

Course title: An Introduction to Learning and Data-Driven Methods

Duration [number of hours]: 24

PhD Program [MERC/MPS/SPACE]: **MERC**

Name and Contact details of unit organizer(s)

Prof. Mirco Musolesi, UCL and University of Bologna (m.musolesi@ucl.ac.uk, <https://www.mircomusolesi.org>)
with seminar lectures from:

Prof. Giovanni Russo,

University of Salerno (giovarusso@unisa.it, <https://sites.google.com/view/giovanni-russo>)

Course Description:

This module introduces students to learning and data-driven methods enabling agents to make decisions. The module starts with motivating the use of these methods across a wide range of application domains that are relevant for risk and complexity. Starting from the motivating examples, the key ingredients of learning and data-driven methods are then introduced. It is then shown how these key ingredients can be *mixed* to obtain classic algorithms enabling agents to make decisions from data. This leads students to familiarize with Dynamic Programming, Gradient Descent and their applications to Reinforcement Learning, focusing on policy-based and value-based resolution methods that require a discretize state/action spaces. Throughout this part of the module, strengths and weaknesses of the methods are discussed and the discussion will smoothly lead to motivate, and introduce, modern data-driven control, deep learning and multi-agent techniques. A key feature of the module is that methodological aspects are complemented with the development of case studies of interest for risk and complexity. By doing so, students will get hands-on-experience in coding their own algorithms, thus leading a computationally-oriented understanding of the methods introduced in the module. Finally, through an interactive discussion, we will outline the open challenges in the fascinating field of data-driven control and reinforcement learning.

Syllabus:

1. Introduction to learning and data-driven methods: motivation
2. Key ingredients: mathematical formulation, optimization and data
3. Value-based methods: Q-learning and SARSA as case studies
4. Policy-based methods: REINFORCE as case study
5. Data-driven methods for control
6. Deep Learning: Applications of Deep Learning to RL and DQN
7. Introduction to multi-agent RL

Assessment: Discussion of project works developed throughout the module

Suggested reading and online resources:

1. Moritz Hardt and Benjamin Recht. Patterns, Predictions, And Actions: A story about machine learning
2. Richard S. Sutton and Andrew G. Barto. Reinforcement Learning: An Introduction
3. Kevin Leyton-Brown and Yoav Shoham. Multiagent Systems, Game-theoretic and Logical Foundations. Cambridge University Press
4. OpenAI. Key papers in Deep RL. <https://spinningup.openai.com/en/latest/spinningup/keypapers.htm>