Ph.D. in Information Technology
Thesis Defense

July 8th, 2026
at 3:00 pm
Room Schiavoni – building 20A

Davide MARAN – XXXVIII Cycle

Theoretical Guarantees for Reinforcement Learning in Continuous State and Action Spaces

Supervisor: Prof. Marcello Restelli

Abstract:

Adapting reinforcement learning to Markov decision processes with continuous state and action spaces is crucial for many real-world applications, particularly in robotics and control. From a theoretical perspective, however, this regime remains poorly understood: while tight guarantees exist for tabular MDPs, a general theory with provably efficient algorithms for the continuous case is still lacking. In this thesis, we introduce a new class of problems, Smooth MDPs, which provides a unifying framework for reinforcement learning in continuous environments. Smooth MDPs are defined through functional regularity assumptions on the Bellman operator and on reward and transition dynamics. This framework rigorously generalizes several models studied in the literature, including continuous bandits, LQRs and their extensions, Lipschitz and kernelized MDPs. Although Sobolev-H\"older-type assumptions are central to nonparametric regression and have recently been explored in bandit problems, prior to this work, a systematic treatment in the context of reinforcement learning was still missing.
From a statistical perspective, we characterize the complexity of Smooth MDPs and demonstrate that, in the presence of a simulator, an algorithm with minimax-optimal sample complexity (with near-log factors) exists. In the online interaction scenario, we propose three algorithms for Smooth MDPs, each associated with a different learning objective, and analyze their performance by obtaining regret bounds that explicitly depend on the degree of process regularity. These results generalize classical bounds for simpler classes and improve the state of the art for kernelized MDPs with Mat\'ern kernels.

Davide SALAORNI – XXXVIII Cycle

Empowering Digital Twins for Battery Energy Storage Systems with Machine Learning: Design, Development and Applications of the Framework

Supervisor: Prof. Marcello Restelli

Abstract:

The global transition toward sustainable energy solutions heavily relies on the widespread deployment of energy storage systems. By mitigating the intermittency of renewable energy sources, batteries have become the cornerstone of modern, decentralized power grids. Essentially, batteries act as energy buffers, storing excess renewable generation and discharging to meet demand when generation falls short. However, optimizing their performance, safety, and longevity remains a critical challenge with profound implications for the entire energy ecosystem.

In this context, digital twins have emerged as a transformative paradigm. By bridging the physical and digital worlds, digital twins provide high-fidelity virtual replicas of real-world systems, enabling risk-free what-if simulations and extensive offline analysis. Furthermore, when coupled with recent advancements in machine learning, digital twins unlock unprecedented capabilities for real-time monitoring and autonomous control, leveraging a virtually unlimited supply of simulated data.

In this dissertation, we explore the intersection of digital twins and the energy storage domain, with a particular focus on integrating advanced machine learning techniques to enhance system monitoring and control. Specifically, we investigate the adoption of a digital twin for a battery storage system, spanning from its core design and implementation to its application across multiple replicas to simulate a renewable energy community. Within this framework, machine learning serves as a pivotal component applied in three distinct forms. First, we employ unsupervised learning and drift detection techniques to enable the digital twin to dynamically adapt to the evolving physical conditions of the battery. Second, we frame the energy management problem of a microgrid as a Markov Decision Process, adopting a reinforcement learning approach to optimize the battery's operational policy. Finally, we extend the analysis to a renewable energy community, proposing a novel multi-agent approach to drive the behavior of self-interested community members toward global cooperation and coordination.

The proposed methodologies are rigorously validated through an extensive experimental campaign using real-world datasets, facilitated by the collaboration with RSE S.p.A., a leading Italian research institution in the energy sector. Ultimately, this research aims to contribute to the advancement of the energy field by demonstrating the immense potential of integrating digital twins with learning systems to foster a more efficient, sustainable, and resilient energy future.

PhD Committee

Prof. Manuel Roveri, Politecnico di Milano

Prof.ssa Giorgia Ramponi, University of Zurich

Prof. Antoine Marot, RTE