Francesco De Lellis

Reinforcement Learning for Control

Tutor: prof. Mario di Bernardo
co-Tutor: prof. Giovanni Russo (University of Salerno)

Cycle: XXXV

Year: 2023

Background information

- MSc degree in Control Engineering on 03/10/2019
- Member of the research group on “Sincronizzazione e Controllo di Reti e Processi” (SINCRO)
- PhD student period: 01/11/2019 – 31/01/2023
- Scholarship: University of Naples Federico II ITEE grant
- Ongoing collaborations:
  - Prof. Giovanni Russo from University of Salerno
  - Prof. Mirco Musolesi from University College London
- Visitor at the University College London (UCL) for the following periods: 26/05/22- 17/06/22 and 27/06/22-04/09/2022

Summary of study activities [1]

- I attended courses, seminars and conferences to keep developing my knowledge of RL and control approaches

- Conferences / events attended:
  - Learning for Dynamics and Control (L4DC), Stanford University
  - UK Mobile, Wearable and Ubiquitous Systems Research Symposium (MOBI UK), University College London
  - Controlling Complexity: From Nonlinear Systems To Complex Networks And Beyond, “Museo Diocesano” of Ischia

Summary of study activities [2]

- Courses :
  - Introduction to Reinforcement Learning and Data-Driven Control for Complex Systems
  - Corso di formazione specifica alla salute e sicurezza sui luoghi di lavoro

- Key seminars:
  - UCL Neuro AI Talk Series
  - UCL Dark Seminar Series
  - Causal foundations for safe AGI by Tom Everitt (Deep Mind)
  - Multi-Agent Reinforcement Learning toward Zero-Shot by Kalesha Bullard (DeepMind)
  - SINCRO research group seminar series

Research areas [1]

- Control Theory:
  - Given a mathematical description of a dynamical systems
  - Control tries to modify the behaviour of such system so that it evolves in a desired way
  - This process passes through the definition of a control law

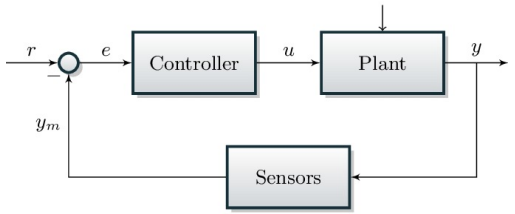

- Reinforcement Learning:
  - It provides a set of algorithms to learn (sub-) optimal solutions
  - It can be used as a controller without requiring a mathematical model of the system dynamics
  - These algorithms tend to be data inefficient and require long trainig sessions

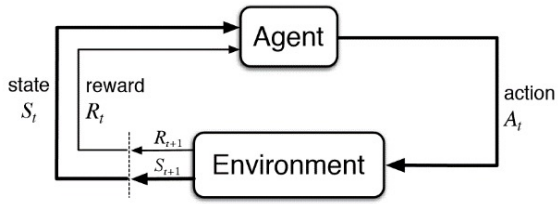

Research areas [2]

- My area of research lies in between Reinforcement Learning and Control

Control Theory

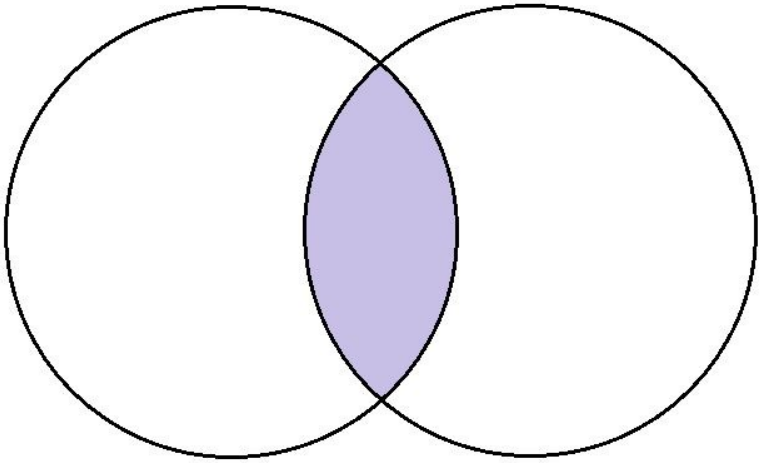

Reinforcement Learning

- In which I formulated reinforcement learning and control in a unified framework …
- … and merged Reinforcement Learning with control approaches in closed loop by:
  - Assisting the learning process with control laws
  - Use control theoretic arguments to provide stability certificate to RL algorithms

Research results [1]

- An optimal control formalism as been adopted to define both Reinforcement Learning and Control problems

- Derivation of tutoring mechanism for reinforcement learning agents

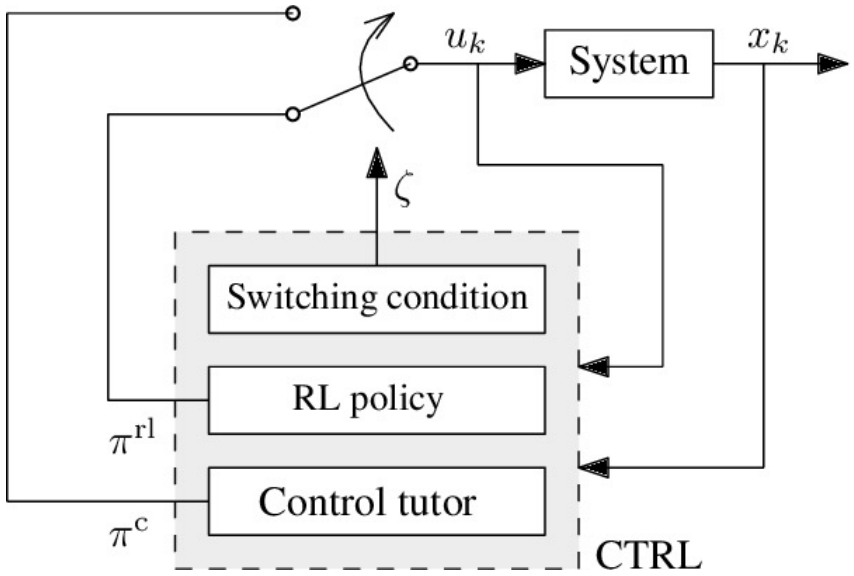

Research results [2]

- Implementation and testing of such algorithms for the case of Q-learning

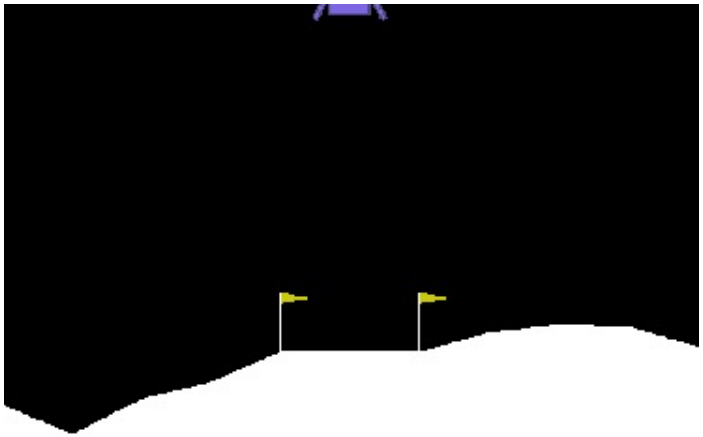

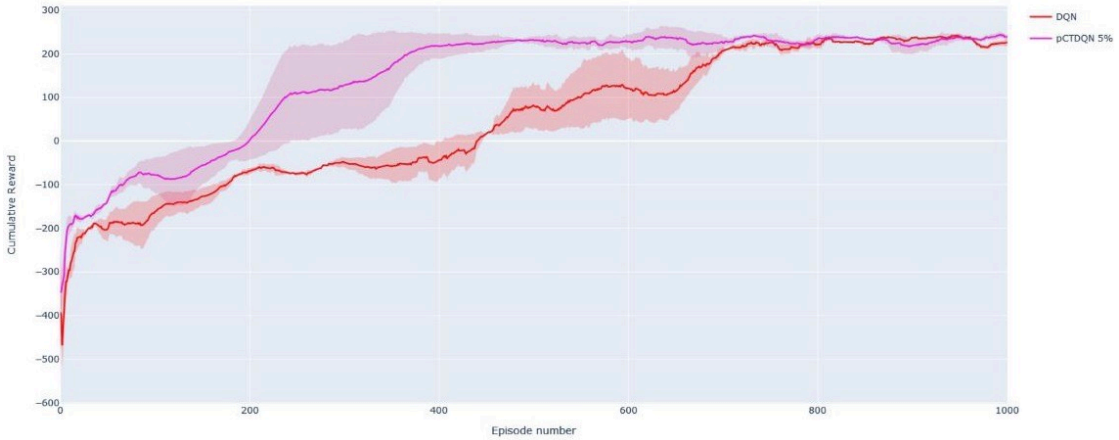

- Theoretical analysis of final control performance of Q-learning agents

Research products (3rd year)

[P1]	**F. De Lellis**, M. Coraggio, G. Russo, M. Musolesi, "CT-DQN: Control-Tutored Deep Reinforcement Learning", submitted to Learning for Dynamics and Control Conference, arXiv:2212.01343 ,2022.
[P2]	**F. De Lellis**, M. Coraggio, G. Russo, M. Musolesi, & M. di Bernardo, "*Control-Tutored Reinforcement Learning: Towards the Integration of Data-Driven and Model-Based Control*". In Learning for Dynamics and Control Conference, pp. 1048-1059, Proceedings of Machine Learning Research (PLMR), 2022.
[P3]	S. M. Brancato, **F. De Lellis**, D. Salzano, G. Russo, M. di Bernardo, "*External control of genetic toggle switch via Reinforcement Learning*", submitted to IEEE European Control Conference, arXiv:2204.04972 ,2022

PhD thesis overview

- Problem statement
  - Can Reinforcement Learning benefit from Control Theory? If so, how?
- Objective
  - Use optimal control to unify Reinforcement Learning and Control Theory
  - Use a control tutor to make RL agents learn faster
  - Provide stability certificate to final learned solution using Control Theory

PhD thesis overview

- Methodology
  - Problem statement

$$\begin{aligned} \max_{\pi} \ & \mathbb{E}[J^{\pi}], \\ \text{s.t.} \ & X_{k+1} = f(X_k, U_k, W_k), \quad k \in \{0, \dots, N-1\}, \\ & U_k = \pi(X_k), \quad k \in \{0, \dots, N-1\}, \\ & x_0 \text{ given}, \end{aligned}$$

  - Definition of evaluation metrics
  - Perform of numerical campaigns on a set of benchmark problems

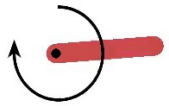

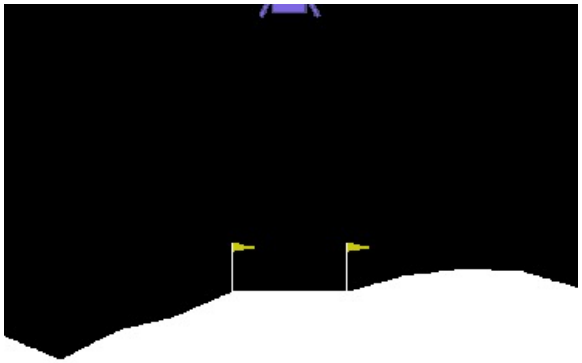

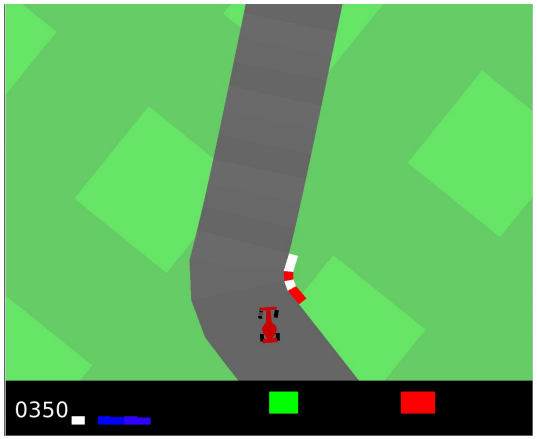

PhD thesis

- I will be showing how the use of control laws can be beneficial in the Reinforcement Learning loop
- The control tutored algorithms proposed are analyzed and tested in depth to clearly show the ups and downs of such approach
- People from both communities can benefit from these result to deliver better learning agents to solve hard control problems

Research products

[P1]	**F. De Lellis**, M. Coraggio, G. Russo, M. Musolesi, "CT-DQN: Control-Tutored Deep Reinforcement Learning", submitted to Learning for Dynamics and Control Conference, arXiv:2212.01343 ,2022.
[P2]	**F. De Lellis**, M. Coraggio, G. Russo, M. Musolesi, & M. di Bernardo, "*Control-Tutored Reinforcement Learning: Towards the Integration of Data-Driven and Model-Based Control*". In Learning for Dynamics and Control Conference, pp. 1048-1059, Proceedings of Machine Learning Research (PLMR), 2022.
[P3]	**F. De Lellis**, G. Russo, and M. di Bernardo. *"Tutoring Reinforcement Learning via Feedback Control",* proc. of IEEE European Control Conference (ECC), 2021.
[P4]	**F. De Lellis**, F. Auletta, G. Russo, P. De Lellis and M. di Bernardo. *"An Application of Control-Tutored Reinforcement Learning to the Herding Problem",* proc. of IEEE International Workshop on Cellular Nanoscale Networks and their Applications (CNNA), 2021.
[P5]	**F. De Lellis**, F. Auletta, G. Russo, P. De Lellis, M. di Bernardo, "*Control-Tutored Reinforcement Learning*", arXiv:2012.06863, 2019.