Internship project, master level, 2026

Evolving large populations of adaptive neural agents in ecologically plausible environments

Supervision: Clément Moulin-Frier. (Note: I am proposing several Master internships and it will be possible to adapt the topics to the scientific interests and the technical skills of the candidates, as well as to the duration of the internship. Integrating ideas from the different projects I propose is also an option).

Research group: BioTiC team, Inria Center of Lyon, France ; in collaboration with the Flowers team, Inria Center of the University of Bordeaux.

Duration: 4 to 6 months (with flexibility).

Level: Master research internship (Licence level could be considered as well).

Keywords: Artificial Intelligence, Artificial Life, Multi-agent systems, meta reinforcement learning, recurrent neural networks, simulation environments, eco-evolutionary dynamics, numerical simulation on GPU.

How to apply: contact clement.moulin-frier@inria.fr ; with a CV and letter of motivation. We also recommend sending documents or reports describing previous projects you have been working on (even if they are not directly related to the topic), as well as your grades and links to code repositories.

Requirements: We are looking for highly motivated MSc students (Master II). Programming skills and prior experience with Python and deep learning frameworks (e.g. Pytorch, JAX) and/or Web technologies with GPU-acceleration (WebGL, WebGPU) are expected (or at least a strong motivation to learn it).

Scientific context

There are striking differences in how adaptation operates in biological versus artificial systems. For instance, Reinforcement Learning (RL) optimizes an action policy in order to maximize reward provided by the environment. In contrast, the natural world does not contain any explicit reward whatsoever: reward systems have instead evolved physiologically in biological organisms. Moreover, agent's training in RL relies on an episodic paradigm where the environment is reset to initial conditions at the start of each new episode. In contrast, natural environments are never reset: they are instead continuously modified by their inhabitants over many generations.

Over the last years, we have developed several models where a large number of agents interact in a large shared environment with complex dynamics and limited resources. In contrast with most work in multi-agent AI, we never reset neither the environment nor the agent population. We study how agents can evolve neural architectures and complex behaviors in such persistent environments, inspired by recent proposals combining evolution and niche construction theories (Laland et al., 2015).

Project

In recent papers (Hamon et al., 2023; Taylor-Davies et al., 2025), we proposed to study the eco-evolutionary dynamics of non-episodic neuroevolution in large multi-agent environments, based on the above principles.

The objective of the proposed internship is to extend these experiments along several possible directions of research (with flexibility depending on the student's background and interest). These include:

- The use of indirect encoding of macro properties of the agents' neural networks using recent methods from neuroevolution (Barandiaran & Stovold, 2025; e.g. Najarro et al., 2023; Zuo et al., 2023), meta reinforcement learning (e.g. Duan et al., 2016) or transformer-based architectures (Lange et al., 2024).
- The implementation of a more complex ecologically plausible simulated environment with realistic physics, compositional dynamics (where resources can be combined together to create new ones, in the spirit of Minecraft, see also Bornemann et al., 2023 from our team), spatio-temporal variability (e.g. seasonal cycles) and multiple co-evolving species (e.g. prey-predator dynamics).
- Thorough qualitative and quantitative evaluation of the system dynamics, both at the macro level (characterizing eco-evolutionary dynamics in the system) and of the agent's evolved behavior, in particular in their ability to adapt during their lifetime (e.g. implementing a battery of test lab environments characterizing their generalization abilities in unknown situations).

References

Note: We do not ask candidates to read all these references. Going through the paper's abstracts can be useful, potentially selecting 2-3 papers to look at in more detail. The references we consider as the most important are highlighted in bold.

Barandiaran, M., & Stovold, J. (2025). Growing Reservoirs with Developmental Graph Cellular Automata. *ALife 2025 Conference*. https://doi.org/10.48550/arXiv.2508.08091

Bornemann, R., Hamon, G., Nisioti, E., & Moulin-Frier, C. (2023). *Emergence of Collective Open-Ended Exploration from Decentralized Meta-Reinforcement Learning*. NeurIPS 2023 - Conference on Neural Information Processing Systems / ALOE Workshop. https://doi.org/10.48550/arXiv.2311.00651

Duan, Y., Schulman, J., Chen, X., Bartlett, P. L., Sutskever, I., & Abbeel, P. (2016). RL^2: Fast

Reinforcement Learning via Slow Reinforcement Learning. *arXiv:1611.02779 [Cs, Stat]*. http://arxiv.org/abs/1611.02779

Hamon, G., Nisioti, E., & Moulin-Frier, C. (2023). **Eco-evolutionary Dynamics of Non-episodic Neuroevolution in Large Multi-agent Environments.** *2022 Genetic and Evolutionary Computation Conference (GECCO 2022)*. https://doi.org/10.48550/arXiv.2302.09334

Laland, K. N., Uller, T., Feldman, M. W., Sterelny, K., Müller, G. B., Moczek, A., Jablonka, E., & Odling-Smee, J. (2015). **The extended evolutionary synthesis: Its structure, assumptions and predictions**. *Proceedings of the Royal Society B: Biological Sciences*, *282*(1813), 20151019. https:**//doi.org/10.1098/rspb.2015.1019**

Lange, R. T., Tian, Y., & Tang, Y. (2024). *Evolution Transformer: In-Context Evolutionary Optimization* (No. arXiv:2403.02985). arXiv. https://doi.org/10.48550/arXiv.2403.02985

Najarro, E., Sudhakaran, S., & Risi, S. (2023, July 24). ***Towards Self-Assembling Artificial Neural Networks through Neural Developmental Programs.*** ALIFE 2023: Ghost in the Machine: Proceedings of the 2023 Artificial Life Conference. https://doi.org/10.1162/isal_a_00697

Taylor-Davies, M., Hamon, G., Boulet, T., & Moulin-Frier, C. (2025). Emergent Kin Selection of Altruistic Feeding via Non-episodic Neuroevolution. In P. García-Sánchez, E. Hart, & S. L. Thomson (Eds.), *Applications of Evolutionary Computation* (pp. 496–509). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-90062-4_31

Zuo, W., Pedersen, J., & Risi, S. (2023). Evolution of an Internal Reward Function for Reinforcement Learning. *Proceedings of the Companion Conference on Genetic and Evolutionary Computation*, 351–354. https://doi.org/10.1145/3583133.3590610