

Verifying Strategic Abilities of Neural-symbolic Multi-agent Systems

M. Akintunde¹, E. Botoeva², P. Kouvaros³ and A. Lomuscio³

¹King's College London, UK

²University of Kent, UK

³Imperial College London, UK

Abstract

We investigate the problem of verifying the strategic properties of multi-agent systems equipped with machine learning-based perception units. We introduce a novel model of agents comprising both a perception system implemented via feed-forward neural networks and an action selection mechanism implemented via traditional control logic. We define the verification problem for these systems against a bounded fragment of alternating-time temporal logic. We translate the verification problem on bounded traces into the feasibility problem of mixed integer linear programs and show the soundness and completeness of the translation. We show that the lower bound of the verification problem is PSPACE and the upper bound is coNEXPTIME. We present a tool implementing the compilation and evaluate the experimental results obtained on a complex scenario of multiple aircraft operating a recently proposed prototype for air-traffic collision avoidance.

The full paper appeared in the Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning (KR20). Rhodes, Greece. IJCAI Press.

The full paper can be accessed at <https://proceedings.kr.org/2020/3/>

The work was partly funded by DARPA under the Assured Autonomy programme (FA8750-18-C-0095), the EPSRC Centre for Doctoral Training in High Performance Embedded and Distributed Systems (EP/L016796/1) and the Royal Academy of Engineering Chair in Emerging Technologies.

Keywords

Neuro-symbolic agents, Verification

