MAP Spring 2027: A Deep Dive into Independent Reinforcement Learners

ELBICA Lab, Professor Eliott

MAP in a nutshell. You will help Professor Eliott revisit the literature on Independent Reinforcement Learners (IL). More importantly, your research will focus on identifying computational mechanisms that may modify independent learners' responses to the same external interaction history.

- **What is an IL?** Take a look at this paper to learn more about it.

Activities

1. On a weekly basis, you will read technical papers and books, analyze, and present/explain to Prof. Eliott and your peers what you learned.
    a. The use of GenAI tools will be very limited, mostly not allowed.
2. There is no coding involved. This MAP focuses heavily on theory at the intersection of Reinforcement Learning (RL) and Multi-Agent Systems (MAS). Game Theory will also be touched upon on an *ad hoc* basis.
    a. Note: this research applies classic RL techniques. Thus, this is NOT a MAP about agentic AI (we are NOT using GenAI systems).

Requisites

- Preferably: successfully completed CSC207 by the end of 2026.
    - You may still apply if you have successfully completed CSC161.
- Be excited about:
    a) Reading technically heavy CS/math content;
    b) Making connections across disciplines (the ELBICA lab has strong roots in neuroscience and philosophy).
    c) Learning more about MAS and connecting it to robotics.
    d) Working as part of a team.

Research Problems

Four students will start from the same reading ground. However, each will pursue a different research question, as below:

1. How does an agent's inner architecture influence learning dynamics under interdependence?
2. What happens to IL when agents use dynamically evolving inner mechanisms to modulate their learning?



3. Under what theoretical assumptions does independent learning become difficult to characterize or guarantee?
4. What is the current state of the literature regarding challenges in IL?

Outcomes – At the end of the MAP, you will have:

1. Acquired a substantial understanding of IL.
2. Compiled your findings in professional formats to share with a technical audience.
3. Developed communication skills to discuss and present technically heavy CS content.
4. Contributed to the ELBICA lab's vision of dissecting and computationally modeling meaningful decision-making.

Expectations – if you apply to this MAP, you agree that:

- The use of GenAI systems will be very limited: this MAP requires ***your*** cognitive effort.
- You will read technically heavy materials and make sense of them. Thus, this MAP sits on the theory side and involves NO coding.

How to Apply?

- Info at the end of this document.

Readings

Students will investigate these papers/books: (others will be added as needed)

1. Coordination of independent learners in cooperative Markov games. Laëtitia Matignon , Guillaume J. Laurent, Nadine Le Fort-Piat
2. Shayak, N. and Eliott, F. Cooperation as Black Box: Conceptual Fluctuation and Diagnostic Tools for Misalignment in MAS, 2025.
3. Situated Learners in a Sequential Decision-Making Setting. Kasimov, T. and Takei, S. and Lu, H. and Lee, M. and Eliott, F. The Fourteenth International Conference on Intelligent Systems and Applications INTELLI 2025.
4. Networked Independent Reinforcement Learners Playing an Evolutionary Game. Xu, Z. and Chen, J. and Eliott, F. In: Quintián, H., et al. Hybrid Artificial Intelligent Systems. HAIS 2024. Lecture Notes in Computer Science, vol 14858. Springer, Cham.
5. Fernanda Eliott and Carlos Henrique Ribeiro. Emergence of cooperation through simulation of moral behavior. In *Hybrid Artificial Intelligent Systems (HAIS 2015). Procs. of the 10th I. Conf. on Hybrid Artificial Intelligence Systems.*
6. Fernanda Eliott and Carlos Henrique Ribeiro. Moral behavior and empathy modeling through the premise of reciprocity. In *Procs. of the 1st International Conference on Human and Social Analytics*. St. Julians, Malta: Huso, 2015.
7. Yu, X., Morri, R. and Eliott, F. EDA, An Empathy-Driven Computational Architecture. Workshop on Goal Reasoning, ACS 2021.
8. Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations. Yoav Shoham (Author), Kevin Leyton-Brown.
9. Multi-Agent Reinforcement Learning: Foundations and Modern Approaches https://www.marl-book.com/
10. C. Boutilier (1999). Sequential optimality and coordination in multiagent systems. In IJCAI'99: Proceedings of the 16th international joint conference on Artificial intelligence.



11. D. Weyns, G. Vizzari, and T. Holvoet (2006) Environments for situated multi-agent systems: beyond infrastructure. In Environments for Multi-Agent Systems II: Second International Workshop, E4MAS 2005.
12. N. Fulda and D. Ventura (2007) Predicting and preventing coordination problems in cooperative q-learning systems. In IJCAI, Vol. 2007, pp. 780–785.
13. L. Sun, Y. Yang, Q. Duan, Y. Shi, C. Lyu, Y. Chang, C. Lin, and Y. Shen (2025) Multi-agent coordination across diverse applications: a survey.
14. Tan, M. (1993). Multi-Agent Reinforcement Learning: Independent Versus Cooperative Agents. Proceedings of the 10th International Conference on Machine Learning (ICML), 330–337. 10.1016/B978-1-55860-307-3.50049-6.
15. Panait, L., & Luke, S. (2005). Cooperative Multi-Agent Learning: The State of the Art. Autonomous Agents and Multi-Agent Systems, 11(3), 387–434. DOI: 10.1007/s10458-005-2631-2.
16. Panait, L., Sullivan, K., & Luke, S. (2006). Lenient Learners in Cooperative Multiagent Systems. Proceedings of AAMAS 2006, 801–803. DOI: 10.1145/1160633.1160776.
17. Matignon, L., Laurent, G. J., & Le Fort-Piat, N. (2007). Hysteretic Q-Learning: An Algorithm for Decentralized Reinforcement Learning in Cooperative Multi-Agent Teams. Proceedings of IEEE/RSJ IROS 2007, 64–69. DOI: 10.1109/IROS.2007.4399095.
18. Palmer, G., Tuyls, K., Bloembergen, D., & Savani, R. (2018). Lenient Multi-Agent Deep Reinforcement Learning. Proceedings of AAMAS 2018, 443–451.
19. Schroeder de Witt, C., Gupta, T., Makoviichuk, D., Makoviychuk, V., Torr, P. H. S., Sun, M., & Whiteson, S. (2020). Is Independent Learning All You Need in the StarCraft Multi-Agent Challenge? arXiv:2011.09533.
20. Papoudakis, G., Christianos, F., Schäfer, L., & Albrecht, S. V. (2021). Benchmarking Multi-Agent Deep Reinforcement Learning Algorithms in Cooperative Tasks. NeurIPS Datasets and Benchmarks Track.
21. Ding, D., Wei, C.-Y., Zhang, K., & Jovanović, M. R. (2022). Independent Policy Gradient for Large-Scale Markov Potential Games: Sharper Rates, Function Approximation, and Game-Agnostic Convergence. Proceedings of ICML 2022, PMLR 162, 5166–5220.
22. Cheng, M., Zhou, R., Kumar, P. R., & Tian, C. (2024). Provable Policy Gradient Methods for Average-Reward Markov Potential Games. Proceedings of AISTATS 2024, PMLR 238, 4699–4707.
23. Jin, R., Chen, Z., Lin, Y., Song, J., & Wierman, A. (2025). Approximate Global Convergence of Independent Learning in Multi-Agent Systems. Proceedings of AISTATS 2025, PMLR 258, 2818–2826.

How to Apply

Answer and email the questions below to Professor Eliott (deadlines will be advertised *via* the CS student list)

1. Academic honesty: you must write the answers yourself – help from GenAI systems is ***not allowed***. Now, answer: Did you get assistance from any AI system to answer the questions in this document?

2. Why are you applying for this MAP, and what are your research expectations?

3. Four research problems are presented. Which one do you prefer, and why?

4. Visit the lab page (https://elbica.cs.grinnell.edu/start-here/) and answer: Why do you consider yourself a good fit for the ELBICA lab?

5. Do you have the preferred prerequisites? If not, that is fine; feel free to discuss other relevant experience.



6. What is the connection between your major/minor/concentrations and previous education with this MAP?

7. What makes your application a good fit for this MAP? What would you like to share to make your application stand out?

8. Passion for doing research is a vital element in the ELBICA lab. What are your thoughts on that?

9. ELBICA lab's projects involve a lot of open-ended tasks: many times, you will be treated as a PhD student. What is your reaction to that? (For context, read the section: ELBICA students reflected on their experiences and shared their advice with future lab members.)

10. How do you approach/navigate stressful academic situations?

11. The ELBICA lab emphasizes team collaboration. For you, what is the importance of team spirit?

12. What does "being coachable" mean to you, and why is it important?

13. Have you ever directly interacted with Professor Eliott? Explain.

14. If invited to this MAP, you MUST reply as soon as possible to help me plan and extend the opportunity to another student (if you reject the invite). Here is the question: Do you understand that:

    a) You MUST acknowledge that you received the invite as soon as possible?

    b) If I do not hear from you, I will extend the opportunity to another student.

