

Model Checking Agents that Learn

Dr Francesco Belardinelli

Studentship: Untaxed bursary of 16,777 per annum (2018/19 figure including London weighting plus home/EU fees)

The Department of Computing is a leading department of Computer Science among UK Universities, and has consistently been awarded the highest research rating. In the 2014 REF assessment, The Department was ranked third (1st in the Research Intensity table published by The Times Higher), and was rated as "Excellent" in the previous national assessment of teaching quality. Applications are invited for a PhD student in *Model Checking Agents that Learn* under the supervision of Dr Belardinelli.

In Reinforcement Learning (RL) autonomous agents have to choose their actions in order to maximise some notion of cumulative reward [1]. Tools and techniques for RL have been applied successfully to domain as diverse as game theory, control theory, multi-agent systems, and swarm intelligence. However, in several cases we do not see our agents simply as reward-maximisers: while learning to perform a certain task (say, exploring an unknown space) we might want our software agents never to enter a "deadlock" state, from which they are not able to get out. Similarly, message-passing agents need always to be responsive, i.e., able to receive and send messages.

These multi-agent scenarios calls for specification languages expressive enough to capture the behaviour of learning agents, as well as model checking methods to tackle the verification of properties such as those above. However, most of the research on formal methods for strategic reasoning has focused on mature (non-learning) agents.

Your contribution: you will first extend Markov decision processes, the formalisms of choice for RL, to multi-agent systems by considering the simultaneous execution of several learning agents. Then, you will develop an interpretation of Alternating-time Temporal Logic [2], one of the most popular logic-based languages to specify the strategic behaviour of agents, on these multi-agent Markov decision processes. In particular, the satisfaction of the agents' goals will be evaluated only w.r.t. executions that maximise the agents' reward. In this context you will analyse the complexity of the corresponding model checking problem. We anticipate that in selected cases the verification task will be decidable. For such cases an implementation might be envisaged as part of the project.

- [1] Sutton, Richard S.; Barto, Andrew G. (1998). Reinforcement Learning: An Introduction. MIT Press. ISBN 0-262-19398-1.
- [2] Alur, R.; Henzinger, T.; and Kupferman, O. 2002. Alternating-time temporal logic. Journal of the ACM 49(5):672713.

To apply for this position, you will need to

- have a strong background in at least one of the following areas: Logics, Formal Methods, and Reinforcement Learning.
- be capable of programming in the most popular programming languages, including C, C++, Java.
- show commitment, team working and a critical mind.

Previous experience in formal verification by model checking and/or logics for multi-agent systems is preferred, but not essential.

Applicants are expected to have a First Class or Distinction Masters level degree, or equivalent, in a relevant scientific or technical discipline, such as computer science or mathematics. Applicants must be fluent in spoken and written English.

The position is fully funded, covering tuition fees, and a stipend/bursary. The position is available to home, and EU students.

How to apply

To apply for this position, please follow the application guidelines at:
<http://www.imperial.ac.uk/computing/prospective-students/courses/phd/phd-application-guidelines/>

In the application form, please write *Model Checking Agents that Learn* in the “Proposed Research Topic” field, and *Dr Francesco Belardinelli* in the Proposed Research Supervisor field.

Early applications are encouraged. Informal inquiries about this position are also encouraged and can be directed to Dr Francesco Belardinelli. For further information see <https://www.doc.ic.ac.uk/~fbelard/>

This position will be based at the South Kensington campus in central London. Applicants are advised to visit <http://www.imperial.ac.uk/computing/prospective-students/courses/phd/> for general information on becoming a PhD student in the Department of Computing.

We are committed to equality and valuing diversity. We are also an Athena SWAN Silver Award winner, a Stonewall Diversity Champion, a Two Ticks Employer, and are working in partnership with GIREs to promote respect for trans people.