

# A Demonstration of Issues with Value-Based Multiobjective Reinforcement Learning Under Stochastic State Transitions

Peter Vamplew  
Federation University Australia  
Mt Helen, Victoria, Australia  
p.vamplew@federation.edu.au

Cameron Foale  
Federation University Australia  
Mt Helen, Victoria, Australia  
c.foale@federation.edu.au

Richard Dazeley  
Deakin University  
Waurin Ponds, Victoria, Australia  
r.dazeley@deakin.edu.au

## ABSTRACT

We report a previously unidentified issue with model-free, value-based approaches to multiobjective reinforcement learning in the context of environments with stochastic state transitions. An example multiobjective Markov Decision Process (MOMDP) is used to demonstrate that under such conditions these approaches may be unable to discover the policy which maximises the Scalarised Expected Return, and in fact may converge to a Pareto-dominated solution. We discuss several alternative methods which may be more suitable for maximising SER in MOMDPs with stochastic transitions.

## KEYWORDS

multiobjective reinforcement learning, multiobjective MDPs, stochastic MDPs

## 1 INTRODUCTION

Multiobjective reinforcement learning (MORL) aims to extend the capabilities of reinforcement learning (RL) methods to enable them to work for problems with multiple, conflicting objectives [12]. RL algorithms generally assume that the environment is a Markov Decision Process (MDP) in which the agent is provided with a scalar reward after each action, and must aim to learn the policy which maximises the long-term return based on those rewards. In contrast MORL algorithms operate within multiobjective MDPs (MOMDPs), in which the reward terms are vectors, with each element in the vector corresponding to a different objective. This creates a number of new issues to be addressed by the MORL agent. Most notably there may be multiple policies which may be optimal (in terms of Pareto optimality), and which policy the agent should learn is not immediately obvious.

In the utility-based paradigm of MORL [12, 29] it is assumed that the preferences of the user can be defined in terms of a parameterised utility function  $f$ , and that the aim of the agent should be to learn the policy which produces vector returns which maximises the utility to the user as defined by  $f$ .

Various approaches have been explored for the form of the utility function – some may be better suited to express the preference of the user within a particular problem domain, while others offer benefits from an algorithmic perspective. A simple weighted linear scalarisation has been widely used because of its simplicity (for example, [1, 3, 8]). Linear scalarisation transforms an MOMDP into an equivalent single-objective MDP, and enables existing RL approaches to be directly applied [12]. However for many tasks this may not be able to accurately represent the utility of the user, and so may fail to discover the policy which is optimal with regards to their

true utility. As a result numerous non-linear scalarisation functions have been explored in the literature (for example, [4, 23, 24]) – these tend to produce algorithmic complications, but also may better represent the true preferences of the user.

As well as the choice of scalarisation function and parameters, a second factor must be considered within this utility-based paradigm – the time-frame over which the utility is being maximised. Roijers et al. [12] identified two distinct possibilities. The agent may aim to maximise the expected scalarised return (ESR). That is, it is assumed the returns are first scalarised, and then this agent aims for the policy which maximises the expected value of that scalar. This ESR approach is suited to problems where the aim is to maximise the expected outcome within any individual episode. For example, when producing a treatment plan for a patient which trades off the likelihood of a cure versus the extent of negative side-effects - any individual patient will only undergo this treatment once, and so they care about the utility obtained within that specific episode.

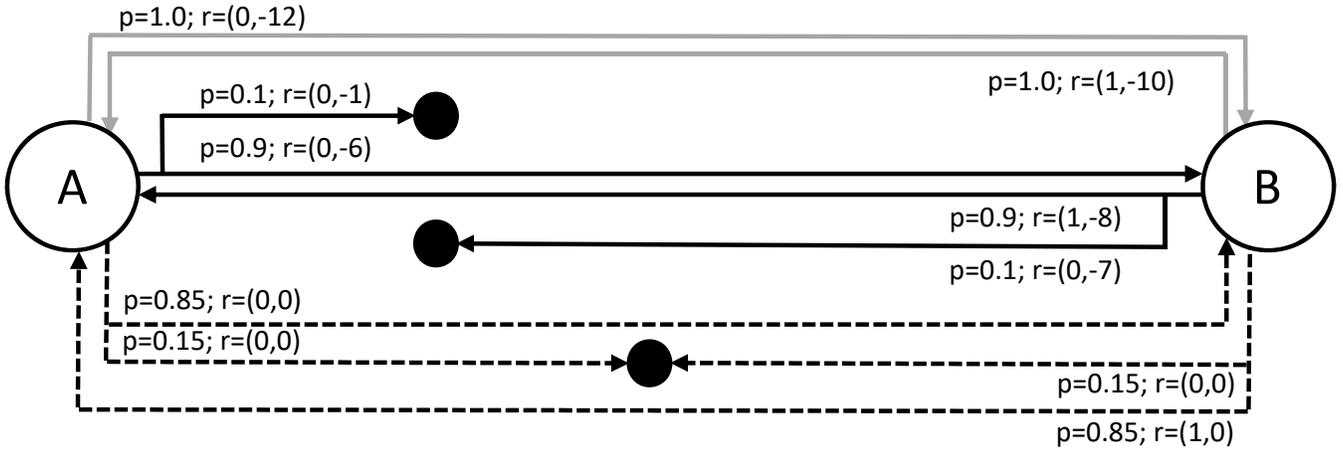
In other contexts we may be concerned about the mean utility received over multiple episodes. In this situation the agent should aim to maximise the scalarised expected return (SER) - that is, it estimates the expected vector return per episode, and then maximises the scalarisation of that expected return. As demonstrated in Roijers et al. [11], the optimal policy for a particular MOMDP under the ESR and SER setting may differ considerably, even if the same scalarisation function and parameters are used in both cases.

As noted by Roijers et al. [11] and Rădulescu et al. [10] much of the existing work in MORL has considered SER optimization, although this has often been implicit rather than explicitly stated. Much of this SER-focused work has been based on benchmark environments such as those of Vamplew et al. [18], the majority of which are deterministic MOMDPs.

In this paper we demonstrate by example that the model-free value-based methods previously widely used in MORL research may fail to maximise the SER utility when applied to MOMDPs with stochastic state transitions.

## 2 SPACE TRADERS: AN EXAMPLE STOCHASTIC MOMDP

As shown in Figure 1 the Space Traders MOMDP is a finite-horizon task with a horizon of 2 time-steps. It consists of two non-terminal states, with three actions available in each state. The agent starts at its home planet (state A) and must travel to another planet (state B) to deliver a shipment, and then return to State A with the payment. The agent receives a reward with two elements - the first is 0 on all actions, except that a reward of 1 is received when the agent successfully returns to state A, while the second element is a negative value reflecting the time taken to execute the action.



**Figure 1: The Space Traders MOMDP.** Solid black lines show the Direct actions, solid grey lines show the Indirect actions, and dashed lines indicate Teleport actions. Solid black circles indicate terminal (failure) states.

State	Action	P(success)	Reward on success	Reward on failure	Mean reward
A	Indirect	1.0	(0,-12)	n/a	(0,-12)
	Direct	0.9	(0,-6)	(0,-1)	(0,-5.5)
	Teleport	0.85	(0,0)	(0,0)	(0,0)
B	Indirect	1.0	(1,-10)	n/a	(1,-10)
	Direct	0.9	(1,-8)	(0,-7)	(0.9,-7.9)
	Teleport	0.85	(1,0)	(0,0)	(0.85,0)

**Table 1: The probability of success and reward values for each state-action pair in the Space Traders MOMDP.**

Policy identifier	Action in state A	Action in state B	Mean return
II	Indirect	Indirect	(1,-22)
ID	Indirect	Direct	(0.9,-19.9)
IT	Indirect	Teleport	(0.85,-12)
DI	Direct	Indirect	(0.9,-14.5)
DD	Direct	Direct	(0.81,-12.61)
DT	Direct	Teleport	(0.765,-5.5)
TI	Teleport	Indirect	(0.85,-8.5)
TD	Teleport	Direct	(0.765,-6.715)
TT	Teleport	Teleport	(0.7225,0)

**Table 2: The mean episodic return vector for each of the nine deterministic policies available for the Space Traders MOMDP.**

There are three possible pathways between the two planets. The direct path (actions shown by solid black lines in Figure 1) is fairly short, but there is a risk of the agent being waylaid by space pirates and failing to complete the task. The indirect path (grey lines) avoids the pirates and so always leads to successful completion of the mission, but takes longer. Finally the recently developed teleportation system (dashed lines) allows instantaneous transportation, but has a higher risk of failure. The figure also details the probability of success, and the reward for the mission-success and time objectives for each action – due to variations in local conditions such as solar winds and the location of the space pirates, the time values for the outward and return journeys on a particular path may vary.

Table 1 summarises the transition probabilities and rewards of the MOMDP, and also shows the mean immediate reward for each action from each state, weighted by the probability of success.

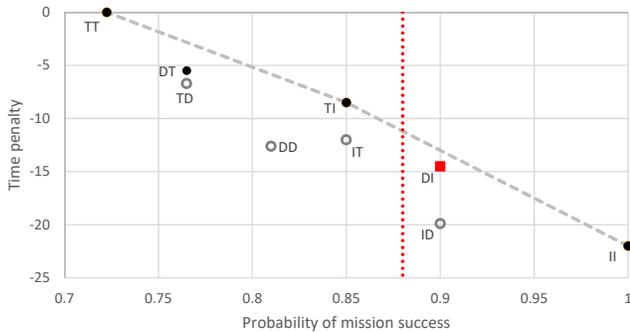
As there are three actions from each state there are a total of nine deterministic policies available to the agent. The mean reward per episode for each of these policies is shown in Table 2 and illustrated in Figure 2. The solid points in the figure highlight the policies which belong to the Pareto front, and the dashed grey line indicates the convex hull (only those policies lying on the convex hull can be

located via methods using linear scalarisation – this set of policies is referred to as the Convex Coverage Set [13]).

For the remainder of the paper we will assume that the agent’s aim is to minimise the time taken to complete the delivery and return to A, subject to having at least an 88% probability of successful completion. That is, the scalarisation function  $f(\vec{v}) = v_2$  if  $v_1 > 0.88$  and  $-\infty$  otherwise. The optimal policy for this aim is to follow the direct path to B and then the indirect path back to A (policy DI).

### 3 APPLYING MODEL-FREE VALUE-BASED MORL METHODS TO SPACE TRADERS

In this section we will discuss how some of the value-based MORL methods previously used in the literature would perform on the Space Traders MOMDP. All the methods discussed are assumed to be based on a multiobjective extension of model-free value-based RL algorithms such as Q-Learning or SARSA – for example see [25, p. 3668]. For the purposes of this section we will restrict discussion to single-policy methods in which the scalarisation function  $f$  is



**Figure 2: The mean return per episode for the nine possible deterministic policies for the Space Traders MOMDP. Each policy’s return is labelled with a bigram specifying its actions. I, D, T refer to the indirect, direct and teleport actions so, for example, policy DI selects the direct action in state A and the indirect action in state B. Solid markers indicate policies which are members of the Pareto-front, and hollow markers indicate dominated policies. The dashed grey lines illustrate the convex hull formed by mixture combinations of the policies which make up the Convex Coverage Set (CCS). The dashed red vertical line indicates the threshold value of 0.88 for the probability of mission success, and the red square marker is the DI policy which is optimal for that setting of the threshold.**

used to filter the multiple Pareto-optimal policies which may be available so as to obtain a single policy which is optimal with regards to  $f$ . Multiple-policy MORL methods will be discussed in Section 5.

All methods learn vector-valued estimated Q-values, but differ in terms of the scalarisation or ordering method used to perform action-selection, and the characteristics on which the Q-value and policy are conditioned.

### 3.1 Linear scalarisation

A simple approach to MORL is to apply a linear weighted scalarisation to the elements of the Q-value vector prior to selecting the greedy action. As mentioned earlier, this converts the MOMDP into an equivalent MDP, and so the Q-values and action-selection need only be conditioned on the current state of the MDP. However it is well-known that methods using linear scalarisation are unable to identify solutions which do not lie on the convex hull of the Pareto front [22]. Clearly from Figure 2 this is the case for policy DI, and so linear methods will not be able to converge to this policy. This result is not surprising and we mention it here simple for the sake of completeness.

### 3.2 Non-linear scalarisation

A variety of non-linear scalarisation methods have been explored in the MORL literature [4, 23, 24]. The non-linear nature of the scalarisation function means that the assumption of additivity underlying the Bellman equation no longer applies. In order to deal with this, both the choice of action and the Q-values must be conditioned not

Action in state A	Policy	$Q(A, a)$
Indirect	ID	(0.9, -19.9)
Direct	DD	(0.81, -12.61)
Teleport	TD	(0.765, -6.715)

**Table 3: The Q-values which will be learned for each action in state A, under the assumption that the Direct action will be selected in State B.**

only on the current state of the environment, but also on rewards received so far by the agent during this episode [5, 11]. That is, if the scalarisation function is  $f$  then at time  $k$  the agent will select the action  $a$  which maximises the value of  $f(Q(s_k, a, \sum_{t=1}^k r_t) + \sum_{t=1}^k r_t)$ .

For the purposes of the following discussion we will assume that  $f$  is the thresholded lexicographic ordering operator (TLO) [4, 6], and that a thresholding parameter of 0.88 is applied to the first element of the Q-value vector. The intention here is to maximise the value of the second objective (i.e. minimise time), subject to achieving the threshold level for the first objective. If this operator could be applied directly to the mean returns of each policy from Table 2, then clearly policy DI would be selected.

However if we consider how the TLO operator selects actions during the execution of a policy, then a different result will emerge. Regardless of the path selected at state A, if state B is successfully reached then a zero reward will have been received by the agent. Therefore the choice of action at state B is independent of the previous action. Looking at the mean action values reported in Table 1, it can be seen that action T will be eliminated as it fails to meet the threshold for the first objective, and that action D will be preferred over I as both meet the threshold, and D has a superior value for the time objective. So it can already be seen that this agent will not converge to the desired policy DI.

Knowing that action D will be selected at state B, we can calculate the Q-values for each action at state A, as shown in Table 3. The TLO action selector will eliminate actions D and T from consideration as neither meets the threshold of 0.88 for the probability of success. Action I will be selected giving rise to the overall policy ID. Not only is this not the desired DI policy, but as is evident from Figure 2 its average outcome is in fact Pareto-dominated by DI.

## 4 THE INTERACTION OF LOCAL DECISION-MAKING AND STOCHASTIC STATE TRANSITIONS

The failure of the non-linear value-based MORL algorithms on the Space Traders MOMDP can be explained by the analysis of stochastic-transition MOMDPs previously carried out by Bryce et al. [2] in the context of probabilistic planning. This analysis has been largely overlooked by MORL researchers so far, and so one of the contributions of this paper is to bring this work to the attention of the MORL research community.

Figure 3 illustrates a simple MDP reproduced from Bryce et al. [2], with a stochastic branch occurring on the transition from the initial state. The table in the lower half of this figure specifies the mean return for the four possible deterministic policies. Keeping in mind that this MOMDP is phrased in terms of minimising cost

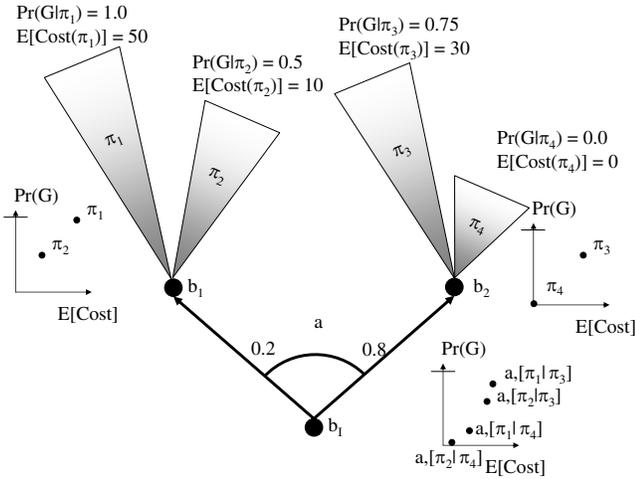


Figure 3: A sample probabilistic planning MOMDP, reproduced from Bryce et al. [2]. Executing action  $a$  from  $b_1$  leads to two branches with probability 0.2 and 0.8. At each of these branches a choice between two sub-plans with different payoffs exists. The aim for the planner is to identify the correct sub-plan to execute at each branch, so as to minimise cost while ensuring successful execution above a fixed probability.

(rather than maximising the inverse of the cost), it can be seen that unlike Space Traders, there are no Pareto-dominated policies for this MOMDP.<sup>1</sup>

The aim of the agent is to minimise the cost, subject to satisfying at least a 0.6 probability of success. Within an ESR formulation of the problem (i.e. ensure the probability of success threshold is achieved in each episode), the optimal policy is to select sub-plan  $\pi_1$  at branch  $b_1$  and  $\pi_3$  at branch  $b_2$  as both of these sub-plans individually satisfy the probability threshold. However if considered from the SER perspective, the optimal plan is to execute  $\pi_2$  at branch  $b_1$  and  $\pi_3$  at branch  $b_2$  – while  $\pi_2$  itself fails to achieve the probability threshold, this branch is executed with a low probability and so the mean outcome of the two sub-plans will achieve the threshold while also producing a significant cost saving.

As identified by [2], whether the overall policy meets the constraints depends on the probability with which each branch is executed as well as the mean outcome of each branch. Determining the correct sub-plan to follow at each branch requires consideration of the sub-plan options available at each other branch in combination with the probability of branch execution.

<sup>1</sup>While clearly illustrating the problem, this MOMDP also lacks the narrative drama of Space Traders!

This requirement is fundamentally incompatible with the localised decision-making at the heart of model-free value-based RL methods like Q-learning, where it is assumed that the correct choice of action can be determined purely based on information available to the agent at the current state. The provision of additional information such as the sum of rewards received so far in the episode as discussed in Section 3.2 is insufficient, as it still only provides information about the branch which has been followed in this episode, rather than all possible branches which might have been executed.

The conclusion to be drawn from both this example and Space Traders is that value-based model-free MORL methods are inherently limited when applied in the context of SER optimisation of non-linear utility on MOMDPs with non-deterministic state transitions. These methods may fail to discover the policy which maximises the SER (i.e. the mean utility over multiple episodes). To the best of our knowledge this limitation has not previously been identified in the MORL literature. It is particularly important as the combination of SER, stochastic state transitions and non-linear utility may well arise in important areas of application such as AI safety [20].

## 5 POTENTIAL SOLUTIONS

In this section we will briefly review and critique various options which may address the issue identified above.

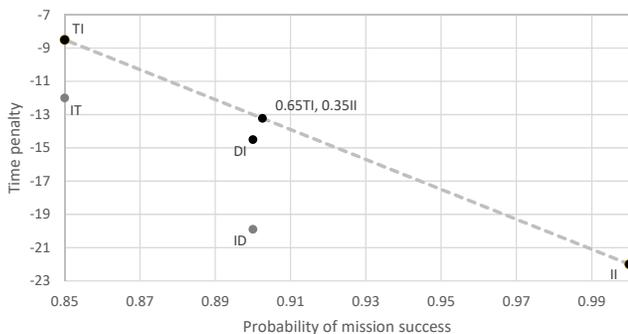
### 5.1 ESR Optimisation

As noted earlier the issue described arises due to the fact that an agent aiming to find a policy optimal with regards to SER must take into account the value which will be received on average by its policy across multiple episodes. Framing the problem in terms of ESR optimisation would eliminate this issue. However ESR is clearly inappropriate for the context of the Space Traders MOMDP. The agent will aim to ensure every episode meets the threshold for the mission-success objective. This can only be achieved by following the strictly safe II policy, which produces results which are far worse for the user’s true utility than the DI policy.

### 5.2 Non-stationary or non-deterministic policies

Previous work has demonstrated that for the SER formulation, or for non-episodic tasks, policies formed from a non-stationary or non-deterministic mixture of deterministic policies can Pareto-dominate deterministic policies [17, 21]. For example, a mixture which randomly selects between policies TI and II with appropriate probabilities at the start of each episode can produce a mean outcome which exceeds that of policy DI, as shown in Figure 4 – the mixture policy which selects TI with probability 0.65 and II with probability 0.35 achieves a mean return of (0.9025, -13.225) which is superior to the deterministic DI policy with regards to both objectives.

However the use of policies which vary so widely may not be appropriate in all contexts – for many problems the more consistent outcome produced by a deterministic policy may be preferable, and so methods to find SER-optimal deterministic policies for stochastic MOMDPs are still required.



**Figure 4: The mean return per episode for a mixture policy formed by selecting between the deterministic policies TI and II with probability 0.65 and 0.35 respectively Pareto-dominates the mean return of deterministic policy DI.**

### 5.3 Multi-policy value-based MORL

As well as the single-policy value-based MORL methods examined in this paper, several authors have proposed multi-policy methods. These operate by retaining multiple value vectors at each state. These can correspond to either all Pareto-optimal values obtainable from that state, or (for purposes of efficiency) be constrained to store only those values which can help construct the optimal value function under some assumptions about the nature of the overall utility function  $f$  [13]. Multi-policy algorithms were first proposed for variants of dynamic programming [26, 27] and more recently have been extended to MORL [14, 25].

By propagating back the coverage set of values available at each successor state, these algorithms would correctly identify all potentially optimal policies available at the starting state, and the optimal policy could then be selected at that point – in the context of Space Traders this would allow for the desired DI policy to be selected. However two issues still need to be addressed. One is ensuring that the agent has a means of determining which action should be performed in each encountered state to align with the initial choice of policy. Existing algorithms do not necessarily provide such a means in the context of stochastic transitions. Second, the existing multi-policy MORL algorithms do not have an obvious extension to complex state-spaces where tabular methods are infeasible. Conventional function-approximation methods can not be applied, as the cardinality of the vectors to be stored can vary between states. Vamplew et al. [19] provides preliminary work addressing this problem, but further work is still required to make this approach practical.

### 5.4 Model-based methods

As well as describing the difficulties faced by probabilistic planning, Bryce et al. [2] also propose a search algorithm known as Multiobjective Looping AO\* (MOLAO\*) to solve such tasks. As a planning method, this assumes an MOMDP with known state transition probabilities and a finite and tractable number of discrete states. It may be possible to extend this approach by integrating it within model-based RL algorithms which can learn to estimate the transition probabilities and to generalise across states. We are not

aware of any prior work which has attempted to do so. However the model-based MORL approach proposed in Wiering et al. [28] may provide a suitable basis for implementing a reinforcement learning equivalent of MOLAO\*.

### 5.5 Policy-search methods

An alternative to value-based approaches is to use policy-search approaches to RL. As these directly maximise the policy as a whole as defined by a set of policy parameters, they do not have the local decision-making issue faced by model-free value-based methods.

Multiple researchers have proposed and evaluated policy-search methods for multiobjective problems [7, 9, 15, 16]. One issue to be addressed however is that these methods most naturally produce stochastic policies and as such may have the same problems as faced by the mixture or non-stationary approaches discussed in Section 5.2, unless they are modified or constrained so as to ensure convergence to a deterministic policy.

## 6 CONCLUSION

We have described a stochastic MOMDP and utility function which, despite their seeming simplicity, are not amenable to solution by the widely-used model-free value-based approaches to MORL. While this issue with MOMDPs with stochastic state transitions has previously been described in the context of probabilistic planning [2], this is the first work to identify the implications for MORL. Our example also demonstrates that under stochastic state-transitions, it is in fact possible for such MORL methods to converge to a Pareto-dominated policy.

The combination of SER optimisation, stochastic state transitions and the need for a deterministic policy are likely to arise in a range of applications (particularly in risk-aware agents), and so awareness of the limitations of some MORL methods to work under these characteristics is important in order to avoid the use of inappropriate methods.

## REFERENCES

- [1] Leon Barrett and Srinivas Narayanan. 2008. Learning all optimal policies with multiple criteria. In *Proceedings of the 25th international conference on Machine learning*. 41–47.
- [2] Daniel Bryce, William Cushing, and Subbarao Kambhampati. 2007. Probabilistic planning is multi-objective. *Arizona State University, Tech. Rep. ASU-CSE-07-006* (2007).
- [3] A Castelletti, Stefano Galelli, Marcello Restelli, and Rodolfo Soncini-Sessa. 2010. Tree-based reinforcement learning for optimal water reservoir operation. *Water Resources Research* 46, 9 (2010).
- [4] Zoltán Gábor, Zsolt Kalmár, and Csaba Szepesvári. 1998. Multi-criteria reinforcement learning. In *ICML*, Vol. 98. 197–205.
- [5] Peter Geibel. 2006. Reinforcement learning for MDPs with constraints. In *European Conference on Machine Learning*. Springer, 646–653.
- [6] Rustan Issabekov and Peter Vamplew. 2012. An empirical comparison of two common multiobjective reinforcement learning algorithms. In *Australasian Joint Conference on Artificial Intelligence*. Springer, 626–636.
- [7] Simone Parisi, Matteo Pirodda, and Jan Peters. 2017. Manifold-based multi-objective policy search with sample reuse. *Neurocomputing* 263 (2017), 3–14.
- [8] Julien Perez, Cécile Germain-Renaud, Balázs Kégl, and Charles Loomis. 2009. Responsive elastic computing. In *Proceedings of the 6th international conference industry session on Grids meets autonomic computing*. 55–64.
- [9] Matteo Pirodda, Simone Parisi, and Marcello Restelli. 2015. Multi-objective reinforcement learning with continuous Pareto frontier approximation. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*.
- [10] Roxana Rădulescu, Patrick Mannion, Diederik M Roijers, and Ann Nowé. 2019. Equilibria in multi-objective games: A utility-based perspective. In *Proceedings of the adaptive and learning agents workshop (ALA-19) at AAMAS*.

- [11] Diederik M Roijers, Denis Steckelmacher, and Ann Nowé. 2018. Multi-objective reinforcement learning for the expected utility of the return. In *Adaptive Learning Agents (ALA) workshop at AAMAS*, Vol. 18.
- [12] Diederik M Roijers, Peter Vamplew, Shimon Whiteson, and Richard Dazeley. 2013. A survey of multi-objective sequential decision-making. *Journal of Artificial Intelligence Research* 48 (2013), 67–113.
- [13] Diederik M Roijers, Shimon Whiteson, and Frans A Oliehoek. 2013. Computing convex coverage sets for multi-objective coordination graphs. In *International Conference on Algorithmic Decision Theory*. Springer, 309–323.
- [14] Manuela Ruiz-Montiel, Lawrence Mandow, and José-Luis Pérez-de-la Cruz. 2017. A temporal difference method for multi-objective reinforcement learning. *Neurocomputing* 263 (2017), 15–25.
- [15] Christian Robert Shelton. 2001. Importance sampling for reinforcement learning with multiple objectives. (2001).
- [16] Eiji Uchibe and Kenji Doya. 2007. Constrained reinforcement learning from intrinsic and extrinsic rewards. In *2007 IEEE 6th International Conference on Development and Learning*. IEEE, 163–168.
- [17] Peter Vamplew, Richard Dazeley, Ewan Barker, and Andrei Kelarev. 2009. Constructing stochastic mixture policies for episodic multiobjective reinforcement learning tasks. In *Australasian joint conference on artificial intelligence*. Springer, 340–349.
- [18] Peter Vamplew, Richard Dazeley, Adam Berry, Rustam Issabekov, and Evan Dekker. 2011. Empirical evaluation methods for multiobjective reinforcement learning algorithms. *Machine learning* 84, 1-2 (2011), 51–80.
- [19] Peter Vamplew, Richard Dazeley, Cameron Foale, and Tanveer Choudhury. 2018. Non-functional regression: A new challenge for neural networks. *Neurocomputing* 314 (2018), 326–335.
- [20] Peter Vamplew, Richard Dazeley, Cameron Foale, Sally Firmin, and Jane Mumery. 2018. Human-aligned artificial intelligence is a multiobjective problem. *Ethics and Information Technology* 20, 1 (2018), 27–40.
- [21] Peter Vamplew, Rustam Issabekov, Richard Dazeley, Cameron Foale, Adam Berry, Tim Moore, and Douglas Creighton. 2017. Steering approaches to Pareto-optimal multiobjective reinforcement learning. *Neurocomputing* 263 (2017), 26–38.
- [22] Peter Vamplew, John Yearwood, Richard Dazeley, and Adam Berry. 2008. On the limitations of scalarisation for multi-objective reinforcement learning of pareto fronts. In *Australasian Joint Conference on Artificial Intelligence*. Springer, 372–378.
- [23] Kristof Van Moffaert, Madalina M Drugan, and Ann Nowé. 2013. Hypervolume-based multi-objective reinforcement learning. In *International Conference on Evolutionary Multi-Criterion Optimization*. Springer, 352–366.
- [24] Kristof Van Moffaert, Madalina M Drugan, and Ann Nowé. 2013. Scalarized multi-objective reinforcement learning: Novel design techniques. In *2013 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL)*. IEEE, 191–199.
- [25] Kristof Van Moffaert and Ann Nowé. 2014. Multi-objective reinforcement learning using sets of pareto dominating policies. *The Journal of Machine Learning Research* 15, 1 (2014), 3483–3512.
- [26] DJ White. 1982. Multi-objective infinite-horizon discounted Markov decision processes. *Journal of mathematical analysis and applications* 89, 2 (1982), 639–647.
- [27] Marco A Wiering and Edwin D De Jong. 2007. Computing optimal stationary policies for multi-objective markov decision processes. In *2007 IEEE International Symposium on Approximate Dynamic Programming and Reinforcement Learning*. IEEE, 158–165.
- [28] Marco A Wiering, Maikel Withagen, and Mădălina M Drugan. 2014. Model-based multi-objective reinforcement learning. In *2014 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL)*. IEEE, 1–6.
- [29] Luisa M Zintgraf, Timon V Kanters, Diederik M Roijers, Frans Oliehoek, and Philipp Beau. 2015. Quality assessment of MORL algorithms: A utility-based approach. In *Benelearn 2015: Proceedings of the 24th Annual Machine Learning Conference of Belgium and the Netherlands*.