

Course title:

Introduction to Reinforcement Learning and Data-Driven Control for Complex Systems

Duration [number of hours]: **24**

PhD Program: 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 by Prof. Giovanni Russo, University of Salerno (giovarusso@unisa.it, <https://sites.google.com/view/giovanni-russo>)

Course Description [max 150 words]:

This module introduces the key concepts of data-driven control and Multi-Agent Reinforcement Learning (MARL). The first part of the module will be devoted to a short overview of Reinforcement Learning (RL) and DDC from the overarching framework of sequential decision making. We will then discuss the extensions of RL to MARL considering different classes of learning, including Online MARL, Co-evolutionary Learning, Swarm Intelligence and Adaptive Mechanism Design. Within the RL framework, the agent learns the task by interacting with an environment. The last part of the module will instead consider situations of practical interest where the task is known but the environment model is not. This leads to formalize the data-driven control problem and, for this problem, we will present a number of resolution methods that allow to solve data-driven control problems in a wide range of applications. Throughout the module, the methodological aspects will be complemented with a discussion on the corresponding computational methods and case studies will be used to illustrate the concepts. We will discuss recent research in MARL and DDC and we will outline the open challenges in this fascinating field.

Syllabus [itemized list of course topics]:

1. Introduction to sequential decision-making: mathematical formulation, key ingredients, historical perspective;
2. The dynamic programming algorithm, value functions, value-based algorithm: examples including Q-learning
3. Policy-based RL: examples, including REINFORCE (with -tutorial)-on REINFORCE
4. Deep Learning: Applications of Deep Learning to Reinforcement Learning
5. DQN and Policy-based methods based on RL
6. Introduction to multi-agent systems and multi-agent RL: Alternative learning paradigms for multi-agent systems
7. MARL algorithms: theoretical and implementation-oriented aspects.
8. Studying cooperation and competition in human, artificial and hybrid societies using ML/RL.
9. What happens when we know the task and the model? MPC Entropic-regularized data driven control and learning
10. fully probabilistic design of control policies

Assessment

Presentation and discussion of a research paper

Suggested reading and online resources:

1. Kevin Leyton-Brown and Yoav Shoham. Multiagent Systems, Game-theoretic and Logical Foundations. Cambridge University Press. 2009.

2. Karl Tuyls and Gerhard Weiss. Multiagent Learning: Basics, Challenges and Prospects. AI Magazine. Volume 33. Issue 3. 2012.
3. Karl Tulys and Peter Stone. Multiagent Learning Paradigms. In Francesco Belardinelli and Estefania Argente, editors. Multi-agent Systems and Agreement Technologies. LNAI Pages 3-21. Springer 2018.
4. OpenAI. Key papers in Deep RL. <https://spinningup.openai.com/en/latest/spinningup/keypapers.htm>
5. N. Matni, A. Proutiere, A. Rantzer, S. Tud, "From self-tuning regulators to reinforcement learning and back again" (available at: <https://arxiv.org/pdf/1906.11392.pdf>)
6. Richard S. Sutton and Andrew G. Barto. Reinforcement Learning: An Introduction.
7. D. Bertsekas, J. Tsitsiklis, "Neuro-dynamic programming" (<http://www.athenasc.com/ndpbook.html>)
8. D. Gagliardi, G. Russo, "On a probabilistic approach to synthesize control policies from example datasets" (<https://arxiv.org/abs/2005.11191>)
9. G. Russo, "On the crowdsourcing of behaviors for autonomous agents" (<https://arxiv.org/abs/2010.04608>)
10. S. Levine, A. Kumar, G. Tucker, J. Fu, "Offline Reinforcement Learning: tutorial, review and perspectives on open problems" (<https://arxiv.org/abs/2005.01643>)
11. E. Garrabè, G. Russo, "Probabilistic design of optimal sequential decision-making algorithms in learning and control" (<https://arxiv.org/abs/2201.05212>)