

Knowledge-based reinforcement learning and knowledge evolution

PhD position / Sujet de thèse

Cultural knowledge evolution and multiagent reinforcement learning share some of their prominent features. Putting explicit knowledge at the heart of the reinforcement process may contribute to better explanation and transfer.

Cultural knowledge evolution deals with the evolution of knowledge representation in a group of agents. For that purpose, cooperating agents adapt their knowledge to the situations they are exposed to and the feedback they receive from others. This framework has been considered in the context of evolving natural languages [Steels, 2012]. We have applied it to ontology alignment repair, i.e. the improvement of incorrect alignments [Euzenat, 2017] and ontology evolution [Bourahla *et al.*, 2021]. We have shown that it converges towards successful communication through improving the intrinsic knowledge quality.

Reinforcement learning is a learning mechanism adapting the decision making process for maximising the reward provided by the environment to the actions performed by agents [Sutton and Barto, 1998]. Many multi-agent versions of reinforcement learning have also been proposed depending on the agent attitude (cooperative, competitive) and the task structure (homogeneous, heterogeneous) [Bučoniū *et al.*, 2010].

From an external perspective, the two approaches operate in a similar manner: agents perceive their environment, perform an action, receive reward or punishment, adapt their behaviour in consequence. However, a look into the inner mechanisms reveals important differences: the emphasis on knowledge quality instead of reward maximisation, the lack of probabilistic or even gradual interpretation, and even the absence of explicit choice in action or adaptation. Hence these two knowledge acquisition techniques are close enough to suggest replacing one by the other and different enough to cross-fertilise.

This thesis position aims at further exploring the commonalities and differences between experimental cultural knowledge evolution and reinforcement learning. In particular, its purpose is to study which features of one technique may be fruitful in the context of the other and which may not.

For that purpose, one research direction is the introduction of knowledge-based reinforcement learning. In knowledge-based reinforcement learning, the decision-making process (the choice of the action to be performed) is obtained through accumulated explicit knowledge. Thus the adaptation performed after reward or punishment will have to directly affect this knowledge. This has the advantage that it allows to explain the decisions made by agents. It will also allow for explicit knowledge exchange among them [Leno da Silva *et al.*, 2018].

This promotes a less utilitarian view of knowledge in which the evaluation of the performance of the system has to be disconnected from reward maximisation but to depend on the quality of the acquired knowledge. Of course, these two aspects need to remain related (the acquired knowledge must be relevant to the environment). This separation between knowledge and reward is useful when agents have to change environment or use their knowledge to perform various tasks.

Another use of reinforcement mechanisms relevant to cultural knowledge evolution is related to the motivation for agents to explore unknown knowledge territories [Colas *et al.*, 2019]. By associating an intrinsic reward to the newly acquired knowledge, agents are able to improve the coverage of their knowledge in a way not guided by the environment. Complementing cultural knowledge evolution with exploration motivation, should make agents more active in their understanding of the environment and knowledge acquisition.

These problems may be treated both theoretically and experimentally

This work is part of an ambitious program towards what we call cultural knowledge evolution partly funded by the [MIAI Knowledge communication and evolution chair](#).

References:

- [Bourahla *et al.*, 2021] Yasser Bourahla, Manuel Atencia, Jérôme Euzenat, Knowledge improvement and diversity under interaction-driven adaptation of learned ontologies, Proc. 20th AAMAS, London (UK), pp242-250, 2021 <https://moex.inria.fr/files/papers/bourahla2021a.pdf>
- [Bučoniū *et al.*, 2010] Lucian Bučoniū, Robert Babuška, Bart De Schutter, Multi-agent reinforcement learning: an overview, Chapter 7 of D. Srinivasan and L.C. Jain, eds., Innovations in Multi-Agent Systems and Applications – 1, Springer, Berlin (DE), pp183–221, 2010 http://www.dcsc.tudelft.nl/~bdeschutter/pub/rep/10_003.pdf
- [Colas *et al.*, 2019] Cédric Colas, Pierre-Yves Oudeyer, Olivier Sigaud, Pierre Fournier, Mohamed Chetouani, Curious: Intrinsically motivated modular multi-goal reinforcement learning, Proc. 36th ICML, Long Beach (CA US), pp1331–1340, 2019 <http://proceedings.mlr.press/v97/colas19a/colas19a.pdf>
- [Euzenat, 2017] Jérôme Euzenat, Communication-driven ontology alignment repair and expansion, Proc. 26th IJCAI, Melbourne (AU), pp185-191, 2017 <https://moex.inria.fr/files/papers/euzenat2017a.pdf>
- [Leno da Silva *et al.*, 2018] Felipe Leno Da Silva, Matthew Taylor, Anna Helena Real Costa, Autonomously reusing knowledge in multiagent reinforcement learning, Proc. 27th IJCAI, pp5487-5493, 2018 <https://www.ijcai.org/proceedings/2018/0774.pdf>
- [Steels, 2012] Luc Steels (ed.), *Experiments in cultural language evolution*, John Benjamins, Amsterdam (NL), 2012

[Sutton and Barto, 1998] Richard Sutton, Andrew Barto, Reinforcement learning: an introduction, The MIT Press, Cambridge (MA US), 1998 (2nd ed. 2018) <http://incompleteideas.net/book/RLbook2020.pdf>

Links:

- This page: <https://moex.inria.fr/training/2021-Th-cultural-reinforcement.html>
 - MIAI Knowledge communication and evolution: <https://moex.inria.fr/cooperation/miai/>
 - mOeX web site: <https://moex.inria.fr>
 - *Lazy lavender*: <https://gitlab.inria.fr/moex/lazylav>
-

Qualification: Master or equivalent in computer science.

Researched skills:

- Curiosity and openness.
- Interaction with other researchers.
- Autonomous researcher.
- Taste for experimentation.
- Knowledge of multi-agent simulation and/or reinforcement learning not required but a plus.
- Innovative.

Doctoral school: [MSTII](#), Université Grenoble Alpes.

Advisor: [Jérôme Euzenat](#) (Jerome:Euzenat#inria.fr) and [Jérôme David](#) (Jerome:David#univ-grenoble-alpes.fr).

Group: The work will be carried out in the [mOeX](#) team common to [INRIA](#) & [LIG](#). mOeX is dedicated to study knowledge evolution through adaptation. It gathers researchers which have taken an active part these past 15 years in the development of the semantic web and more specifically ontology matching and data interlinking.

Place of work: The position is located at [INRIA Grenoble Rhône-Alpes](#), Montbonnot a main computer science research lab, in a stimulating research environment.

Hiring date: October 2023.

Duration: 36 months

Salary: 1982€/month (gross, before social contributions and taxes).

Deadline: as soon as possible.

Contact: For further information, contact us.

Procedure: Contact us (do not wait) and apply to the [INRIA campaign](#) (<https://jobs.inria.fr/public/classic/fr/offres/2023-05854>). See also [here](#).

File: Provide Vitæ, motivation letter and references. It is very good if you can provide a Master report and we will ask for your marks in Master, so if you have them, you can join them.