment, making them well-suited for handling complex and uncertain systems like energy communities. RL models can dynamically adjust their control policies based on real-time feedback, allowing them to adapt to changing conditions and optimize energy flows accordingly. RL-based approaches are also capable of discovering non-intuitive control strategies that may not be evident through traditional optimization techniques. Furthermore, RL's exploration-exploitation mechanism enables it to continuously explore different control actions and learn from their outcomes, gradually improving its decision-making capabilities over time. In contrast, optimal control approaches often rely on predefined control strategies derived from mathematical models, limiting their ability to adapt to unforeseen scenarios or incorporate new information. Another advantage of RL is its flexibility in handling diverse and complex objective functions. Energy communities often have multiple competing objectives, such as cost minimization, renewable energy integration, and carbon emissions reduction. RL allows for the incorporation of various reward functions, enabling the optimization of multiple objectives simultaneously or dynamically adjusting priorities based on changing circumstances. However, it is important to note that RL also presents its own challenges, such as the need for substantial computational resources and potentially longer training times. Additionally, the interpretability of RL models can be a concern, as their decision-making processes are often seen as "black boxes" due to their complex nature. Addressing these challenges and developing techniques for model interpretability are active areas of research within the RL community.

Case Studies and Experimental Results

In our previous publication [4], we addressed the challenge of maximizing the social welfare of a renewable energy community that receives incentives for virtual self-consumption within the Italian energy framework. A key component in this optimization problem is the battery energy storage system (BESS), which plays a crucial role in balancing supply and demand. To tackle this problem, we proposed a novel Reinforcement Learning-based BESS controller that operates in real-time using only available data at the current time-step. Our approach aims to maximize the community's social welfare by effectively managing the energy flows within the community. Through extensive simulations across various scenarios, we demon-strated the effectiveness of our RL-based approach, showcasing its ability to surpass the performance of a state-of-the-art rule-based controller. Moreover, we evaluated our proposed approach by comparing its performance with an ideal optimal control policy [5] based on an oracle that possesses perfect knowledge of future data. While the oracle-based approach represents an unattainable ideal, it provides a benchmark for assessing the performance of our RL-based controller. The results demonstrated that our RL-based approach achieved com- petitive performance, even when compared to the ideal control policy. By leveraging RL's ability to learn from experience and make real-time decisions based on available data, our proposed approach showcases the potential of reinforcement learning in optimizing the management of renewable energy. communities. The results of our study validate the effectiveness of our RL-based BESS controller in maximizing social welfare, indicating its superiority over both rule-based controllers and a hypothetical optimal control policy based on perfect future data knowledge.

Challenges And Future Directions

While reinforcement learning holds great promise for optimal control in energy communities, several challenges must be addressed for wider adoption. These challenges include the need for robust and explainable RL models, handling large- scale systems, incorporating uncertainty and risk management, and considering the socioeconomic aspects of energy communities. Addressing these challenges will require interdisciplinary collaborations between computer scientists, energy experts, policymakers, and stakeholders. We claim that, while optimal control approaches have been traditionally employed for energy management, RL-based approaches demonstrate clear advantages in handling the complexities, uncertainties, and evolving nature of energy communities. RL's ability to learn from experience, adapt to changing conditions, explore novel control strategies, and optimize multiple objectives make it a promising and superior choice for managing energy flows in these interconnected systems. Reinforcement learning algorithms offer a powerful approach to optimal control in managing energy flows within energy communities. By leveraging RL's ability to learn from experience and optimize decision-making, energy communities can achieve higher energy efficiency, improved renewable energy integration, and reduced environmental impact. As we continue to develop and refine RL-based algorithms, interdisciplinary efforts and collaboration will play a pivotal role in advancing sustainable energy solutions for future energy communities. Ongoing research aims to address challenges related to interpretability, scalability, and real-world constraints to further enhance the state of the art in RL-based energy management for energy communities. Currently, we are facing the emerging challenge of adapting algorithms to accommodate the increasing scale of Community Energy Resources (CERs), which, as observed in other countries, can reach magnitudes of tens of thousands of units. The centralized solution previously proposed encounters implementation complexities as the number of nodes associated with a CER grows. Our objective is to overcome these limitations by advocating a decentralized approach based on the promising technique of federated learning, specifically Federated Reinforcement Learning (FRL).

FRL allows multiple nodes to anonymously share their information at the edge side to train a single machine learning model at the cloud side. Therefore, our goal is to develop and test FRL-based algorithms for energy optimization within the CER framework, ensuring both the confidentiality of CER user data and the scalability of the algorithms through a distributed edge-cloud architecture. Our ultimate aim is to effectively govern the real-time energy flows of CERs, enhance energy efficiency, and maximize the utilization of renewable energy sources. This can be achieved by deploying FRL algorithms that are disseminated among the community nodes through dedicated hardware/software kits. In our study, we have demonstrated that our Reinforcement Learning-based approach achieves energy management performance comparable to that of Optimal control. This indicates its capability to handle situations of "ordinary" uncertainty, where the consumption production dataset exhibited unpredictability but lacked drastically critical episodes or dramatic variations. However, as we scale up and encounter a larger number of nodes, the frequency of critical events, anomalies, and exceptional circumstances increases. Consider more complex scenarios, such as controlled areas of the grid experiencing emergencies or unexpected changes. To address these challenges, we will investigate the application of federated learning techniques. Therefore, a promising future perspective is to assess the extent to which Federated Reinforcement Learning can effectively handle such scenarios. Furthermore, we will not only evaluate the performance of the Federated Reinforcement Learning approach in situations with ordinary variability arising from changes in energy production and consumption but also assess its effectiveness in managing emergencies and significant shifts in production or consumption.

References

1. Directive (EU) (2018) 2018/2001 of the European Parliament and of the Council on the promotion of the use of energy from renewable sources. Official Journal of the European Union.
2. S Parhizi, H Lotfi, A Khodaei, and S Bahramirad (2015) State of the art in research on microgrids: A review. IEEE Access 3(1): 890-925.
3. M F Zia, E Elbouchikhi, and M Benbouzid (2018) Microgrids energy man- agement systems: A critical review on methods, solutions, and prospects. Applied Energy 222(6): 1033-1055.
4. L Guiducci, G Palma, M Stentati, A Rizzo, S Paoletti (2023) A re- inforcement learning approach to the management of renewable energy communities. Mediterranean Conference on Embedded Comput- ing (MECO).
5. M Stentati, S Paoletti, A Vicino (2022) Optimization of energy communi- ties in the italian incentive system. IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT-Europe) pp. 1-5.

DOI: 10.32474/CTCSA.2023.02.000145