

Cooperative Multi-Agent Learning in a Complex World: Challenges and Solutions

Yali Du

Department of Informatics
King's College London
yali.du@kcl.ac.uk

Over the past few years, artificial intelligence (AI) has achieved great success in a variety of applications, such as image classification and recommendation systems. This success has often been achieved by training machine learning models on static datasets, where inputs and desired outputs are provided. However, we are now seeing a shift in this paradigm. Instead of learning from static datasets, machine learning models are increasingly being trained through feedback from their interactions with the world. This is particularly important when machine learning models are deployed in the real world, as their decisions can often have an impact on other agents, turning the decision-making process into a multi-agent problem. As a result, multi-agent learning in complex environments is a critical area of research for the next generation of AI, particularly in the context of cooperative tasks. Cooperative multi-agent learning is an essential problem for practitioners to consider as it has the potential to enable a wide range of multi-agent tasks.

In this presentation, we will review the background and challenges of cooperative multi-agent learning. We will survey our research that aims to address challenges in the following areas:

- Scalability and cooperation in multi-agent learning, including studies on (Han et al. 2019; Du et al. 2019; Xu et al. 2020; Du et al. 2021a; Fang et al. 2019; Mguni et al. 2021).
- Evaluating AI agents, including studies on (Yan et al. 2022; Du et al. 2021b).
- Human-in-the-loop interactive learning for robust and trusted decision making, including studies on (Kazantzidis et al. 2022; Gu et al. 2022; Liu et al. 2022).
- Various applications of cooperative multi-agent learning, including (Ruan et al. 2022; Xu et al. 2020; Wang et al. 2021; Chen et al. 2021).

We will also discuss future directions for research in this field.

References

Chen, X.; Du, Y.; Xia, L.; and Wang, J. 2021. Reinforcement Recommendation with User Multi-aspect Preference. In Leskovec, J.; Grobelnik, M.; Najork, M.; Tang, J.; and Zia, L., eds., *WWW*, 425–435. ACM / IW3C2.

Du, Y.; Han, L.; Fang, M.; Dai, T.; Liu, J.; and Tao, D. 2019. LIIR: learning individual intrinsic reward in multi-agent reinforcement learning. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, 4403–4414.

Du, Y.; Liu, B.; Moens, V.; Liu, Z.; Ren, Z.; Wang, J.; Chen, X.; and Zhang, H. 2021a. Learning Correlated Communication Topology in Multi-Agent Reinforcement learning. In *AAMAS*, 456–464.

Du, Y.; Yan, X.; Chen, X.; Wang, J.; and Zhang, H. 2021b. Estimating alpha -Rank from A Few Entries with Low Rank Matrix Completion. In *ICML*, 2870–2879. PMLR.

Fang, M.; Zhou, T.; Du, Y.; Han, L.; and Zhang, Z. 2019. Curriculum-guided Hindsight Experience Replay. In Wallach, H. M.; Larochelle, H.; Beygelzimer, A.; d’Alché-Buc, F.; Fox, E. B.; and Garnett, R., eds., *NeurIPS*, 12602–12613.

Gu, S.; Yang, L.; Du, Y.; Chen, G.; Walter, F.; Wang, J.; Yang, Y.; and Knoll, A. 2022. A Review of Safe Reinforcement Learning: Methods, Theory and Applications. *arXiv preprint arXiv:2205.10330*.

Han, L.; Sun, P.; Du, Y.; Xiong, J.; Wang, Q.; Sun, X.; Liu, H.; and Zhang, T. 2019. Grid-wise control for multi-agent reinforcement learning in video game ai. In *ICML*, 2576–2585. PMLR.

Kazantzidis, I.; Du, Y.; Freeman, C.; and Norman, T. 2022. How to train your agent: Active Learning from Human Preferences and Justifications in Safety-critical Environments. In *AAMAS*.

Liu, R.; Bai, F.; Du, Y.; and Yang, Y. 2022. Meta-Reward-Net: Implicitly Differentiable Reward Learning for Preference-based Reinforcement Learning. In *NeurIPS*.

Mguni, D.; Wu, Y.; Du, Y.; Yang, Y.; Wang, Z.; Li, M.; Wen, Y.; Jennings, J.; and Wang, J. 2021. Learning in Nonzero-Sum Stochastic Games with Potentials. In *ICML*, 7688–7699.

Ruan, J.; Du, Y.; Xiong, X.; Xing, D.; Li, X.; Meng, L.; Zhang, H.; Wang, J.; and Xu, B. 2022. GCS: Graph-Based Coordination Strategy for Multi-Agent Reinforcement Learning. In *AAMAS*.

Wang, X.; Du, Y.; Zhu, S.; Ke, L.; Chen, Z.; Hao, J.; and Wang, J. 2021. Ordering-Based Causal Discovery with Reinforcement Learning. In *IJCAI*, 3566–3573.

Xu, Y.; Fang, M.; Chen, L.; Du, Y.; Zhou, J. T.; and Zhang, C. 2020. Deep Reinforcement Learning with Stacked Hierarchical Attention for Text-based Games. In *NeurIPS*.

Yan, X.; Du, Y.; Ru, B.; Wang, J.; Zhang, H.; and Chen, X. 2022. Learning to Identify Top Elo Ratings: A Dueling Bandits Approach. In *AAAI*, 6375–6383. AAAI Press.