DIEDERIK M. ROIJERS

DIEDERIK.ROIJERS@HU.NL, HTTP://ROIJERS.INFO

On the necessity of using multiple objectives in future AI

ABOUT ME

Vrije Universiteit Brussel

AI Research Group

HU University of Applied Science Utrecht

Microsystems Technology, Institute of ICT

City of Amsterdam

Innovation Team

MY LITTLE ROBOT: A SIMPLE OPTIMISATION PROBLEM?

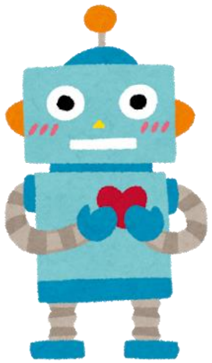

MY LITTLE ROBOT

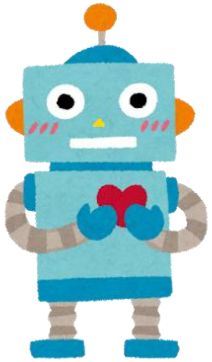

MY LITTLE ROBOT

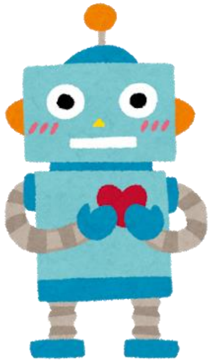

MY LITTLE ROBOT

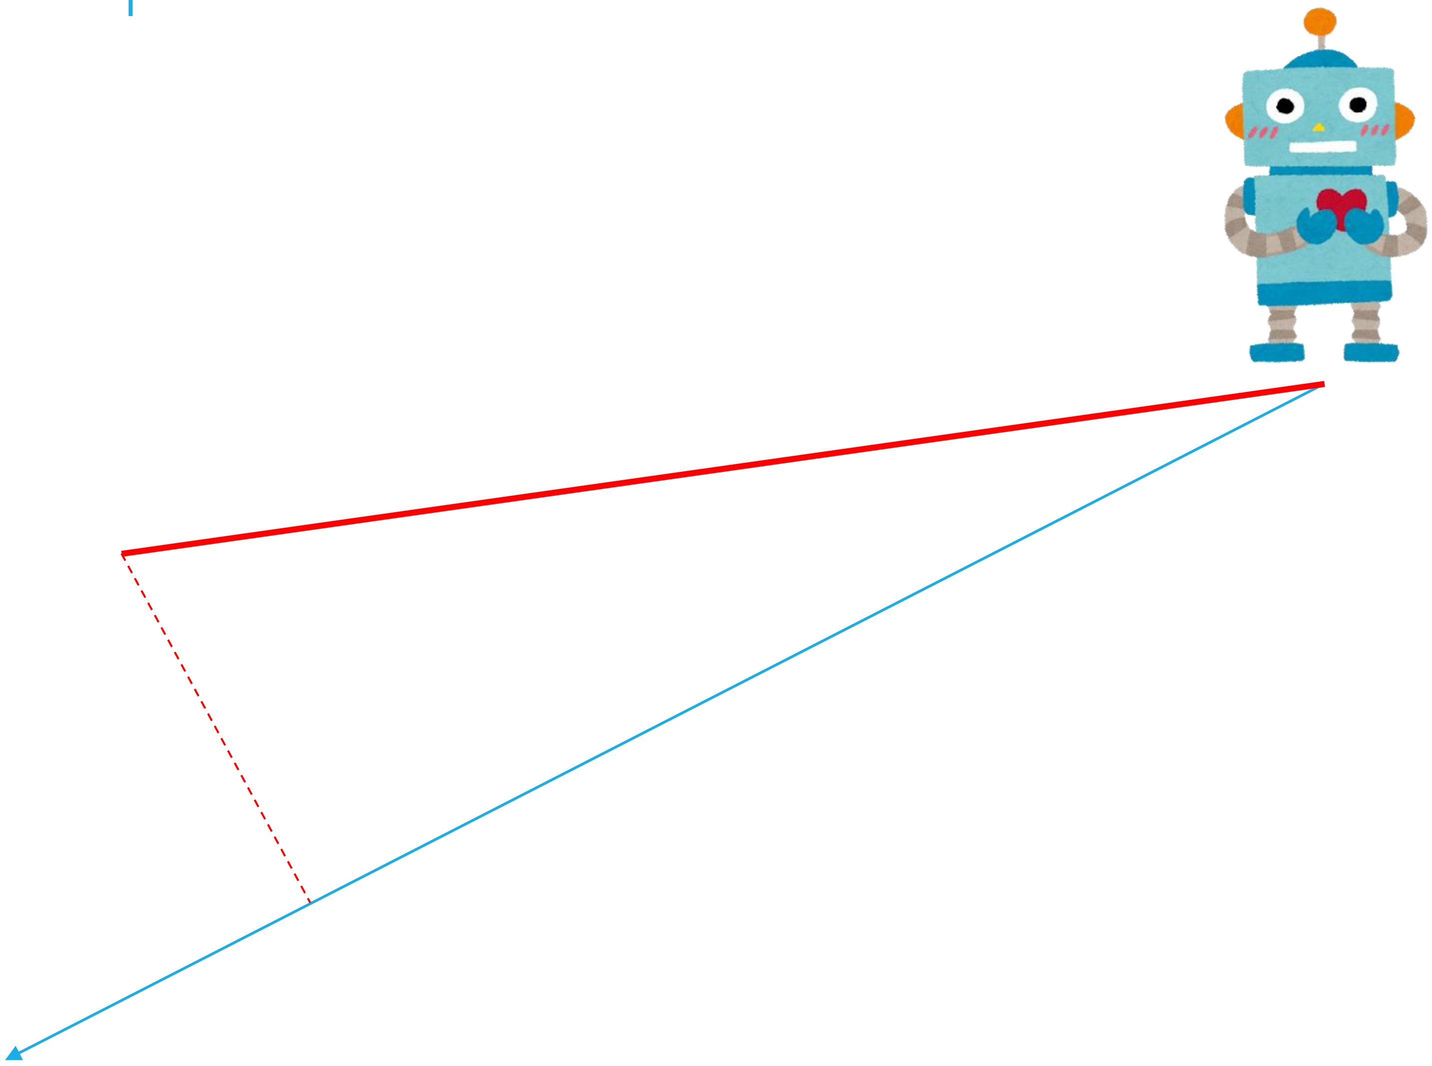

MY LITTLE ROBOT

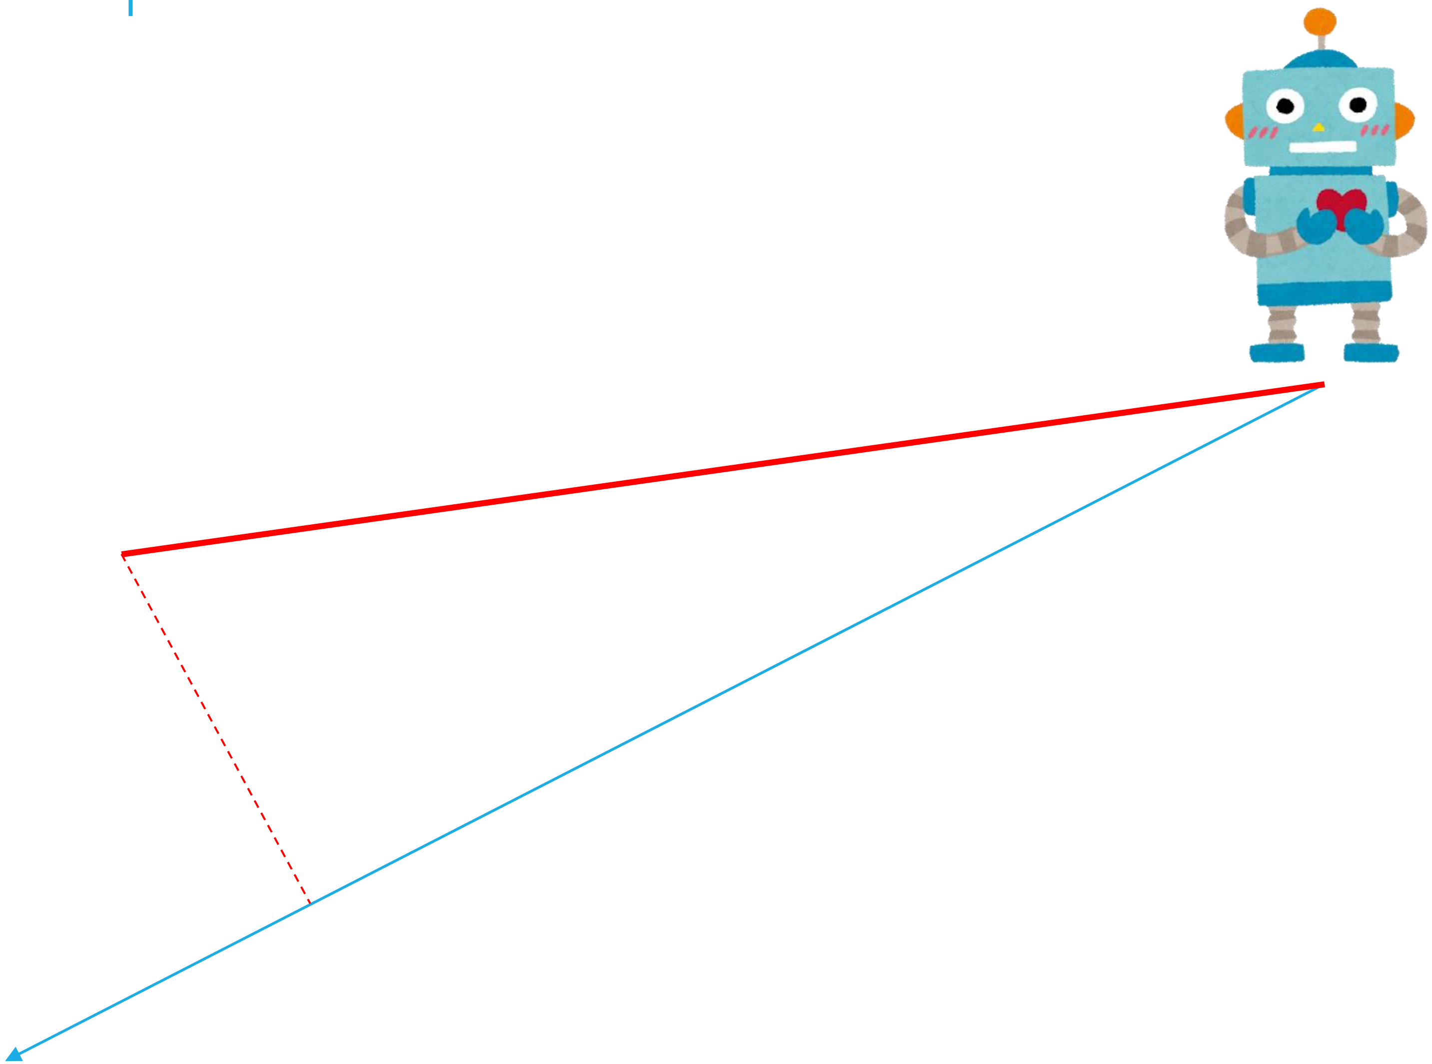

MY LITTLE ROBOT

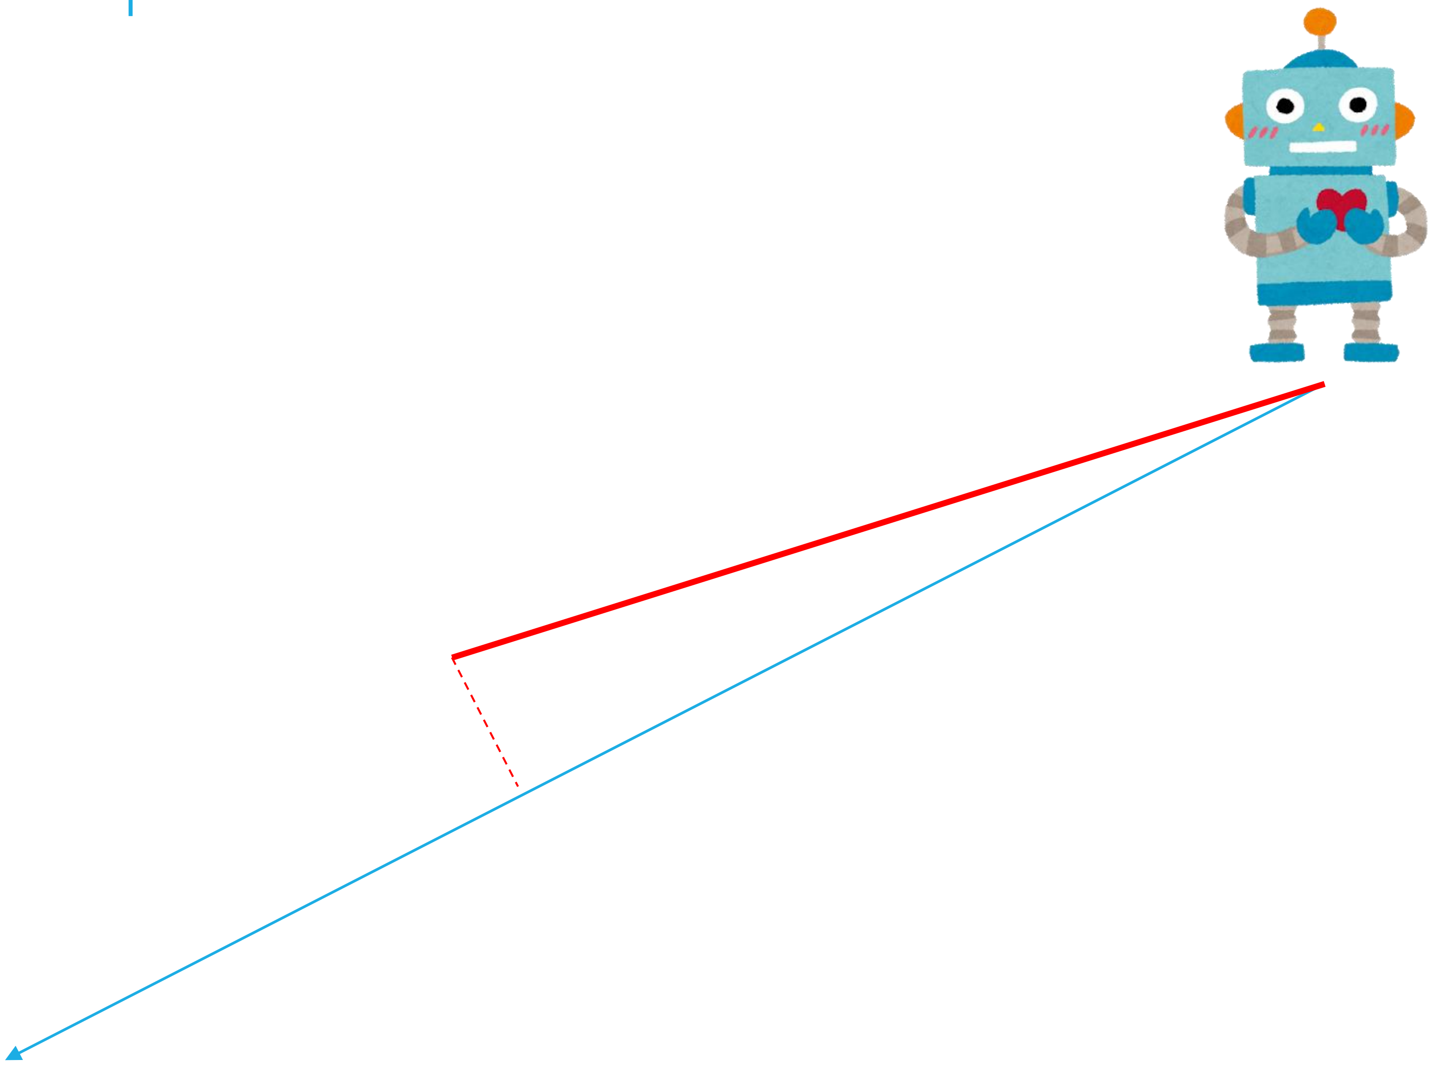

MY LITTLE ROBOT

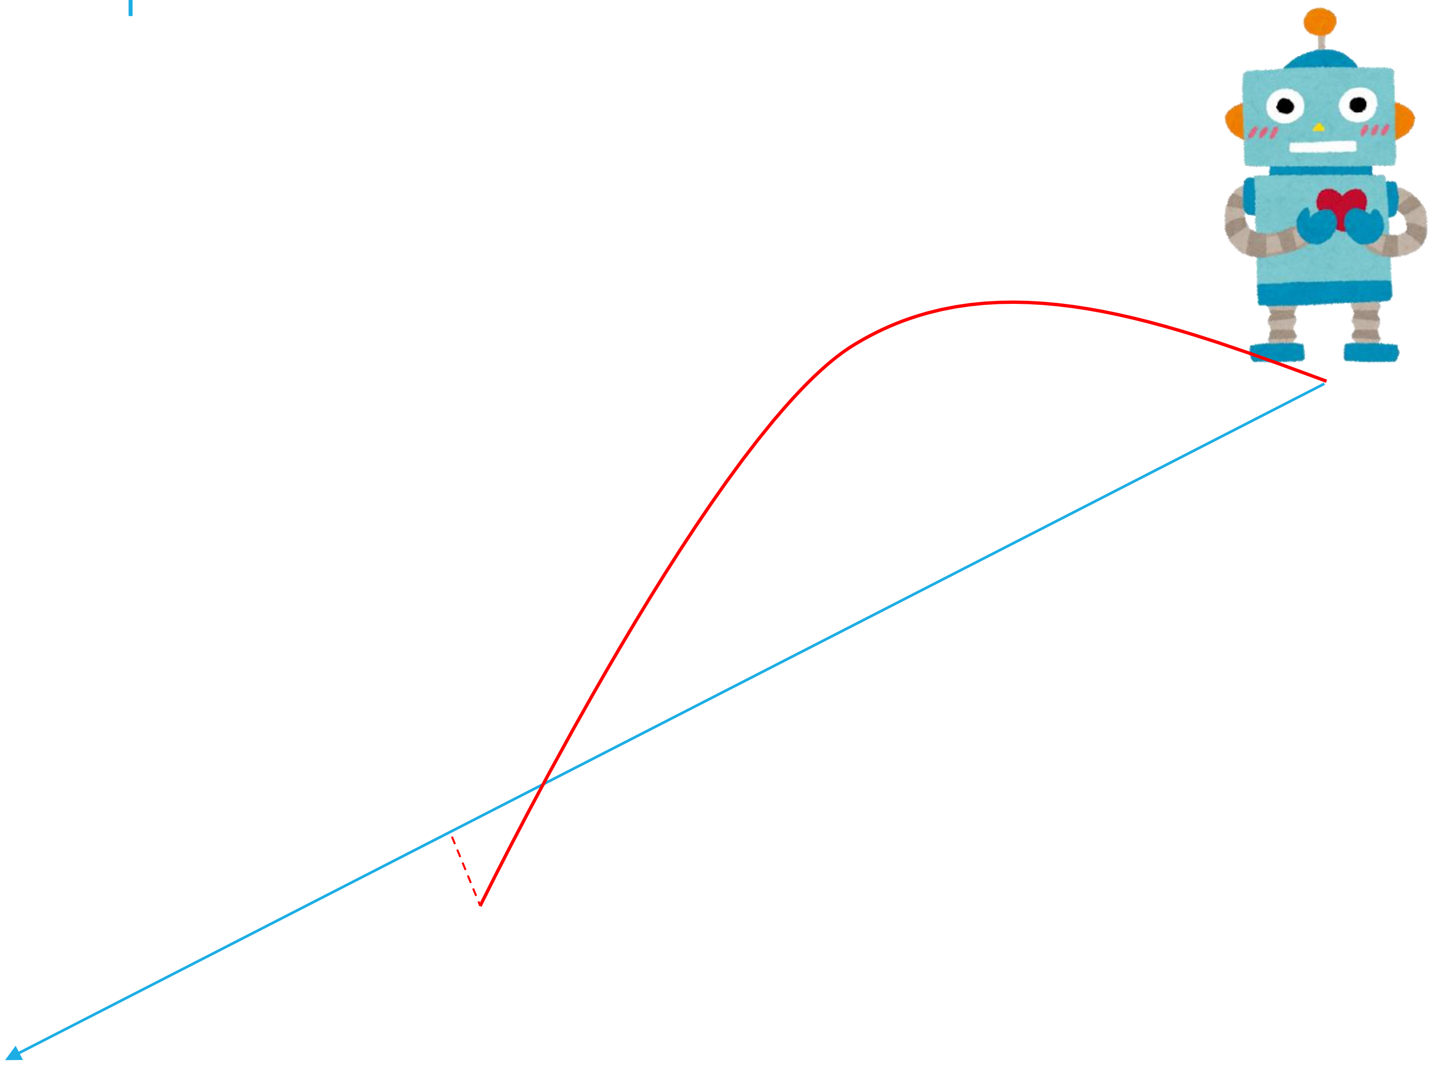

MY LITTLE ROBOT

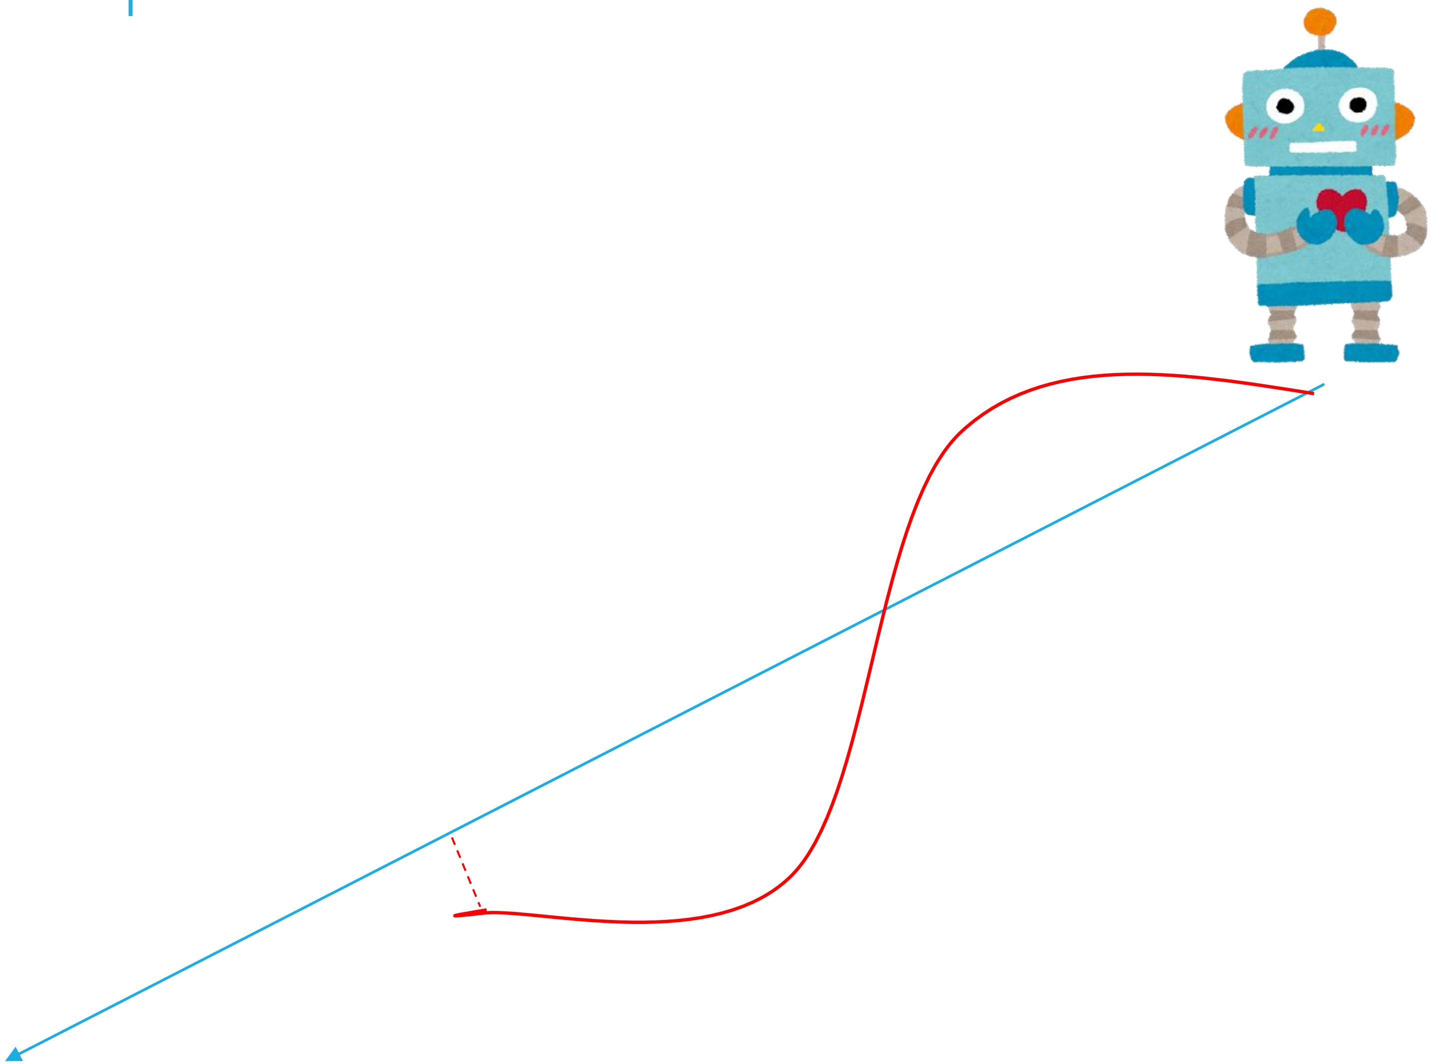

MY LITTLE ROBOT

It's never easy!

Move 30 sec in direction x: three objectives?

1. max projected length
2. min angle end point
3. min path length to get to end point

Engineering the a reward function until…

ROBOT

It's never easy!

Move 30 sec in direction x: three objectives?

1. max projected length
2. min angle end point
3. min path length to get to end point

Engineering the a reward function until… it works…

ROBOT

It's never easy!

Move 30 sec in direction x: three objectives?

1. max projected length
2. min angle end point
3. min path length to get to end point

Engineering the a reward function until… it works… sort of…

ROBOT

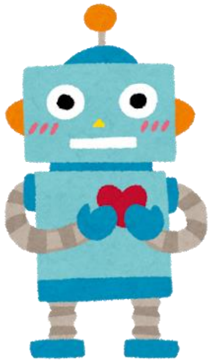

WAS THIS REALLY THE BEST?!

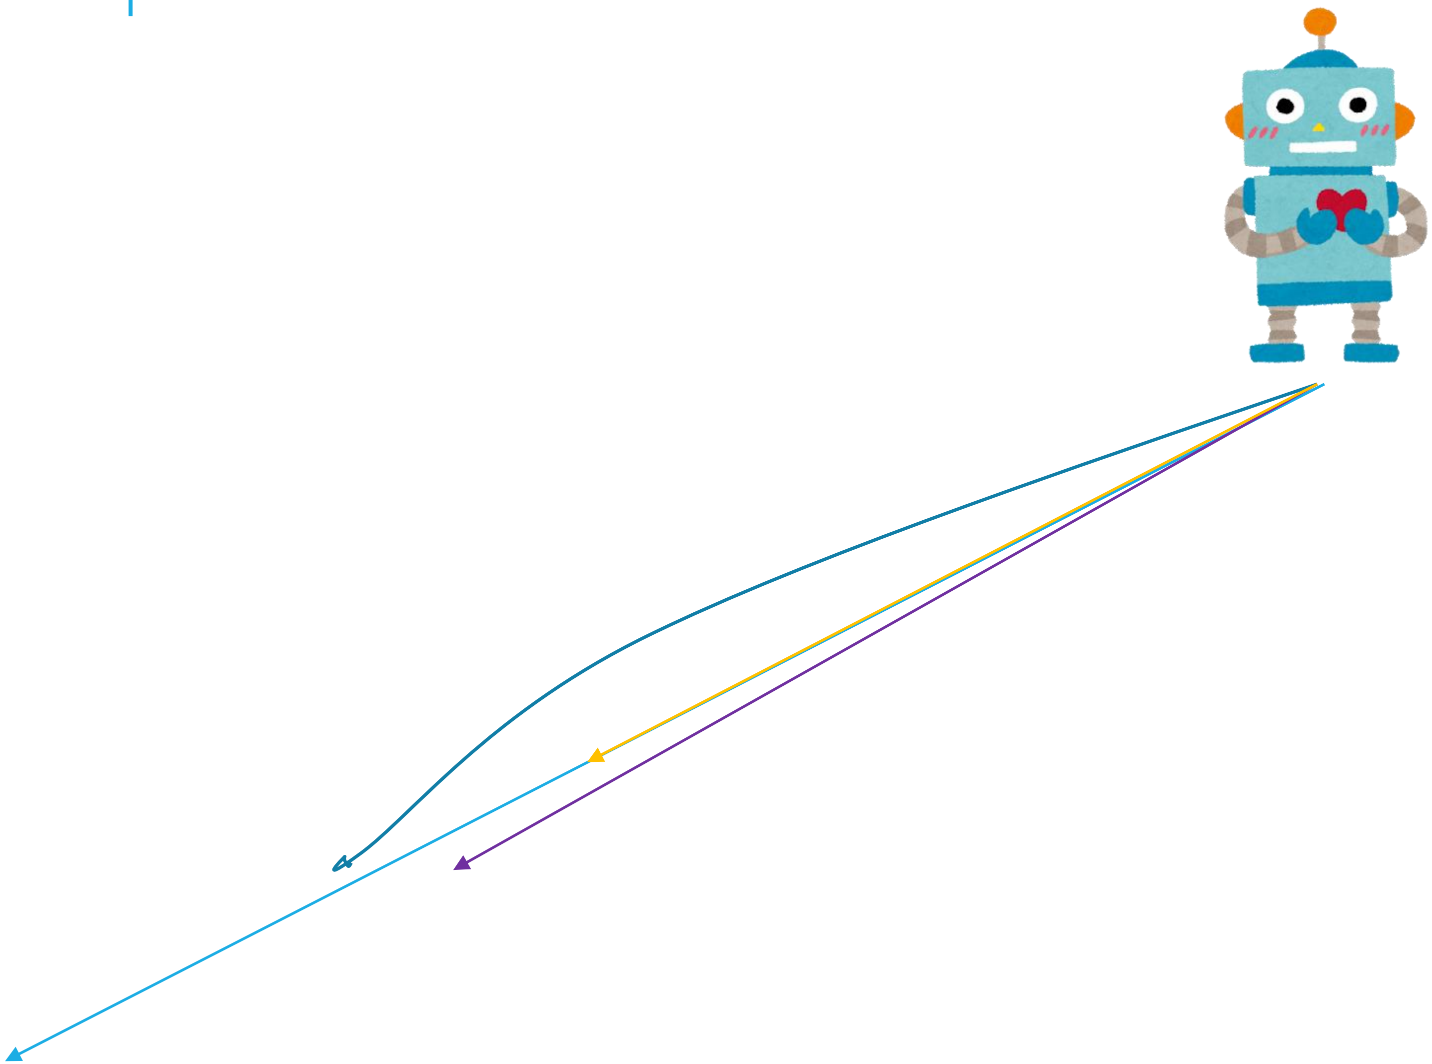

ROBOT STORY MORAL

Even simple problems have multiple objectives

- Bryce et al 2007: probabilistic planning *is* multi-objective

Engineering single-objective reward function is a semi-blind process

Single-objective reward functions make implicit decisions about what is optimal (without explicitly reasoning about it)

ROBOT STORY MORAL

Even simple problems have multiple objectives

- Bryce et al 2007: probabilistic planning *is* multi-objective

Engineering single-objective reward function is a semi-blind process

Single-objective reward functions make implicit decisions about what is optimal (without explicitly reasoning about it)

… might be okay, but we don’t know?

SELF-DRIVING CAR

- Detect objects to avoid collisions

- When unavoidable, limit the damage

- Mixture of policy optimisation/planning and ML/RL

SELF-DRIVING CAR: MULTI-OBJECTIVE

- Not all misclassifications of objects are equally damaging
  - Some concern the potential for loss of life or long-term damage to a person, others concern just property damage.
- Policy choices (optimisation) make trade-offs between key objectives

SELF-DRIVING CAR: ACCIDENT AVOIDANCE

AI takes risks with driver's life to save the life of a child running onto the street, and may cause damage to parked vehicles

What is fair?

How much risk is acceptable?

SELF-DRIVING CAR: ACCIDENT AVOIDANCE

AI takes risks with driver's life to save the life of a child running onto the street, and may cause damage to parked vehicles

What is fair?

How much risk is acceptable?

I don't know!

SELF-DRIVING CAR: ETHICS DOMAIN

Ethically very difficult choices

Some have suggested that using AI is unethical; while I'd argue that not using AI is unethical

Difficult choices have to be made *by people*

We need information about the values for different objectives to be able to make these choices

MORAL IMPLICATIONS

Medical treatment optmisation? Policy optimisation for robots in human environments? Insurance intake?

Is it even ethical to take a single-objective approach?

Human-aligned AI is a multi-objective problem (Vamplew et al., 2018)

PARTICULARIST ETHICS AND MULTIPLE OBJECTIVES

What the ethically optimal course of action is, is determined by the particular relevant factors in each situation. It is always possible to add factors that change the optimal action.

"Pessimistic" view: the utility function depends on the domain and situations in which we apply the AI

PRACTICAL NECESSITY

AI has an ever stronger impact

PRACTICAL NECESSITY

AI has an ever stronger impact

I don't trust researchers and engineers
to make the trade-offs between important objectives

PRACTICAL NECESSITY

AI has an ever stronger impact

I don’t trust researchers and engineers
to make the trade-offs between important objectives

And I don’t trust anybody to get it right in one go

DESIGNING AND MAINTAINING AI

We really need to see the alternatives

The responsible people need to take the shots, not the AI (researchers)

We really don’t want the designers/engineers of algorithms deciding what the (ethically / socially) right thing to do is

We need to be able to adjust in the face of new situations

(Hayes et al. 2022, A practical guide to multi-objective reinforcement learning and planning)

SELF-DRIVING CAR: ACCIDENT AVOIDANCE

We need decision-makers that are not the people that design the algorithms.

But algo's do need to take immediate action

It will make trade-offs between objectives

Were those okay?

Review and adjust

DESIGNING AND MAINTAINING AI

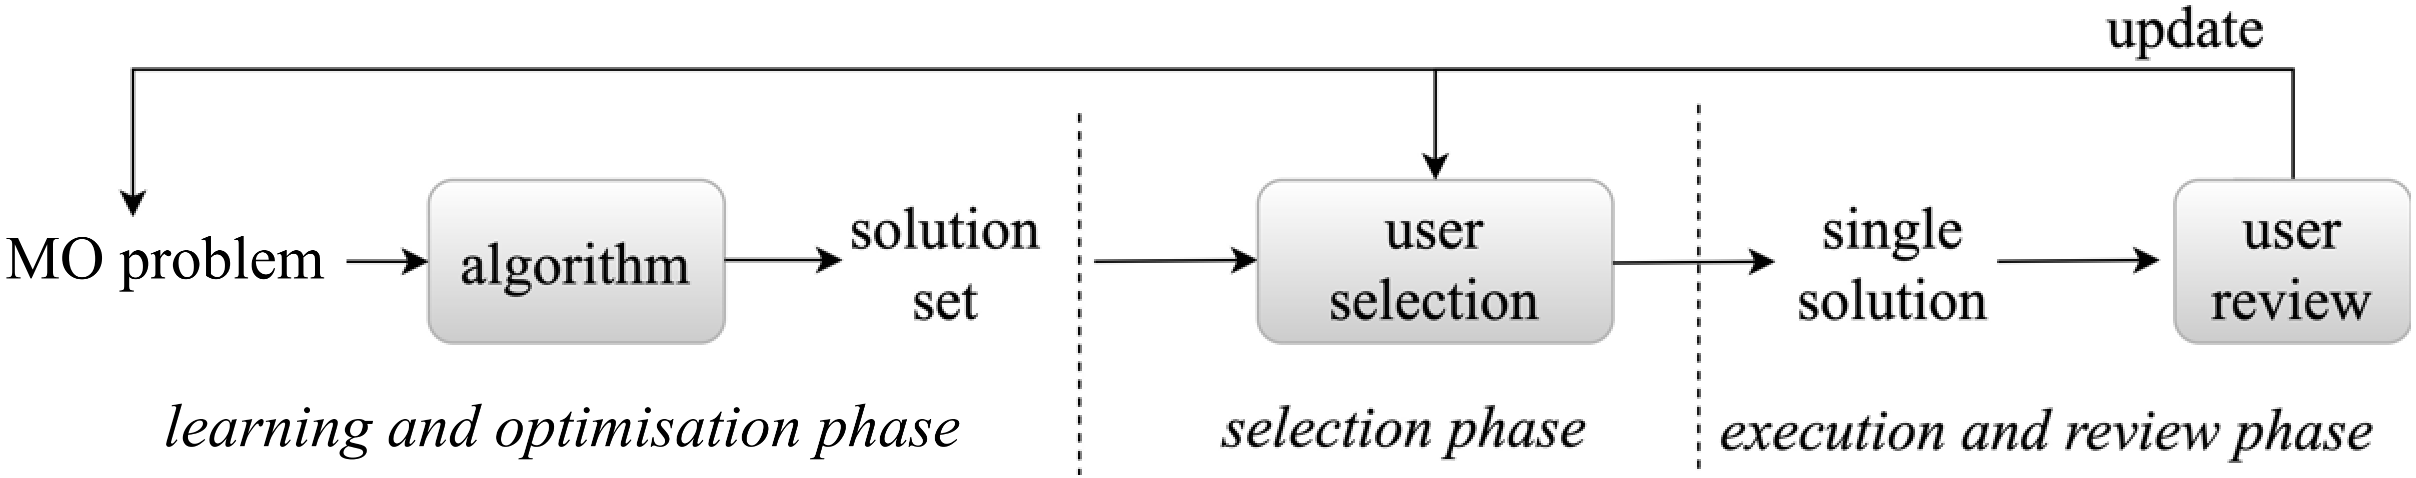

(Hayes et al. 2022, A practical guide to multi-objective reinforcement learning and planning)

SUMMARY OF FIRST PART

0) Multiple objectives are essential to many – if not most – real-world problems

1) Explicitly modelling multiple objectives is essential for explainable AI as well as human-aligned AI

2) That multiple objectives will help us make AI systems better maintainable

THE UTILITY-BASED APPROACH

to multi-objective learning and optimisation

FROM THE MORALS TO MORL

- Vector-valued value/payoff/fitness functions
- Meaningful objectives:
  - easy to define
  - easy to interpret the results

$$\mathbf{V} : \Pi \to \mathbb{R}^n$$

DECISION MAKERS

“Owners” of the utility

Utility-based approach $u :: \mathbb{R}^n \to \mathbb{R}$

Utility function can be implicit or explicit

Monotonically increasing in all objectives

UTILITY-BASED APPROACH

Derive your optimal set from:

- What you know about *u*
- How *u* is applied to derive the utility
- Which solutions are allowed

Useful methods, theory, and tricks can be used depending on correct positioning resulting from this derivation

ALTERNATIVE: AXIOMATIC APPROACH

Just assume you need the Pareto Front

It’s the most general solution set (minimal number of assumptions on *u*)

This can hurt efficiency, and does not consider all factors

CCS VERSUS PCS (PARETO FRONT)

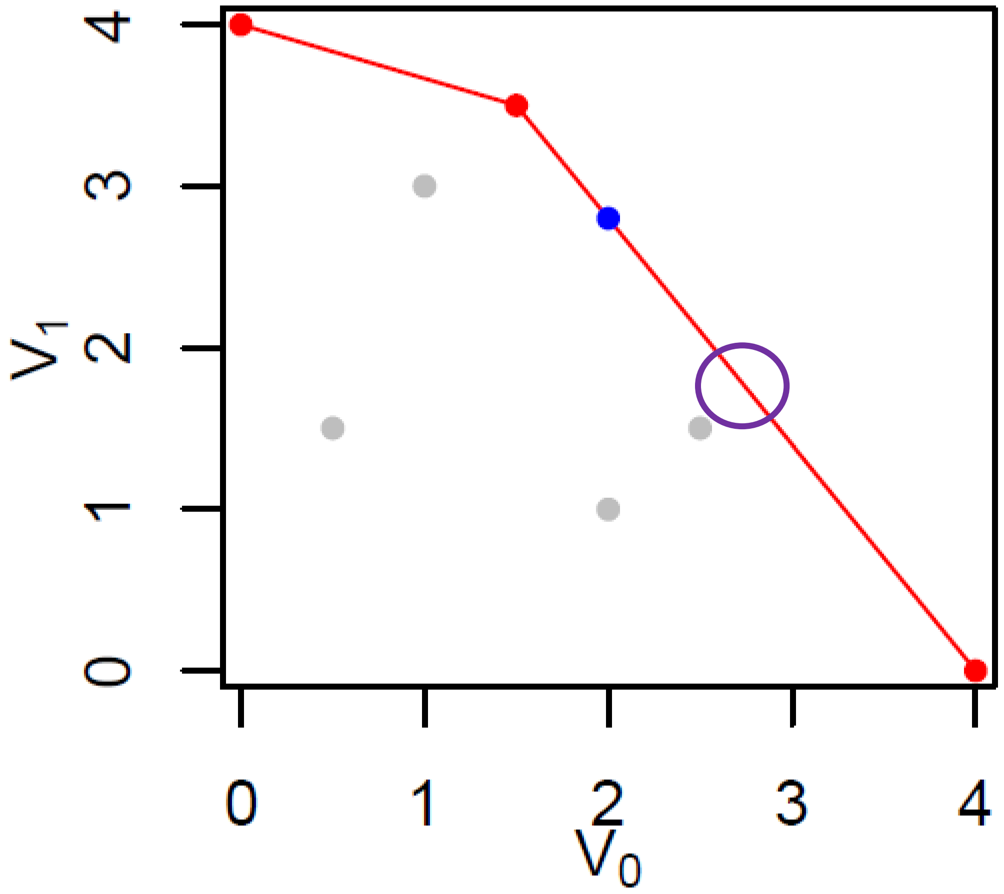

Linear *u*: Convex Coverage Set
(much smaller, easier to obtain)

CCS VERSUS PCS (PARETO FRONT)

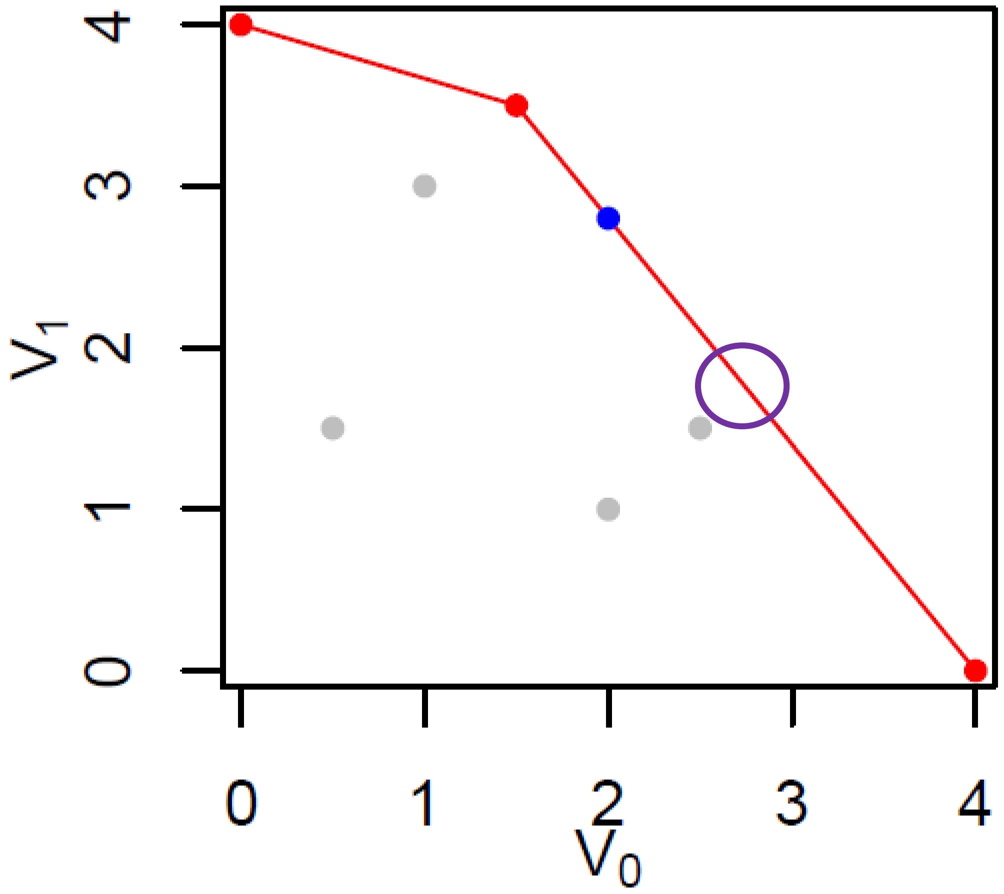

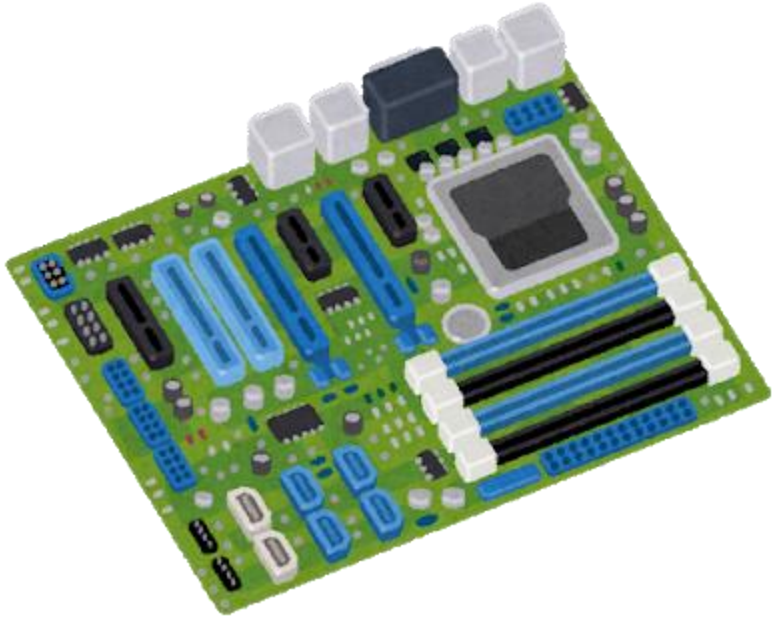

Stochastic policies are often OK

No Pareto front needed

FOR PRECISE DEFINITIONS SEE

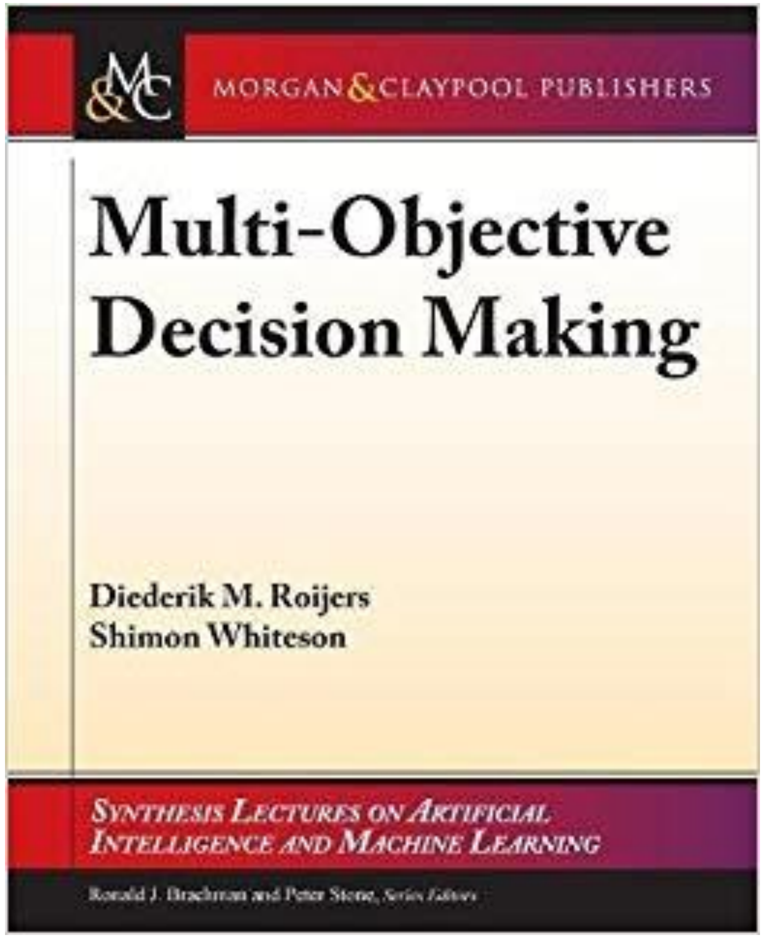

Diederik M. Roijers, Peter Vamplew, Shimon Whiteson, and Richard Dazeley - A Survey of Multi-Objective Sequential Decision-Making. *Journal of Artificial Intelligence Research*, 48:67–113, 2013.

CONVEX COVERAGE SETS

Viable in a lot of problems if stochastic policies are allowed

Linear utility functions distribute over expectations: for known weights single-objective methods still work. Very convenient!

For sequential planning under linear utility functions, POMDPs are a mathematically equivalent superclass of MOMDPs. No need to prove much (!) (convergence, etc.) Can take inspiration from POMDP methods.

INNER LOOP VERSUS OUTER LOOP

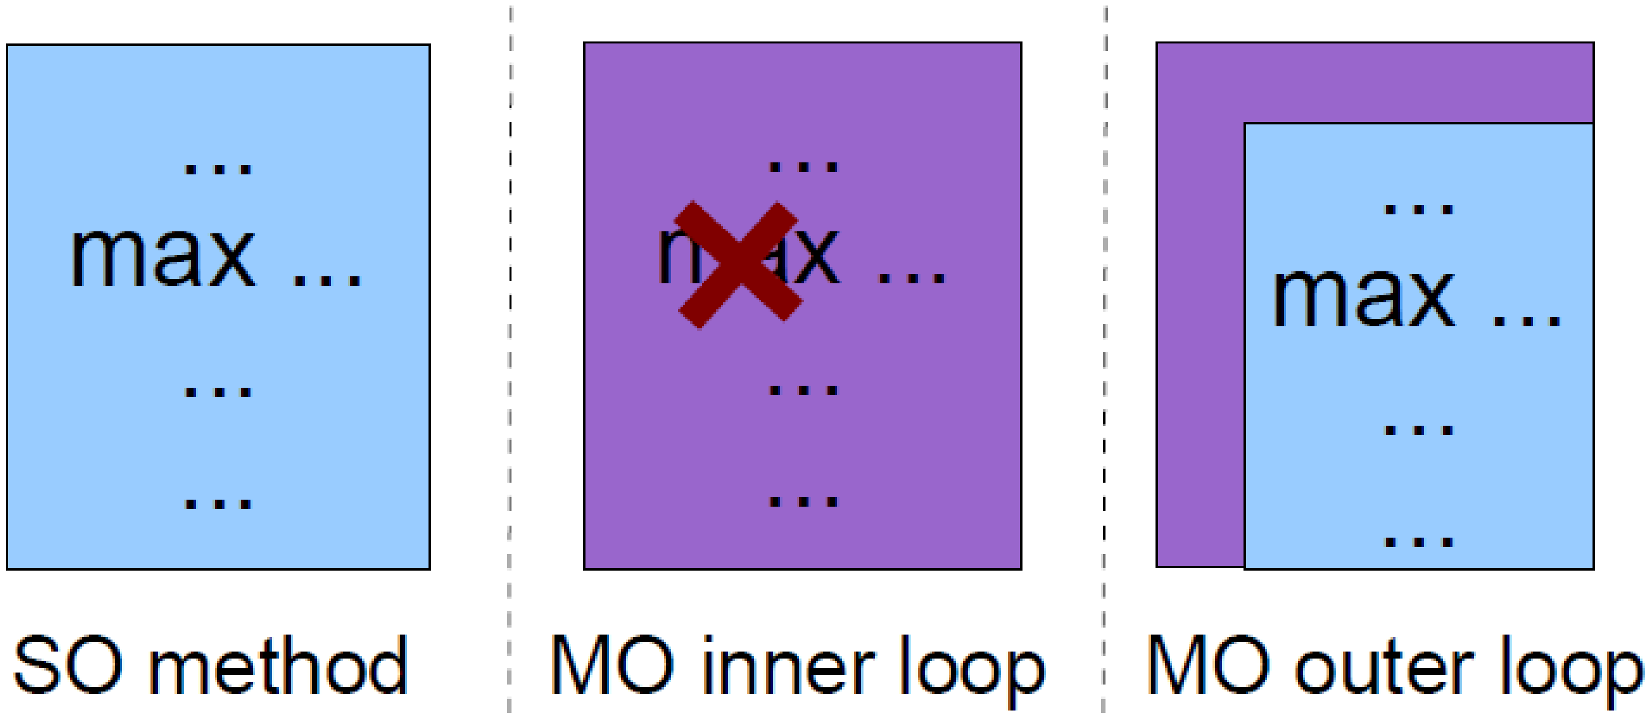

INNER LOOP VERSUS OUTER LOOP

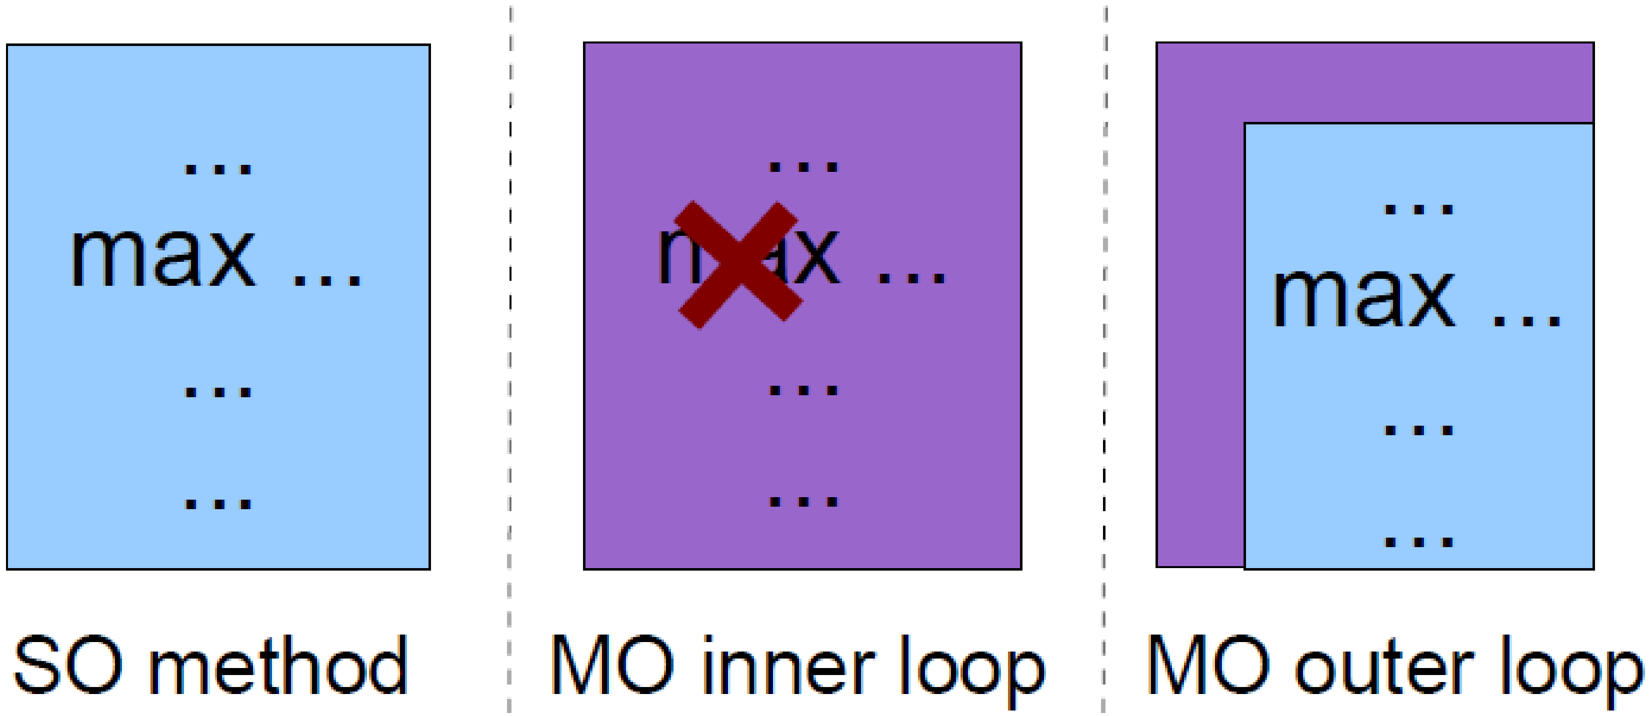

Outer loop methods (CCS) are easier, and faster for 2, 3 objectives

Inner loop methods scale better in the numbers of objectives

MEDICAL: SER? (WHEN TO APPLY *U*)

$$V_u^\pi = u\left(\mathbb{E}\left[\sum_{t=0}^{\infty}\gamma^t \mathbf{r}_t \mid \pi, \mu_0\right]\right)$$

MEDICAL: ESR!

$$V_u^\pi = \mathbb{E}\left[u\left(\sum_{t=0}^{\infty}\gamma^t \mathbf{r}_t\right) \mid \pi, \mu_0\right]$$

MEDICAL: ESR!

Setting can fundamentally change optimality

CONCLUSIONS AND FUTURE OUTLOOK

What we can do

What we should do

WHAT DOES MO ENABLE US TO DO

Reason about problems in a natural way (in meaningful statistics)

Inform human decision makers about viable alternatives

Adjust to changes in utility judgements

Helps us engineer human-aligned explainable AI solutions

Helps us make application of AI viable and flexible

Helps us make AI long-lived

MO CHALLENGES

We need:

- systems that model objectives explicitly

- that can interact with decision makers

- who may change the definition, and even the number of objectives

MO CHALLENGES

We need:

- systems that model objectives explicitly
- that can interact with decision makers
- who may change the definition, and even the number of objectives
- we cannot currently do this
- we need to extend our test horizons, long-term utility

MANY THANKS

Roxana Rădulescu, Willem Röpke, Conor F. Hayes, Zoltan Istvan Zardai, Matthieu Reymond, Matthew Macfarlane, Patrick Mannion, Ann Nowé, Peter Vamplew, Richard Dazeley, Luisa M. Zintgraf, Frans Oliehoek, Shimon Whiteson, Denis Steckelmacher, Eugenio Bargiacchi, Hélène Plisnier, Pieter Libin, Timothy Verstraeten, Matthijs T.J. Spaan, Mathijs de Weerdt, Joris Scharpff, Dirk Sierag, Maarten Inja, Chiel Kooijman, Maarten de Waard, Joost van Doorn, Daan Odijk, Maarten de Rijke, Gongjin Lan, Axel Abels, Tom Lennaerts, Felipe Leno Da Silva, Cyntia E.H. Nishida, Anna H. Reali Costa, Shang Wang, Xiaodong Nian, Athirai A. Irissappane, Ayumi Igarashi, Yijie Zhang, Dean Webb, Hossam Mossalam, Yannis Assael, Roberta Piscitelli, … and so many more people I've worked with over the years.

FAQ SLIDES

Multi-agent settings
Acknowledgements
SER vs ESR
Interactive settings

MULTI-AGENT SETTINGS

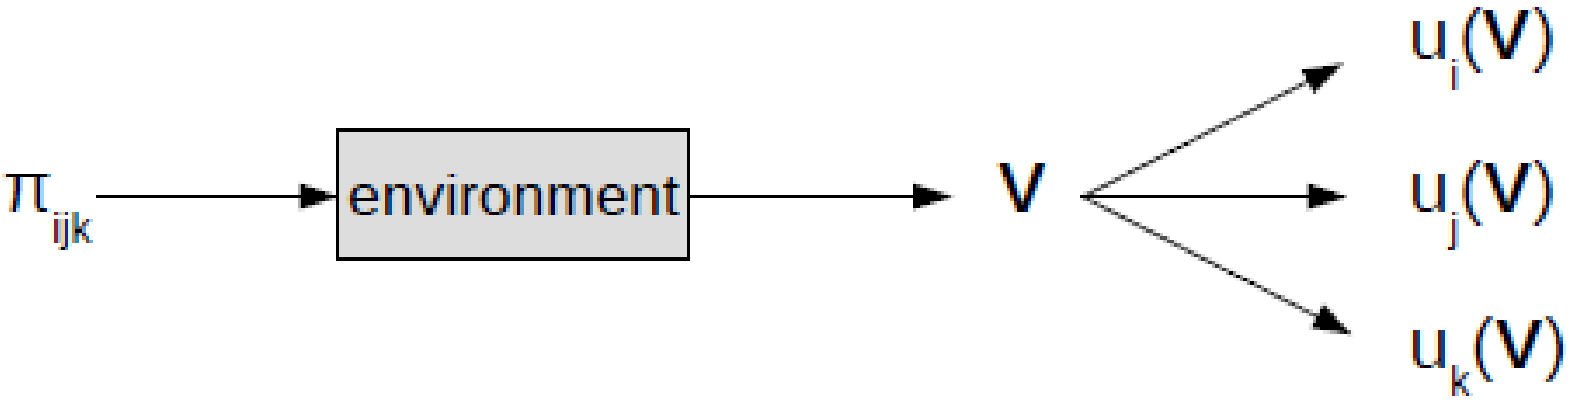

INTERACTIVE DECISION SUPPORT

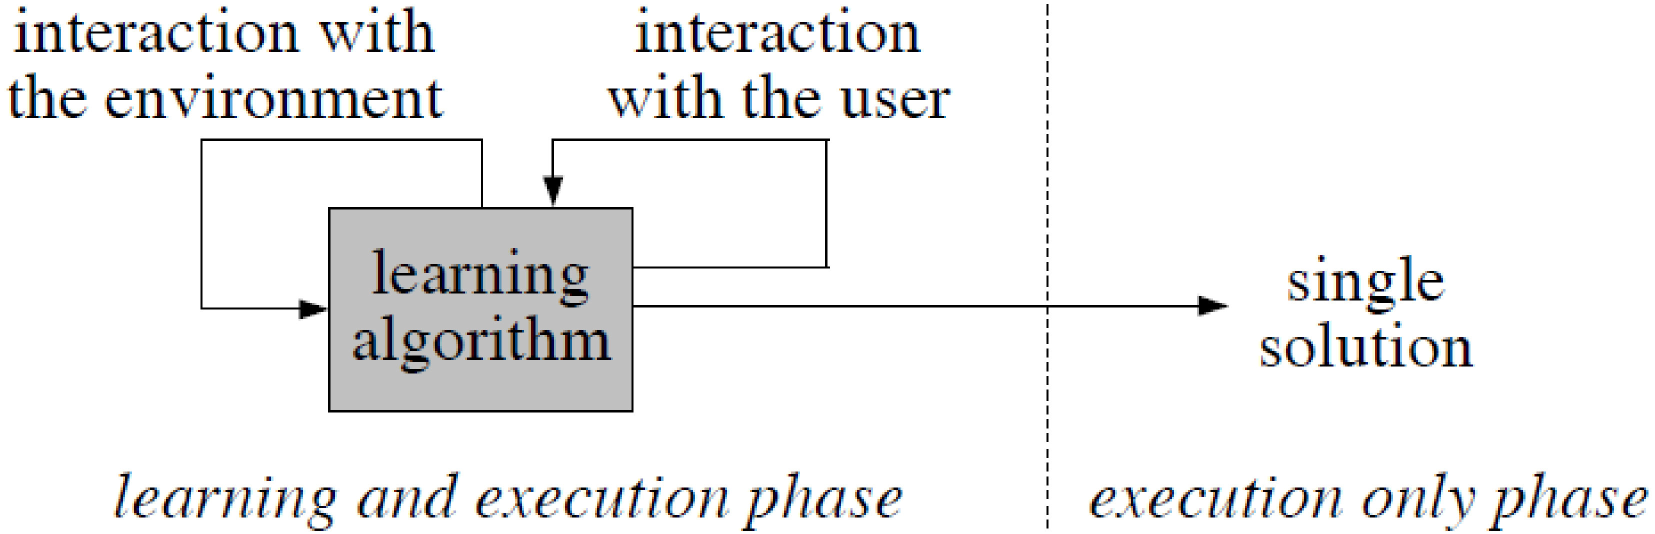

DYNAMIC WEIGHTS