Recent Advances in Multi-Objective Multi-Agent Decision Making

Roxana Rădulescu[1], Patrick Mannion[2], Diederik M. Roijers[1,3], and Ann Nowé[1]

[1] Artificial Intelligence Lab, Vrije Universiteit Brussel, Belgium
[2] School of Computer Science, National University of Ireland Galway, Ireland
[3] Microsystems Technology, HU Univ. of Applied Sciences Utrecht, the Netherlands

Numerous real-world problems involve both multiple actors and objectives that should be taken into account when making a decision. Multi-objective multi-agent systems (MOMAS) represent an ideal setting to study such problems, but given the increasingly complex dimensions involved, it still remains an under-studied domain despite its high relevance. We present here a short overview of our recent advances in multi-objective multi-agent decision making settings.

MOMAS Taxonomy In MOMAS the reward signal for each agent is a vector, where each component represents the performance on a different objective. We consider that compromises between competing objectives should be made on the basis of the utility that these compromises have for the users. In other words, we assume there exists a utility function that maps the vector value of a compromise solution to a scalar utility.

In order to offer a unified view of the field, we build a taxonomy (Figure 1) of what constitutes a solution for a multi-objective multi-agent decision problem based on reward and utility functions. More details on each setting and solution concept can be found in [2].

Reward \ Utility	Team	Social choice	Individual
Team	Coverage sets	Mechanism design	Coverage sets (+ Negotiation) Equilibria & stability concepts
Individual		Mechanism design	Equilibria & stability concepts Coverage Sets as best responses

Fig. 1: Multi-objective multi-agent decision making taxonomy and mapping of solution concepts.

Another factor we identify is the difference between the optimisation criteria: expected scalarised returns (ESR) and scalarised expected returns (SER) [1].

This roughly distinguishes settings where either the utility of a single outcome (ESR) or the utility of the average outcome over multiple runs (SER) matters.

Learning in MONFGs We have studied multi-objective normal form games under the SER optimisation criterion with non-linear utility functions [3]. We show by example that while Nash equilibria (NE) need not exist, correlated equilibria (CE) can still be present when optimising with respect to a single given signal (i.e., single-signal CE).

Opponent modelling in MONFGs When the same multi-objective reward vector leads to different utilities for each user, it becomes essential for an agent to learn about the behaviour of other agents in the system. In [4] we present the first study of the effects of opponent modelling (OM) on MONFGs with non-linear utilities, under the SER criterion. We demonstrate that OM can alter the learning dynamics in this setting: when there are no NE, OM can have adverse effects on utility, or a neutral effect at best; when equilibria are present, OM can confer significant benefits (Figure 2).

	L	M	R
L	$(4,1)$	$(1,2)$	$(2,1)$
M	$(3,1)$	$(3,2)$	$(1,2)$
R	$(1,2)$	$(2,1)$	$(1,3)$

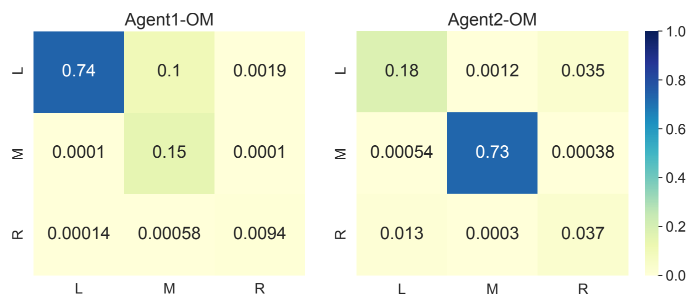

Fig. 2: Empirical outcome distributions when agents are using opponent modelling with utility functions $u_1(\mathbf{p}) = p^1 \cdot p^1 + p^2 \cdot p^2$ and $u_2(\mathbf{p}) = p^1 \cdot p^2$. Opponent modelling allows each agent to steer the outcome towards its preferred NE. Agent 1 obtains the highest SER under (L,L), while for Agent 2 that is (M, M).

References

1. Roijers, D.M., Steckelmacher, D., Nowé, A.: Multi-objective reinforcement learning for the expected utility of the return. In: Proceedings of the Adaptive and Learning Agents workshop at FAIM. vol. 2018 (2018)
2. Rădulescu, R., Mannion, P., Roijers, D.M., Nowé, A.: Multi-objective multi-agent decision making: a utility-based analysis and survey. Autonomous Agents and Multi-Agent Systems **34** (2020)
3. Rădulescu, R., Mannion, P., Zhang, Y., Roijers, D.M., Nowé, A.: A utility-based analysis of equilibria in multi-objective normal-form games. The Knowledge Engineering Review **35**, e32 (2020)
4. Zhang, Y., Rădulescu, R., Mannion, P., Roijers, D., Nowé, A.: Opponent modelling for reinforcement learning in multi-objective normal form games. In: Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems. pp. 2080–2082 (2020)