

# Automatic Curricula in Deep Multi-Agent Reinforcement Learning

Thore Graepel  
Google DeepMind  
thore@google.com

## ABSTRACT

Multi-agent systems are emerging as a crucial element in our pursuit of designing and building intelligent systems. In order to succeed in the real world artificial agents must be able to cooperate, communicate, and reason about other agents' beliefs, intentions and behaviours. Furthermore, as system designers we need to think about composing intelligent systems from intelligent subsystems, a multi-agent approach inspired by the observation that intelligent agents like organisations or governments are composed of other agents. Last but not least, as a product of evolution intelligence did not emerge in isolation, but as a group phenomenon. Hence, it seems plausible that learning agents require interaction with other agents to develop intelligence.

In his talk, he will discuss the exciting role that deep multi-agent reinforcement learning can play in the design and training of intelligent agents. In particular, training RL agents in interaction with each other can lead to the emergence of an automatic learning curriculum: From the perspective of each learning agent, the evolving behaviours of the other learning agents constitute a challenging environment dynamics and pose ever evolving tasks. He will present three case studies of deep multi-agent RL with auto-curricula: (i) Learning to play board games at master level with AlphaZero, (ii) Learning to play the game of Capture-The-Flag in 3d environments, and (iii) Learning to cooperate in social dilemmas.

## KEYWORDS

deep multi-agent reinforcement learning

### ACM Reference Format:

Thore Graepel. 2020. Automatic Curricula in Deep Multi-Agent Reinforcement Learning. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020)*, Auckland, New Zealand, May 9–13, 2020, IFAAMAS, 1 pages.

## BIOGRAPHY

Thore Graepel works as a research group lead at Google DeepMind and holds a part-time position as Chair of Machine Learning at University College London. In support of responsible innovation in artificial intelligence, Thore also serves as a Member of the Board of Directors at Partnership on AI. Thore studied physics at the University of Hamburg, Imperial College London, and Technical University of Berlin, where he also obtained his PhD in machine learning in 2001. After holding post-doctoral positions at ETH Zurich and Royal Holloway College, University of London, Thore joined Microsoft Research in Cambridge in 2003, where he co-founded the Online Services and Advertising group. Major applications of Thore's work include Xbox Live's TrueSkill system for ranking and matchmaking and the AdPredictor framework for click-through rate prediction in Bing. Furthermore, Thore's work on the predictability of private attributes from digital records of human behaviour has been the subject of intense discussion among privacy experts and the general public. At DeepMind, Thore has returned to his original passion of understanding and creating intelligent systems, and recently contributed to creating AlphaGo, the first computer program to defeat a human professional player in the full-sized game of Go.