MERC PhD Project Proposal 2023/2024

Title of the research project:

Towards Multi-Agent Learning: bridging the gap between RL, control and complex systems

Keywords (up to five)

Data-driven control & optimization, complex systems, multi-agent reinforcement learning

Supervisors (at least two from two different areas):

Supervisor 1: *Mirco Musolesi (m.musolesi@ucl.ac.uk)*
Dept. of Computer Science at University College London, MERC PhD Board Member
www.mircomusolesi.org
Human-centred Artificial Intelligence, Reinforcement Learning, Modelling of social systems (human, artificial, hybrid)

Supervisor 2: *Giovanni Russo*
Dept. of Information and Electrical Engineering & Applied Mathematics (DIEM) at University of Salerno
www.sites.google.com/view/giovanni-russo
Reinforcement Learning, Data-Driven Control, Control Theory, Complex Systems, Optimization

Project description (max 5000 characters)

The vision and the context of the project. This project is concerned with learning and decision-making in complex, multi-agent, systems. While data-driven approaches have achieved several ground-breaking results in simulated/controlled environments, it is also apparent that new breakthrough is needed to efficiently deploy these algorithms into real-world, potentially safety-critical, applications. In this context, the project will investigate the development of new adaptation and learning mechanisms/models inspired for complex systems. We envision that, by bridging the gap between AI/RL, data-driven control and complex systems, the project will lead to the design of novel algorithms enabling autonomous agents to efficiently learn and compute actions directly from the observed data in a coordinated way. We will investigate machine learning models of increasing complexity, including foundation models. The results will be relevant to the design of networks of agents interacting with humans, the analysis/design of large-scale biochemical networks and financial/economic multi-agent systems.

State of the art. It is undeniable that model-based approaches to decision-making have been successful in many frontier applications. The design of small synthetic biological circuits and the design of platoons of autonomous cars are just few success stories of this approach. Unfortunately, reliable first-principles models in the form of e.g. ordinary difference/differential are not always available and, even when available, model parameters might be hard to identify. The limits of model-based approaches to decision-making become even more apparent in a multi-agent setting, spanning from finance/economics to social networks and human-robot collaboration, where autonomous agents coexist, and indeed interact, with humans. Hence, a popular paradigm to design autonomous systems is to feed them with data, using Reinforcement Learning (RL) to learn policies with recent successes of this data-driven approach include Deepmind's results in playing Go and beyond. There is, however, no free lunch and a crucial aspect is that results are achieved by failing: this, combined with the poor adaptability/robustness issues, is often unacceptable in real-world applications. More recently, there has also been an increasing interest in Multi-agent Learning and, in particular, in the area of cooperation and coordination of multi-agent systems. This is a very open area, with many basic challenges (such as how to deal with non-stationarity and high variance in the system under consideration). This project also aims at making substantial contributions in that area.

Key research questions. With this project we seek to overcome some of the current limitations of the above approaches by investigating an approach combining in a novel way AI/RL, neural networks and complex systems science. Ultimately, we will investigate how collaborative mechanisms can be designed to enable collaboration within a complex network of these agents making decisions from local observations through e.g., neural networks.

Project objectives. The detailed research program will be shaped based on the interests of students. A preliminary list of research objectives is given below:

O1 – Physics-aware decisions. Mechanisms will be designed, which embed physics in various ways in the learning and decision-making process. We will explore the use of e.g., state-space and foundation models within the control/learning process, investigating if these models can be used to learn efficiently and to guarantee safety. We will also explore how these models can be integrated into the actuation pipeline via end-to-end neural networks.

O2 – Learning in multi-agent settings. This objective investigates how decision-makers from O1 can collaborate and learn how to perform a collaborative task. The focus will be on Multi-agent Reinforcement Learning frameworks and possibility to apply foundation models to also tackle this class of collaborative problems.

O3 – From theory to the real world (and back). The specific application, which will involve collaborative multi-agent tasks, will be picked jointly with interested students. We expect an iterative process between theory and applications, where application aspects inform the development of the tools.

Workplan. The project will be developed in incremental tasks and periodic meetings will be scheduled with the supervisors. During the first part of the PhD, the student will start by reading and summarising the existing literature in the area. Afterwards, the student will start considering some small-scale problems, which are representative of specific classes of applications. In the final part of their PhD, the student will apply the theoretical framework developed during the PhD to a selection of real-world applications, such as those in the area of cooperative AI.

See the list of references for further details on the different aspects of the project.

Relevance to the MERC PhD Program (max 2000 characters)

This is a multi-disciplinary project that involves the use of approaches from a broad spectrum of disciplines such as physics modelling, machine learning, control and complexity science. We expect that the potential application of the results will be relevant to wide range of applications across the MERC program. In fact, we will ultimately design agents that, by interacting across a complex network, will be able to perform cooperative tasks that each agent alone could not successfully complete. Through tasks execution the agents must guarantee safety (both for themselves, the other agents and the environment) and this key requirement will be embedded into the application of choice that will be used to illustrate the effectiveness of the results.

See the list of references for further details on the different aspects of the project.

Key references

Stefano V. Albrecht and Peter Stone. Autonomous Agents Modelling Other Agents: A Comprehensive Survey and Open Problems. Artificial Intelligence. Volume 258. 2018.

Lucian Busoniu, Robert Babuska, Bart De Schutter. A Comprehensive Survey of Multiagent Reinforcement Learning. In IEEE Transactions on Systems, Man and Cybernetics. Volume 38. Issue 2. March 2008.

Karl Tuyls and Gerhard Weiss. Multiagent Learning: Basics, Challenges and Prospects. AI Magazine. Volume 33. Issue 3. 2012.

Karl Tuylis and Peter Stone. Multiagent Learning Paradigms. In Francesco Belardinelli and Estefania Argente, editors, Multi-agent Systems and Agreement Technologies. Lecture Notes in Artificial Intelligence. Pages 3-21. Springer 2018.

A. Dafoe et al. Open Problems in Cooperative AI. arXiv:2012.08630.

N. Anastassacos, S. Hailes and M. Musolesi. Partner Selection for the Emergence of Cooperation in Multi-agent Systems using Reinforcement Learning. Proceedings of AAAI'20. 2020.

E. Hughes et al. Inequity aversion improves cooperation in intertemporal social dilemmas. In Proceedings of NeurIPS'18, 2018.

Rupam Mahmood, D. Rorenkevych, G. Vasan, W. Ma and J. Bergstra. Benchmarking Reinforcement Learning Algoritms on Real-world Robots. Proceedings of CoRL 2018. 2018.

Henderson et al. Deep Reinforcement Learning that Matters. In Proceedings of AAAI 2018. 2018.

E. Theodorou, J. Buchli, S. Schall. A generalized Path Integral Control Approach to Reinforcement Learning. In Journal of Machine Learning Research. Volume 11. November 2010.

M. Karny. Towards fully probabilistic control design. In Automatica. Volume 32. Numer 12. December 1996

M. Sznaier. Control oriented learning in the era of big data. IEEEControl Systems Letters. Volume 5. Number 6. December 2020

R. M. Murray, K. J. Astrom, S. P. Boyd, R. W. Brockett, G. Stein. Future directions in control in an information-rich world. IEEE Control Systems Magazine. Volume 23. March 2003.

E. Garrabé, G. Russo, Probabilistic design of optimal sequential decision-making algorithms in learning and control, Annual Reviews in Control, vol. 54, pp. 81-102, 2022.

Joint supervision arrangements

Meetings will be scheduled on an as-needed basis, in order to ensure the effective development of the project. As a minimum, supervisor(s) will meet students at least once a week.

Location and length of the study period abroad (min 12 months)

The candidate will be able to spend a research period (or research periods) at Prof. Musolesi's lab at University College London. Other research visits at leading labs working at the forefront of the topics covered by the project can be arranged.

Any other useful information

The project can be tailored towards both applied and methodological aspects. Potential candidates are encouraged to contact their potential supervisors for more details.

Please return this form via email by no later than 9th February 2024 to merc@ssmeridionale.it