Title: **Safe Learning of Multi-Agent Action Models from Concurrent Joint Action Observations**

Authors:
Argaman Mordoch (Ben Gurion University)
Ori Karat (Ben Gurion University)
Lea Shmilovich (Ben Gurion University)
Yarin Benjamin (Ben Gurion University)
Brendan Juba (Washington University in St. Louis)
Roni Stern (Ben Gurion University)

Abstract: https://jair.org/index.php/jair/article/view/20494
Multi-Agent Planning (MAP) involves coordinating the actions of multiple autonomous agents to achieve shared objectives. A prevalent formalism for MAP is the Multi-Agent Planning Domain Definition Language (MA-PDDL). While effective, existing MA-PDDL solvers typically require complete access to agents' action models—specifically their preconditions and effects. However, manually creating these models is often intractable, requiring exhaustive domain expertise. This work explores an alternative approach: automatically learning agents' action models from observed transitions. Since learned models may be inaccurate, planning with them can yield invalid or non-executable sequences. To mitigate this, we formalize a requirement for *safety*, ensuring that plans generated via the learned model remain sound with respect to the real unknown action model.
Previous research introduced the Safe Action Model Learning (SAM) algorithm for single-agent domains. However, SAM is not suitable for MA-PDDL environments where observations include concurrently executed actions, since it cannot naturally disambiguate the individual contributions of the agents to the observed effects. To address this, we introduce Multi-Agent Safe Action Model Learning (MA-SAM), a safe action model learning algorithm designed to handle concurrent multi-agent observations. For scenarios where individual action effects remain ambiguous, we further propose MA-SAM+ , which learns the preconditions and effects of *macro-actions* representing concurrent execution of subsets of actions. We evaluate both algorithms on domains from the Competition of Distributed and Multi-Agent Planners (CoDMAP) benchmarks and a novel MAP domain inspired by the game *Overcooked*. We establish a theoretical lower bound on the sample complexity for learning safe action models in multi-agent settings. We prove that MA-SAM does not achieve this lower bound in all cases, identifying specific conditions under which its sample complexity may become unbounded. Empirically, both MA-SAM and MA-SAM+ significantly outperform SAM-based baselines in coverage and applicability rates. While their performance is comparable in many settings, MA-SAM+ highly outperforms MA-SAM in some of the evaluated domains. In conclusion, we present the first algorithms capable of learning safe MA-PDDL action models from concurrently executed actions, providing both theoretical foundations and empirical validation across diverse planning benchmarks.