Internship position: Human-in-the-loop Multi-Agent Reinforcement Learning for Self-Adaptive Energy Systems

Context & objectives

Many problems in energy management cannot be solved by individual components or systems working in isolation because they require different expertise, resources, and information. Energy systems are, therefore, commonly composed of interacting components and subsystems, owned by different actors and stakeholders with their individual objectives, constraints, and policies. Due to historical, political, physical, or social constraints, it may be impractical or undesirable to permanently synthesise the necessary components into a single entity, thus cooperation among the independent components is required. This implies decentralised decision making and control. Multi-agent systems (MAS) provide an efficient approach to model and design such systems, enabling the loosely coupled components represented by agents to interact to achieve their individual goals as well as system's objectives [1-4].

Dynamics and uncertainty present an integral challenge. The variability of the systems and the effects of the decisions made by different agents make it impractical, if not impossible, to determine complete strategies at design time. Often, it is necessary to adopt some sort of learning capability to improve the strategy at run time. Systems need to adapt themselves (e.g., their decision making) as they observe and learn about the environment. Reinforcement learning allows agents to explore the unknown environment and to improve their behaviour through trial-and-error interaction with their dynamic environment. Reinforcement learning focuses on single-agent learning. Its extension in a multi-agent context considering multiple learning agents is known as multi-agent reinforcement learning (MARL) [5-8].

In principle, learning agents are designed to learn autonomously based on their reward function and learning algorithm. However, fully autonomous learning is often infeasible in the presence of highly complex problems or potentially dangerous outcome, in which case interactive guidance from humans is necessary. Human in the loop learning (HILL) addresses the challenge by involving humans in the learning process of the agents where humans can intervene and give advice [9-12].

The mission of this internship is to investigate the potentials of employing MARL and HILL in MAS to design self-adaptive energy systems that can learn to adapt themselves in response to dynamics. The proposed solutions will be applied in the management of local energy communities that are equipped with renewable energy sources (RES). More precisely, we aim at increasing RES penetration and electricity cost saving by optimising energy consumption to improve self-consumption in the communities and to support grid constraints (e.g., through demand response schemes), while ensuring user comfort. MARL and HILL will be employed to learn user preferences from data or observations and to consider dynamic changes in the system, which aids agents in their decision making. Various on-going projects in the lab on energy communities could provide further specifics on the use-cases and data for this internship.

Work environment

ENGIE Lab CRIGEN is the corporate research centre of the ENGIE Group dedicated to new gases, new uses of energy, digital solutions, and emerging technologies. Located in the Paris region at Stains (93), it has 200 employees. It provides tested, proven, and marketable industrial applications, as well as new offers based on the development and pooling of innovative ideas, scientific knowledge, and technical expertise. Its ability to innovate is a key advantage for the ENGIE Group. A part of CRIGEN, the CSAI Lab (Computer Science & Artificial Intelligence) carries out work in the fields of Computer Science and Artificial Intelligence (Machine Learning), Interoperability for Digital Ecosystem, and Distributed Autonomous System. The intern will be integrated into CSAI and work closely with scientists in machine learning and MAS.

Tasks and responsibilities

- Prepare a state of the art on MARL, HILL, and their applications in energy management
- Propose relevant strategies to integrate MARL and HILL into MAS to improve energy management (e.g., learn user preferences and learn to improve decision making)
- Design and implement a proof of concept to validate the proposed solutions
- Present and report findings orally and in writing
- Contribute to scientific papers

Education: Bac+5 (Master 2 or final year of engineering school) in Computer Science or related fields

Requirements

- Knowledge in machine learning and deep learning; experiences in reinforcement learning or HILL is a plus;
- Knowledge in MAS such as agent models and architectures, coordination, and agent-oriented programming
- Experiences in performing innovative research in applied machine learning and AI
- Fluency in English (spoken and written)
- Pro-active, independent, and comfortable working in teams
- Knowledge in optimisation and energy domain is appreciated.

Technical skills

- Advanced skills in Java and Python
- Software development (analysis, design, implementation, test) and Git
- Knowledge and skills in machine learning frameworks (scikit-learn, TensorFlow, or PyTorch)
- Experiences with agent platform (e.g., JADE, JaCaMo, …) is a plus.

Work conditions

- **Location**: ENGIE Lab CRIGEN, 4 Rue Joséphine Baker, 93240 Stains (Paris region, RER D, station Pierrefitte - Stains)
- **Duration**: up to 6 months
- **Start date**: as soon as possible

Application: Interested applicants may send their **CV, cover letter, and transcripts of the last 2 years** to the following people:

Sarra BEN ABBES: sarra.ben-abbes@external.engie.com
Rim HANTACH: rim.hantach@external.engie.com
Oudom KEM: oudom.kem@external.engie.com

Bibliography

[1] Stephen DJ McArthur, Euan M Davidson, Victoria M Catterson, Aris L Dimeas, Nikos D Hatziargyriou, Ferdinanda Ponci and Toshihisa Funabashi. Multi-agent systems for power engineering applications—Part I: Concepts, approaches, and technical challenges. IEEE Transactions on Power systems, vol. 22, no. 4, pages 1743–1752, 2007.

[2] Stephen DJ McArthur, Euan M Davidson, Victoria M Catterson, Aris L Dimeas, Nikos D Hatziargyriou, Ferdinanda Ponci and Toshihisa Funabashi. Multi-agent systems for power engineering applications—Part II: echnologies, standards, and tools for building multi-agent systems. IEEE Transactions on Power Systems, vol. 22, no. 4, pages 1753–1759, 2007.

[3] Vitor N Coelho, Miri Weiss Cohen, Igor M Coelho, Nian Liu and Frederico Gadelha Guimarães. Multi-agent systems applied for energy systems integration: State-of-the-art applications and trends in microgrids. Ap- plied energy, vol. 187, pages 820–832, 2017.

[4] Gregor Rohbogner, Ulf JJ Hahnel, Pascal Benoit and Simon Fey. Multi-agent systems' asset for smart grid applications. Computer Science and Information Systems, vol. 10, no. 4, pages 1799–1822, 2013.

[5] Lucian Buşoniu, Robert Babuška and Bart De Schutter. Multi-agent reinforcement learning: An overview. Innovations in multi-agent systems and applications-1, pages 183–221, 2010.

[6] Liviu Panait and Sean Luke. Cooperative multi-agent learning: The state of the art. Autonomous agents and multi-agent systems, vol. 11, no. 3, pages 387–434, 2005.

[7] Esmat Samadi, Ali Badri and Reza Ebrahimpour. Decentralized multi-agent based energy management of microgrid using reinforcement learning. International Journal of Electrical Power & Energy Systems, vol. 122, page 106211, 2020.

[8] Rui Yang and Lingfeng Wang. Development of multi-agent system for building energy and comfort management based on occupant behaviors. Energy and Buildings, vol. 56, pages 1–7, 2013.

[9] Huanghuang Liang, Lu Yang, Hong Cheng, Wenzhe Tu and Mengjie Xu. Human-in-the-loop reinforcement learning. In 2017 Chinese Automation Congress (CAC), pages 4511–4518. IEEE, 2017.

[10] David Abel, John Salvatier, Andreas Stuhlmüller and Owain Evans. Agent-Agnostic Human-in-the-Loop Reinforcement Learning. ArXiv, vol. abs/1701.04079, 2017.

[11] Ruohan Zhang, Faraz Torabi, L. Guan, Dana H. Ballard and Peter Stone. Leveraging Human Guidance for Deep Reinforcement Learning Tasks. In IJCAI, 2019.

[12] Yuchen Cui, Pallavi Koppol, Henny Admoni, Reid G. Simmons, Aaron Steinfeld and Tesca Fitzgerald. Understanding the Relationship between Interactions and Outcomes in Human-in-the-Loop Machine Learning. In IJCAI, 2021.