

The Research Group

## Artificial Intelligence Lab

has the honor to invite you to the public defense of the PhD thesis of

# Roxana-Teodora Rădulescu

to obtain the degree of Doctor of Sciences

Title of the PhD thesis:

**Decision Making in Multi-Objective Multi-Agent Systems - A Utility-Based Perspective**

Promotor:

**Prof. dr. Ann Nowé**

Co-promotors:

**Dr. Diederik M. Roijers**

**Dr. Patrick Mannion**

The defense will take place on

**Friday, September 3, 2021 at 14h00**

The defense can be followed through a live stream. Contact [Roxana.Radulescu@vub.be](mailto:Roxana.Radulescu@vub.be) for more information

### Members of the jury

Prof. dr. Coen De Roover (VUB, chair)

Prof. dr. Bart Bogaerts (VUB, secretary)

Prof. dr. Vincent Ginis (VUB)

Prof. dr. Peter Vamplew (Federation University Australia)

Prof. dr. Frans Oliehoek (Technische Universiteit Delft, the Netherlands)

Prof. dr. Katja Hofmann (Microsoft Research Lab, UK)

### Curriculum vitae

Roxana Rădulescu graduated as MSc in Computer Science with a specialisation in Artificial Intelligence at the VUB in 2015. She then joined the Artificial Intelligence Lab, VUB as a PhD student and teaching assistant.

Her research focuses on decision-making in settings involving multiple agents as well as multiple objectives. She has published 18 peer-reviewed articles in top-tier conferences and journals.

She has been an involved community member co-organising 3 times the Adaptive and Learning Agents workshop at AAMAS and being part in numerous program committees.

Her teaching tasks also allowed her to supervise 7 master and 10 bachelor students.

### Abstract of the PhD research

The prevalence of artificial agents in our world raises the need to ensure that they are able to handle the salient properties of the environment, in order to plan or learn how to solve specific tasks.

A first important aspect is the fact that real-world problems are not restricted to one agent, and often involve multiple agents acting in the same environment. Such settings have already proven to be challenging to solve, with a few examples including traffic systems, electricity grids, or warehouse management. Furthermore, the majority of these multi-agent system implementations aim to optimise the agents' behaviour with respect to a single objective, despite the fact that many problem domains inherently involve multiple objectives. By taking a multi-objective perspective on decision-making problems, complex trade-offs can be managed; e.g., supply chain management involves a complex coordination process for optimising the information and material flow between all the components of the supply chain, while minimising overall costs and complying with the conflicting demands of the involved partners.

In this work, we focus on these highlighted aspects and discuss how the process of decision-making and learning of artificial agents can be formalised and approached when there are multiple agents involved, and multiple objectives that need to be considered in the process. To analyse such problems, we adopt a utility-based perspective, and advocate that compromises between competing objectives should be made on the basis of the utility that these compromises have for the users, in other words, it should depend on the desirability of the outcomes.

As a first contribution, we develop a novel taxonomy to classify these settings. This allows us to offer a structured view of the field, to clearly delineate the current state-of-the-art in multi-objective multi-agent decision making approaches and to identify promising directions for future research.

As a second contribution, we proceed to analyse and investigate game theoretic equilibria under different multi-objective optimisation criteria and provide theoretical results concerning the existence and conditions for arriving to such solutions in these scenarios. We additionally show that it is possible for Nash equilibria to not exist in certain multi-objective multi-agent settings.

As a final contribution, we present the first study of the effects of opponent modelling on multi-objective multi-agent interactions. We contribute novel reinforcement learning algorithms for this setting, along with extensions that incorporate opponent behaviour reconstruction and learning with opponent learning awareness (i.e., learning while anticipating one's impact on the opponent's learning step).