

Roxana Rădulescu

Multi-objective learning agents

Most real-world problems involve multiple, potentially conflicting objectives; for example, safety versus fuel efficiency versus speed in autonomous driving, or treatment effectiveness versus side effects in medical treatment planning. Tackling such problems using reinforcement learning (RL) methods either requires an a-priori scalarisation of the reward signal, or involves applying multi-objective RL. In this talk I compare these approaches, and take a deeper dive into multi-objective RL. The goal is to highlight practical considerations, theoretical results, and additional challenges and benefits, as well as to delineate how and when it is appropriate to use multi-objective RL.