

Multiagent Learning: From Fundamentals to Foundation Models

Keynote Talk

Karl Tuyls

DeepMind

London, UK

karltuyls@deepmind.com

ABSTRACT

Research in multiagent learning has come a long way over the past few decades, from learning in abstract normal-form games such as Rock-Paper-Scissors, to learning in complex worlds such as Humanoid Soccer, Capture the Flag, Gran Turismo racing, and recently board games such as Diplomacy and Stratego. In this talk I will take you on a journey that starts in the mid 90's and sheds light on algorithmic progress over the years in multiagent learning systems, uncovering game-theoretic fundamentals for reinforcement learning, adaptability, and decision-making. There have been two major research eras in the field thus far, the pre-deep multiagent learning and deep multiagent learning periods. I believe we are now at the verge of a third period, multiagent learning with foundation models. We will connect old and new ideas of the first two periods, and lay out interesting challenges ahead of us for the coming era. Specifically, we consider the ways in which the cornerstone ideas of the first two periods may inform the development of generally capable multi-agent foundation models in the future.

Keywords

Multiagent learning; game theory; replicator dynamics; foundation model

ACM Reference format:

Karl Tuyls. 2023. Multiagent Learning: From Fundamentals to Foundation Models. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 2 pages.

BIOGRAPHY

Karl Tuyls (FBCS) is a research director at DeepMind where he leads the Game Theory & MultiAgent Team. He is also an honorary professor of Computer Science at the University of Liverpool, UK, and a Guest Professor at the University of Leuven, Belgium. Previously, he held academic positions at the Vrije Universiteit Brussel, Hasselt University, Eindhoven University of Technology, and Maastricht University. Prof. Tuyls has received several awards with his research, amongst which: the Information Technology prize 2000 in Belgium, best demo award at AAMAS'12, winner of various Robocup@Work competitions ('13, '14), and he was a co-author of the runner-up best paper award at ICML'18. He co-invented DeepNash, the first AI agent to reach human expert-level performance in the imperfect information game Stratego. Furthermore, his research has received substantial attention from international press and media, most recently his work on football analytics and Graph Imputer featured in Wired UK and Nature. He is a fellow of the British Computer Society (BCS), is on the editorial board of the Journal of Autonomous Agents and Multi-Agent Systems, and is (co)-editor-in-chief of the Springer briefs series on Intelligent Systems. Prof. Tuyls is also an emeritus member of the board of directors of the International Foundation for Autonomous Agents and Multiagent Systems.

ACKNOWLEDGEMENTS

The ideas presented in this talk are the result of joint work with many colleagues over the years, foremost at DeepMind, the University of Liverpool, Maastricht University, Vrije Universiteit Brussel and KU Leuven. Special thanks to Edward Hughes for being such a great sparring partner in developing this talk.