

An Analytical Study of Utility Functions in Multi-Objective Reinforcement Learning



UNIVERSITAT DE
BARCELONA

Manel Rodríguez-Soto,¹ Juan Antonio Rodríguez-Aguilar,¹ and Maite López-Sánchez,²

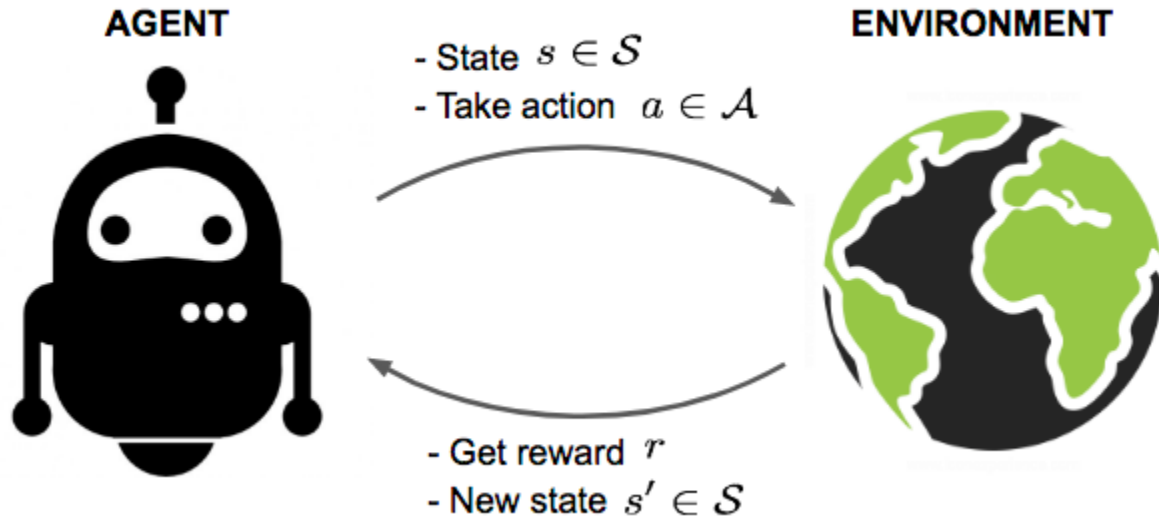
¹ Artificial Intelligence Research Institute (IIIA-CSIC). ² Universitat de Barcelona (UB).

Motivation

Multi-Objective reinforcement learning (MORL) is a *promising* research field.

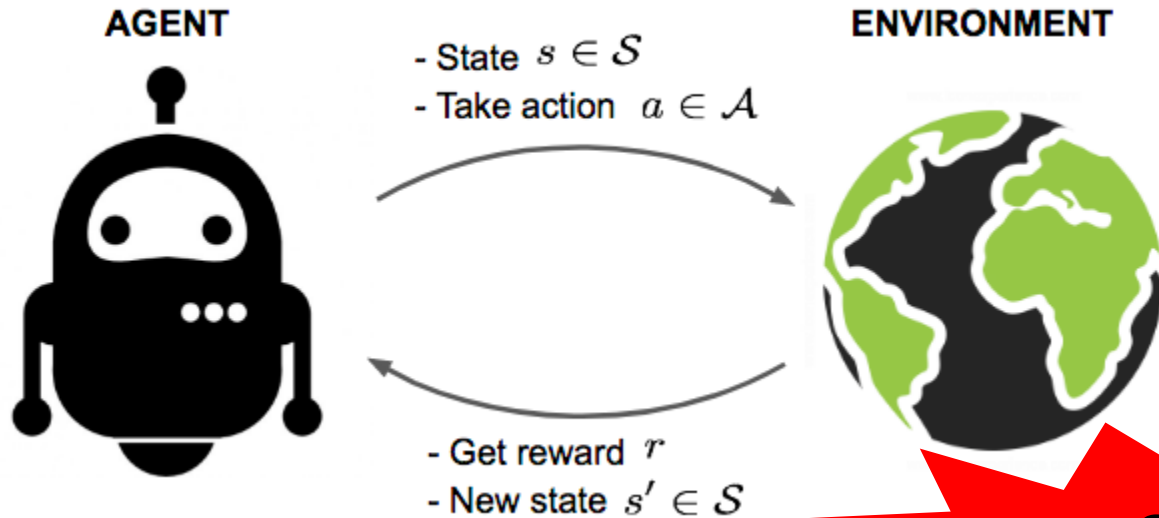
But there are still *crucial theoretical analyses missing...*

MORL background



**Multi-Objective
Markov Decision Process**

MORL background



Multi-Objective
Markov Decision Process

**REWARDS ARE
VECTORIAL**

MORL background



Policy


The agent's behaviour

MORL background



Policy

The agent's behaviour



Utility function


Compacts all objectives into one

MORL background



Policy

The agent's behaviour



Utility function


Compacts all objectives into one

MORL background



Policy

The agent's behaviour



Utility function

Compacts all objectives into one



Preference relation

Expresses the preferred outcomes

Research questions

- Which **utility functions** have **policies** optimising them?

Research questions

- Which **utility functions** have **policies** optimising them?

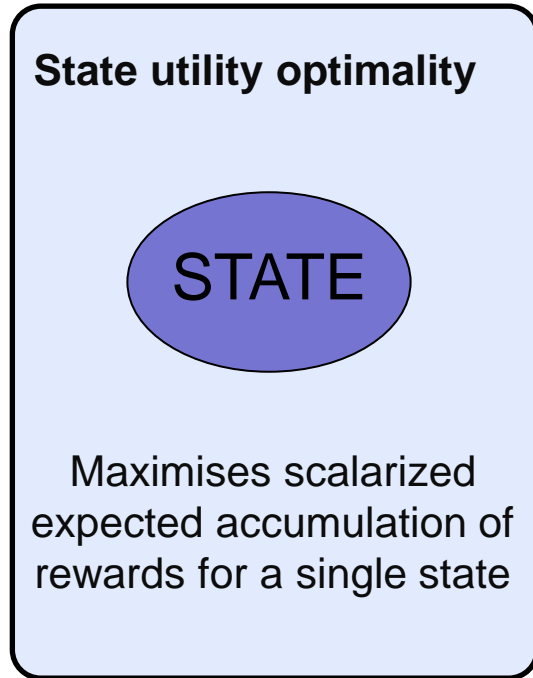
- Which **preferences** can be expressed as **utility functions**?

Research questions

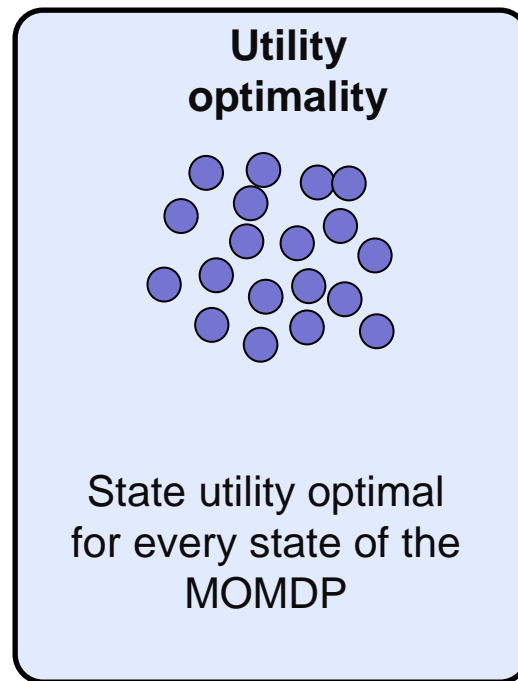
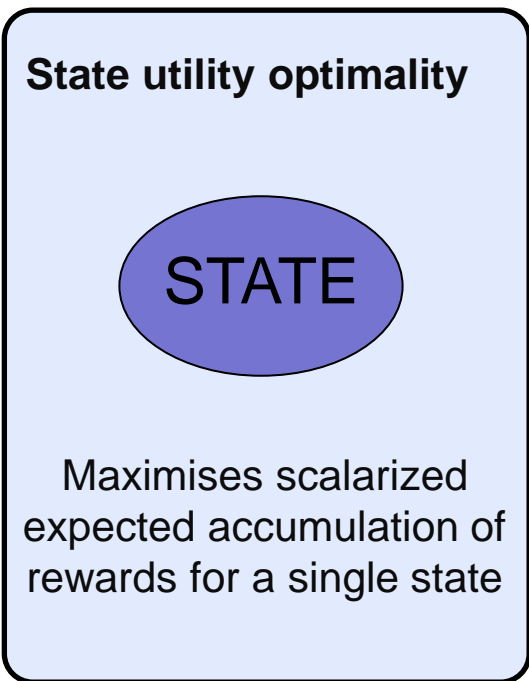
- Which **utility functions** have **policies** optimising them?

- Which **preferences** can be expressed as **utility functions**?

Our novel MORL definitions



Our novel MORL definitions



Our theoretical results

Stationary state utility optimal policies are guaranteed to exist for...

- **Continuous utility functions**

Our theoretical results

Stationary state utility optimal policies are guaranteed to exist for...

- **Continuous utility functions**

Stationary utility optimal policies are guaranteed to exist for...

- **Composition of affine utility functions and strictly monotonically increasing functions**

Our theoretical results

Stationary state utility optimal policies are guaranteed to exist for...

- **Continuous utility functions**

Stationary utility optimal policies are guaranteed to exist for...

- **Composition of affine utility functions and strictly monotonically increasing functions**

In conclusion...

- **It may happen that the policy your algorithm learns is only state utility optimal at the initial state, but not at the rest!!!**

Take-home message

When using MORL, be wary of whether your algorithm is really learning **utility optimal** policies. Otherwise, you may learn a **contradictory policy**!

Take-home message

When using MORL, be wary of whether your algorithm is really learning **utility optimal** policies. Otherwise, you may learn a **contradictory policy**!

Contact:

manel.rodriguez@iia.csic.es