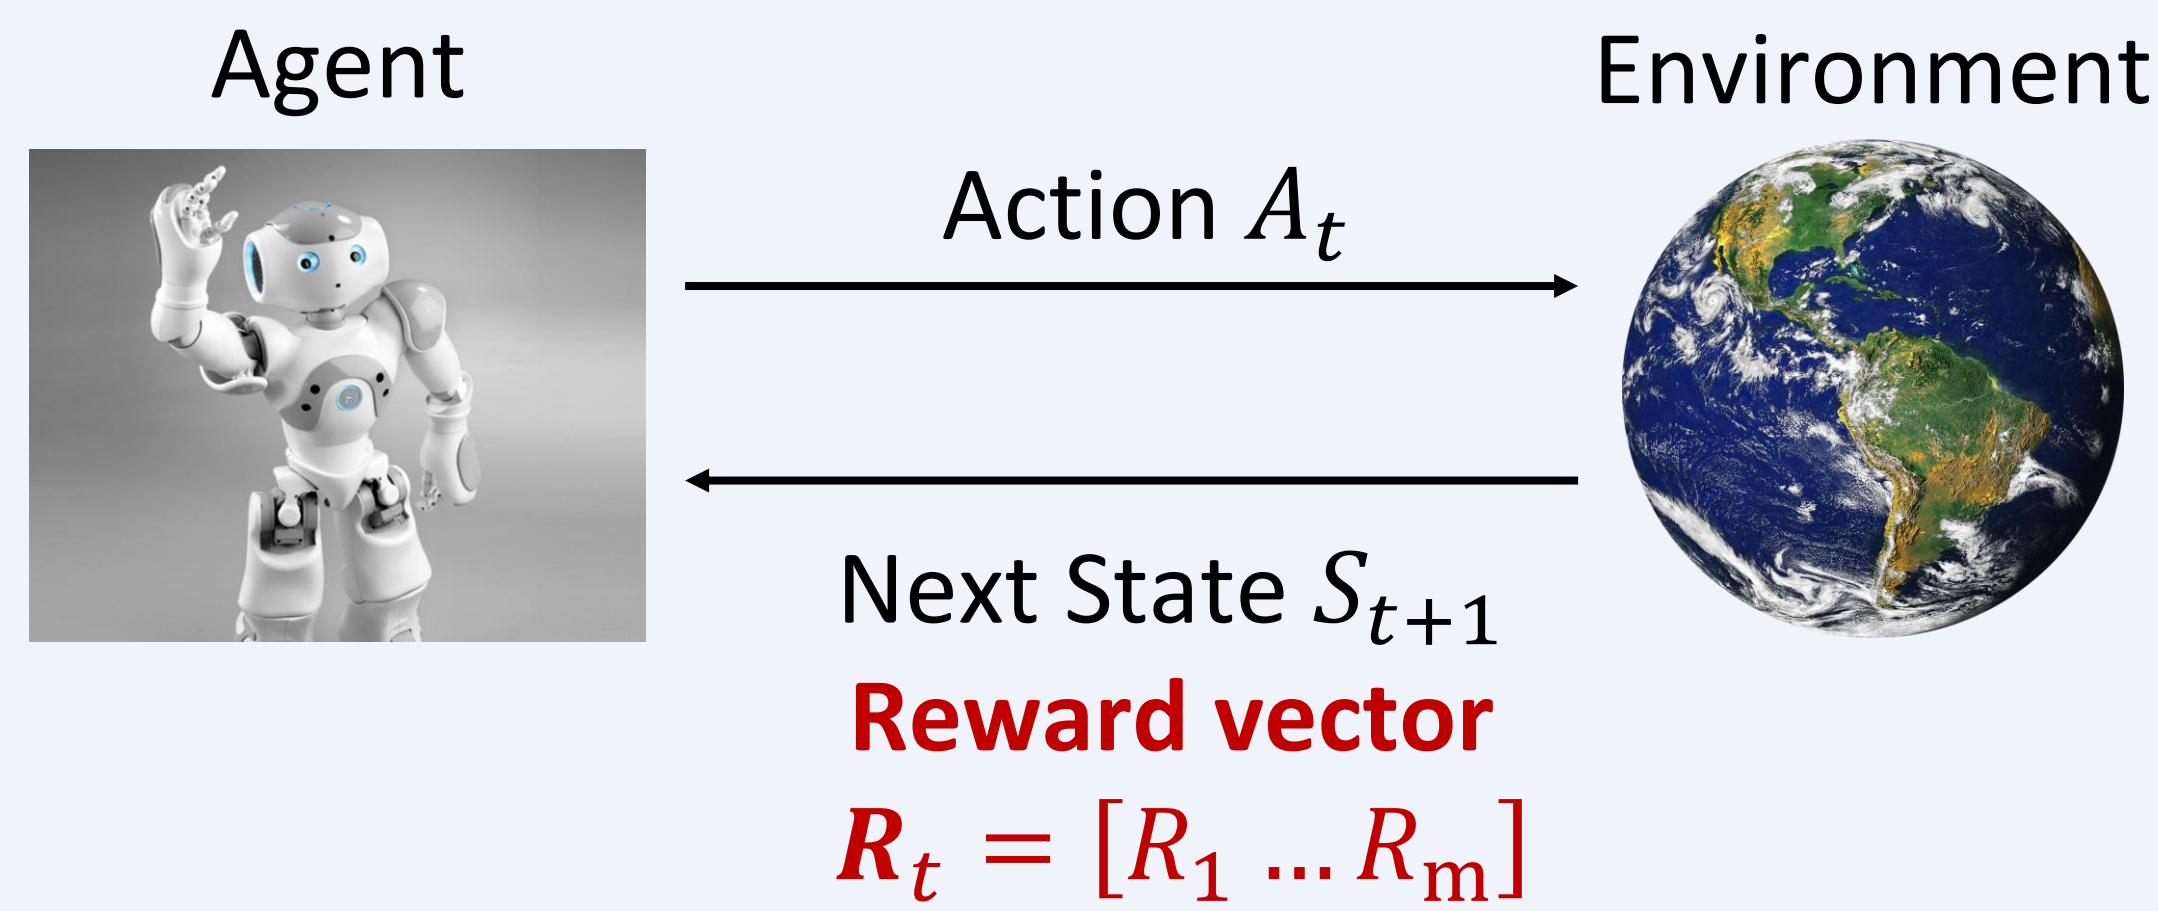


# A Toolkit for Reliable Benchmarking and Research in Multi-Objective Reinforcement Learning

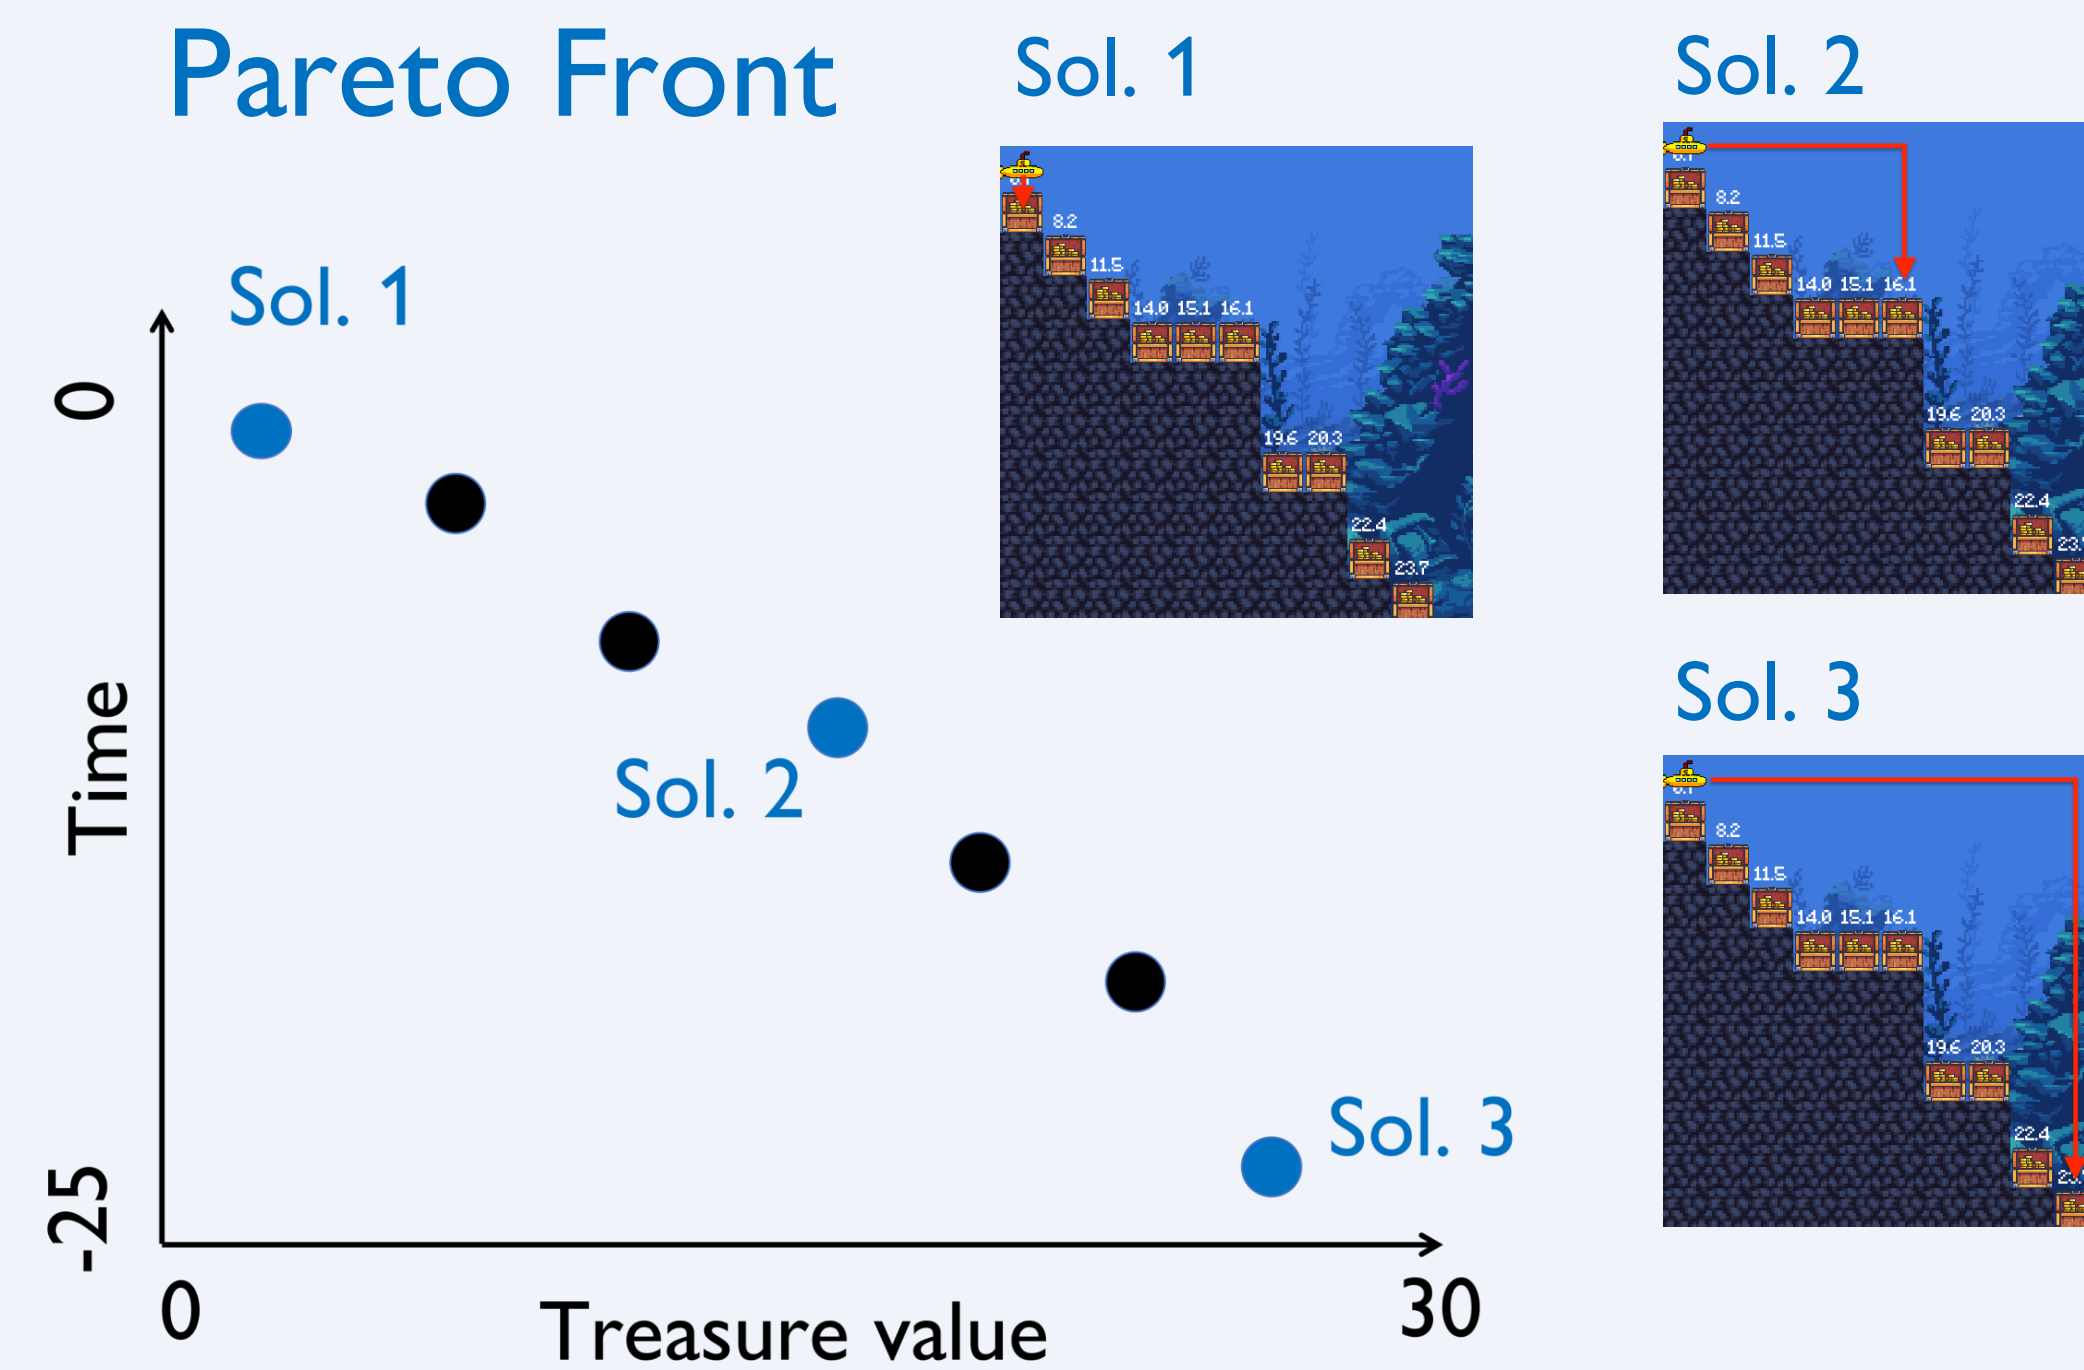
Florian Felten\* Lucas N. Alegre\* Ann Nowé Ana L. C. Bazzan  
 El Ghazali-Talbi Grégoire Danoy Bruno C. da Silva

florian.felten@uni.lu @FlorianFelten1  
 Inalegre@inf.ufrgs.br @Inalegre

## Multi-Objective RL



**Goal:** Learn a set of policies, each specialized in a different preference over the  $m$  conflicting objectives.



## Motivation

Current challenges in MORL:

1. Lack of standardized environments
2. Missing ready-to-use algorithms
3. Need to re-run expensive baselines for comparison

## 1. Environments

### MO-Gymnasium

- Standard API extending Gymnasium's API
- Part of the Farama Foundation
- MORL-specific wrappers
- Available on Pypi:

```
pip install mo-gymnasium
```

Over 20 environments available!



#### Gymnasium's Half-Cheetah

```
ctrl_cost = self.control_cost(action)
forward_reward = self._forward_reward_weight * x_velocity
reward = forward_reward - ctrl_cost
```

#### MO-Gymnasium's Half-Cheetah

```
reward = np.array([forward_reward, -ctrl_cost])
```

## 2. Algorithms

### MORL-Baselines

- Over 10 MORL algorithms implemented
- Clean, tested, and maintained code
- Automated logging and metrics
- Utilities for designing new algorithms

Algorithm	Single or multi-policy	Utility function	Observation space	Action space
MOQL (Van Moffaert et al., 2013a)	Single	Linear	Disc.	Disc.
EUPG (Rojiers et al., 2018a)	Single	Non-linear, ESR	Disc.	Disc.
MPMOQL (Van Moffaert et al., 2013a)	Multi	Linear	Disc.	Disc.
PQL (Van Moffaert and Nowé, 2014)	Multi	Non-linear, SER (*)	Disc.	Disc.
OLS (Rojiers, 2016)	Multi	Linear	/ (**)	/ (**)
Envelope (Yang et al., 2019)	Multi	Linear	Cont.	Disc.
PGMORL (Xu et al., 2020)	Multi	Linear	Cont.	Cont.
PCN (Reymond et al., 2022)	Multi	Non-linear, ESR/SER (*)	Cont.	Disc.
GPI-LS & GPI-PD (Alegre et al., 2023)	Multi	Linear	Cont.	Any
CAPQL (Lu et al., 2023)	Multi	Linear	Cont.	Cont.

## 3. Benchmark Results

- All MORL-Baselines trained on all MO-Gymnasium environments
- Weights & Biases dashboard
- Tracked code version and hyperparameters
- Integrated with openrlbenchmark providing CLI for plotting

