Scalable Decision-Theoretic Planning in Open and Typed Multiagent Systems

Adam Eck¹, Maulik Shah², Prashant Doshi², and Leen-Kiat Soh³ ¹Oberlin College, ²University of Georgia, ³University of Nebraska

Problem

Open multiagent systems: agents join and leave the environment over time



Models only some peers and extrapolates their behaviors to the entire system Relies on statistical sampling theory similar to survey/polling research in the social sciences

Solution



Important for real-world problems

Wildfire suppression, autonomous ridesharing, cybersecurity

I-POMCP_o: principled *online MCTS planning algorithm* for many-agent environments for I-POMDP(–Lite) problem models

Scalable Decision-Theoretic Planning in Open and Typed Multiagent Systems

Adam Eck¹, Maulik Shah², Prashant Doshi², and Leen-Kiat Soh³ ¹Oberlin College, ²University of Georgia, ³University of Nebraska

Theoretical Results

Analytically establish bounds on:

- the error in an agent's estimated proportion of its neighbors that choose each action
- the error in the agent's estimated likelihoods of counts of how many neighbors will perform all actions
- the **regret** incurred from modeling only some neighbors



Experimental Results



I-POMCP₀ outperformed all baselines Due to better

coordination between agents by working together to put out fires in more locations

Modeling <u>fewer</u> neighbors led to <u>greater</u>

accumulated rewards Due to more sampled trajectories in the fixed time budget