

Utility-Based Reinforcement Learning: Unifying Single-objective and Multi-objective Reinforcement Learning

Blue Sky Ideas Track

Peter Vamplew
Federation University Australia
Ballarat, Australia
p.vamplew@federation.edu.au

Cameron Foale
Federation University Australia
Ballarat, Australia
c.foale@federation.edu.au

Conor F. Hayes
Lawrence Livermore National
Laboratory
Livermore, USA
hayes56@llnl.gov

Patrick Mannion
University of Galway
Galway, Ireland
patrick.mannion@universityofgalway.ie

Enda Howley
University of Galway
Galway, Ireland
enda.howley@universityofgalway.ie

Richard Dazeley
Deakin University
Geelong, Australia
r.dazeley@deakin.edu.au

Scott Johnson
Deakin University
Geelong, Australia
s223861809@deakin.edu.au

Johan Källström
Linköping University
Linköping, Sweden
johan.kallstrom@liu.se

Gabriel Ramos
Universidade do Vale do Rio dos Sinos
São Leopoldo, Brazil
gdoramos@unisinis.br

Roxana Rădulescu
Vrije Universiteit Brussels, Belgium
Utrecht University, The Netherlands
roxana.radulescu@vub.be

Willem Röpke
Vrije Universiteit Brussel
Brussels, Belgium
willem.ropke@vub.be

Diederik M. Roijers
Vrije Universiteit Brussel
Brussels, Belgium
diederik.roijers@vub.be

ABSTRACT

Research in multi-objective reinforcement learning (MORL) has introduced the *utility-based* paradigm, which makes use of both environmental rewards and a function that defines the utility derived by the user from those rewards. In this paper we extend this paradigm to the context of single-objective reinforcement learning (RL), and outline multiple potential benefits including the ability to perform multi-policy learning across tasks relating to uncertain objectives, risk-aware RL, discounting, and safe RL. We also examine the algorithmic implications of adopting a utility-based approach.

KEYWORDS

reinforcement learning; utility

ACM Reference Format:

Peter Vamplew, Cameron Foale, Conor F. Hayes, Patrick Mannion, Enda Howley, Richard Dazeley, Scott Johnson, Johan Källström, Gabriel Ramos, Roxana Rădulescu, Willem Röpke, and Diederik M. Roijers. 2024. Utility-Based Reinforcement Learning: Unifying Single-objective and Multi-objective Reinforcement Learning: Blue Sky Ideas Track. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024)*, Auckland, New Zealand, May 6 – 10, 2024, IFAAMAS, 5 pages.



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand. © 2024 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

1 INTRODUCTION

Multi-objective reinforcement learning (MORL) has emerged as an important sub-field of reinforcement learning (RL) [13, 24]. So far knowledge has flowed primarily from single-objective RL (SORL) into MORL, with SORL algorithmic innovations being adapted to the context of multiple objectives [2, 6, 22, 34]. This paper counters that trend, as we argue that the utility-based paradigm widely used in MORL [5, 13, 21], has relevance and benefits to SORL. We present a general framework for utility-based RL (UBRL), unifying SORL and MORL, and discuss benefits for single-objective problems – in particular the potential for *multi-policy learning* which allows greater flexibility and decision-maker control over the behaviour of agents. We also highlight lessons from MORL research regarding the algorithmic implications of a utility-based approach as a guide to future single-objective UBRL research.

2 FORMALISING UTILITY-BASED RL

2.1 MDPs, MOMDPs and Optimisation criteria

SORL and MORL share a common foundation in the assumption that the environment can be represented as some form of Markov Decision Process (MDP). In the single objective case, a MDP is represented by the tuple $\langle S, A, T, \gamma, \mu, R \rangle$, where:

- S is the state space and A is the action space
- $T: S \times A \times S \rightarrow [0, 1]$ is a probabilistic transition function
- $\gamma \in [0, 1)$ is a discount factor
- $\mu: S \rightarrow [0, 1]$ is a probability distribution over initial states
- $R: S \times A \times S \rightarrow \mathbb{R}$ is a scalar-valued reward function

For a single-objective problem of this nature, the agent’s aim is to discover a policy π which maximises either the finite-horizon undiscounted reward (Equation 1), the infinite-horizon cumulative discounted reward (Equation 2), or the average reward (Equation 3). We note that the choice of criteria to optimise is made by the system designer, rather than being an inherent property of the MDP itself.

$$V_\pi = \mathbb{E} \left[\sum_{i=0}^{N-1} r_i \mid \pi, \mu \right] \quad (1)$$

$$V_\pi = \mathbb{E} \left[\sum_{i=0}^{\infty} \gamma^i r_i \mid \pi, \mu \right] \quad (2)$$

$$V_\pi = \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E} \left[\sum_{i=0}^{N-1} r_i \mid \pi, \mu \right] \quad (3)$$

MORL differs from SORL in being based on the concept of a Multi-Objective Markov Decision Process (MOMDP). This is identical to a MDP in all respects other than the reward, where rather than a scalar reward $R: S \times A \times S \rightarrow \mathbb{R}$, the MOMDP has a vector-valued reward function $\mathbf{R}: S \times A \times S \rightarrow \mathbb{R}^d$ which specifies the immediate reward for each of the d objectives.

The optimisation criteria in Equations 1–3 are insufficient for a MOMDP, as they only provide a partial-ordering over the vector values. Therefore the utility-based paradigm for MORL assumes the existence of a *utility function* (sometimes called a *scalarisation function*) u , which maps the multi-objective value of a policy to a scalar value [13]. Note that the details of this utility function may or may not be available in advance – in the latter case the agent may need to identify a set of policies which would be optimal under different parameterisations of the utility function. We will denote a parameterised utility function as u_ω .

Given the existence of u we can define two alternative optimisation criteria, depending on the stage at which u is applied (the reasoning for this distinction will be discussed further in Section 5)¹. Equation 4 defines the scalarised expected reward (SER) where u is applied to the expected vector reward, whereas Equation 5 defines the expected scalarised reward, where u is applied inside the expectation operator. In either case, these equations provide a scalar value which defines a total ordering over policies.

$$V_u^\pi = u \left(\mathbb{E} \left[\sum_{i=0}^{\infty} \gamma^i \mathbf{r}_i \mid \pi, \mu \right] \right) \quad (4)$$

$$V_u^\pi = \mathbb{E} \left[u \left(\sum_{i=0}^{\infty} \gamma^i \mathbf{r}_i \right) \mid \pi, \mu \right] \quad (5)$$

2.2 Utility-based RL as a general framework

In the previous subsection we highlighted the differences between conventional single-objective RL and utility-based MORL. We now consider the commonalities between these approaches, and argue

¹Note that these equations relate to the infinite-horizon discounted return which has been the most widely studied criteria in the MORL literature. Similar utility-based equations can readily be defined for the finite-horizon and average-reward criteria.

that utility-based RL provides a general framework which encompasses both SORL and MORL².

We start with the straightforward observation that MDPs are a subclass of MOMDPs. Any MDP can be mapped to an equivalent MOMDP with a reward function which consists of a one-dimensional vector $\mathbf{R} = [R]$.

Having done this conversion, then utility-based methods such as those defined in terms of the optimisation criteria in Equations 4 and 5 can be applied to the resultant MOMDP. More generally we might consider any utility function which maps from the expected mean rewards or the expected distribution of rewards to a scalar value – we will explore examples of this in Sections 4.2 and 4.3. In addition the utility function may itself be parameterised, so that a set of policies would be required in order to optimise the criteria across all possible instantiations of u .

We note that some authors have previously criticised approaches such as MORL and risk-aware RL as “special cases” [25], and we anticipate that similar criticisms may be levelled at UBRL. To the contrary, we argue that the UBRL framework presented here is a strict generalisation of prior RL frameworks. It can encompass standard scalar-reward RL by simply setting the number of objectives $n=1$ and the utility function u to the identity function, while also supporting other more complex forms of RL, such as MORL.

3 MOTIVATION FOR UTILITY-BASED SORL

While it is clear that SORL problems can be represented within the UBRL framework, the question remains as to whether this is an unnecessary complication, or whether it provides tangible benefits.

One advantage of a utility-based paradigm is that it can simplify the task of reward engineering [8] – that is, designing rewards in such a way that maximising them induces the desired behaviour from an RL agent. In many cases it may be easy to identify events (significant changes in state) within the MDP for which rewards should be provided (e.g. picking up an object). But the task of tailoring the magnitude of those rewards and the choice of optimisation criteria to engender the desired behaviour is non-trivial (see for example [16]). Consider an environment with gold nuggets of differing size at varying distances from the starting location. Whether an agent prioritises collecting nearby low-value objects over more distant, higher-valued objects will depend both on the relative magnitude of the rewards and factors of the optimisation criteria, such as the choice of discounting parameter. The utility-based framework separates the specification of the rewards within the MDP (i.e. defining the environment) from the definition of utility (i.e. defining the desired outcome), which may make the reward engineer’s task easier.

This is particularly true considering that MORL research has found that a key benefit of a utility-based approach is that it allows for the creation of multi-policy algorithms [13]. Rather than aiming to learn a single policy optimal for a specific definition of utility, multi-policy algorithms instead consider a set of possible utility functions $U = \{U_1, \dots, U_m\}$, and produce a *coverage set* of policies such that the agent has an optimal policy for every member of U .

²The name ‘utility-based’ is not intended to imply that standard SORL is *not* maximising a measure of utility, but to emphasise our framework’s *explicit* consideration of user-utility. In standard SORL the choice of optimisation criteria is made by the developer whereas in multi-policy UBRL the final decision rests with the end-user.

Learning multiple policies in this way is not possible within conventional SORL, and so represents a novel contribution of the UBRL framework that greatly enhances the RL process. We can define an MDP with a simple reward structure, learn multiple policies for that MDP based on a variety of different definitions of utility, and then select the utility function which produces the most desirable outcomes. This will be significantly easier than trying to specify *a priori* the rewards required to produce the desired behaviour. It also means that the final decision can be left in the hands of the ultimate user of the system, rather than inappropriately requiring it to be made by the system engineers [32]. In addition this provides flexibility should the desired behaviour change over time, as a new policy can be selected without any need for further learning.

The benefits of multi-policy learning may also be realised at relatively little additional cost. The utility-based approach allows for *inner-loop* multi-policy methods, in which the multiple policies are learned in parallel [13]. For example, the Conditioned Network algorithm trains a single Deep Q-Network, conditioned on both the current MOMDP state and the values of the parameters ω of the utility function. Experiences gained while following the policy for one value of ω are leveraged to update Q-values for policies conditioned on different values of ω . This greatly increases the sample-efficiency of these methods.

4 POTENTIAL SINGLE-OBJECTIVE APPLICATIONS OF UTILITY-BASED RL

In this section we will discuss several potential applications of the utility-based paradigm within SORL. These are intended as representative examples where UBRL provides benefits, rather than an exhaustive list. These represent both a reframing of existing concepts such as risk-aware RL to the UBRL framework, as well as more novel concepts made possible by the adoption of this framework.

4.1 Multi-policy methods for hard-to-define objectives

Consider the scenario of a mining company which has the choice between carrying out its usual operations or following a riskier course of action that can potentially lead to a lot of the resource being mined (e.g., committing most of the workforce to exploratory excavation). The company has outstanding contracts that commit it to sell a given amount to certain costumers at a set price and leading to an incurred penalty if these amounts are not delivered.

There are various ways to address this using standard RL:

- We might model this as constrained RL [10], where the fulfillment of contracts is an inviolable constraint. But perhaps this is overly restrictive, as the management may be prepared to breach contracts if the potential pay-off is sufficiently high?
- It could be modelled as a SORL problem, with a positive reward for the value of the resource, and a negative penalty for unfulfilled contracts. The agent might select the riskier action if the potential gain offsets the contractual penalty. But such an approach might fail to account for longer-term reputational damage arising from unfulfilled contracts.
- MORL methods could model reputational harm as a further objective, but this may be difficult to define quantitatively.

As a result, the preferred approach of the company’s management might be to generate a set of alternative policies using UBRL, and make an *a posteriori* decision about which policy to follow. This could be accomplished by defining the reward in terms of the quantity of the resource mined, and then specifying the utility function as in Equation 6, where d is the monetary value of each unit of the resource, $c(\cdot)$ is a binary function returning 1 if the contract terms are breached and 0 otherwise, p is the penalty for breaching the contract, and h is an estimate of the financial impact of the reputational harm caused by a contract breach.

$$u_h = d * \sum_{i=0}^{N-1} r_i - c\left(\sum_{i=0}^{N-1} r_i\right) * (p + h) \quad (6)$$

Note that the utility term u is conditioned on the value of h . Rather than learning an optimal policy for a single estimate of cost of the reputational harm, this utility-based formulation enables the UBRL agent to learn multiple policies, each optimal with respect to a different estimated value of h . The multiple policies produced via this process would then be presented to the management, allowing them to make an informed decision about the best policy to execute. Existing multi-policy MORL methods such as the Conditioned Network approach of [1] could readily be adapted to this task by conditioning Q-values and policies on h .

4.2 Multi-policy risk-aware RL

Many authors have investigated risk-aware RL agents – that is, agents which apply some awareness of risk during their decision-making rather than simply aiming to maximise the expected return. A common approach to address this problem is to use a distributional form of RL [3, 15, 17, 19], which learns the distribution of future returns from any state rather than just the mean expected return. These approaches can take into account various aspects of this distribution when identifying the optimal policy (e.g. best worst case, more than 60% chance of profit, etc.) [7, 9, 11, 14, 33].

Again we would argue that this approach is already utility-based, as the choice of constraints applied to the distribution represents the definition of the utility function. Conventional SORL approaches to risk-awareness assume that these constraints are fixed *a priori*, and learn a single policy which is optimal under those constraints. In contrast a UBRL approach would learn, in parallel using multi-policy methods, a set of policies which are optimal under different risk-sensitivity preferences. These could then be presented to a human decision-maker for selection of the policy to be executed.

For example, consider the conditional value at risk (CVaR) (Equation 7), which is widely-used as an optimisation criteria in risk-aware RL [26]. Z is the distribution of future returns, and $\alpha \in [0, 1]$ is a parameter controlling the risk-sensitivity of the agent. In conventional risk-aware RL, the value of α is fixed and a single optimal policy is found. If we instead treat $CVaR_\alpha$ as a utility function parameterised by α , we can apply multi-policy methods to find policies with a diverse range of different sensitivities to risk, and then allow a human decision-maker to select their preferred policy.

$$CVaR_\alpha(Z) = \mathbb{E}[Z | Z \leq VaR_\alpha(Z)] \quad (7)$$

4.3 Multi-policy discounting

We noted earlier in Section 2 that the choice of optimisation criteria used in SORL (undiscounted sum, discounted sum, or average reward) is made by the system designer/user rather than being an inherent property of the MDP itself. As such, this decision could in itself be regarded as a definition of user utility, although it usually is not framed as such. Similarly we would contend that the discounting term γ used in Equation 2 can be viewed as a form of utility definition. This term does not originate from the MDP itself, and two agents using different discounting rates may well derive differing optimal policies from the same MDP. Conventional SORL requires that the discounting rate be fixed prior to learning, and the impact of this choice on the final policy is not evident to the user. The only way to gain insight into the effect of the discounting rate on the agent’s behaviour is to subsequently train another agent on the same MDP using a different value for γ which is inefficient.

Alternatively we could parameterise the utility function on γ as in Equation 8, and use multi-policy UBRL to simultaneously learn optimal policies for multiple values of γ . These can then be presented to the user, allowing them to make a fully informed choice of the appropriate policy³

$$u_\gamma = \sum_{i=0}^{\infty} \gamma^i r_i \quad (8)$$

4.4 Satisficing agents

AI safety researchers have argued that ‘hard optimizers’ like conventional RL can be unsafe, as their focus on maximising their reward may lead to adverse side-effects if the rewards are not fully aligned with the desired behaviour [28, 29]. This has led to interest in *satisficing agents* which are not incentivised to over-optimize [12]. For example, in a multi-agent resource-gathering task we might disincentivise an agent from collecting more resources than it actually needs so as to avoid adversely impacting others. In the context of MORL satisficing has been implemented by causing the agent to switch emphasis between objectives once a suitable level of return has been achieved for each objective. This can either use hard thresholding as in [31] or a non-linear utility which reduces the weight of objectives as higher returns are achieved [27].

These approaches are not directly applicable to SORL due to the lack of alternative objectives. For example, reducing the gradient of the utility function as in [27] would not disincentivise the agent from still trying to maximise it. So for safe SORL we might need to consider the concept of a non-monotonic utility function where the utility actually falls in value after a satisfactory amount of reward has been received, as shown in Equation 9.

$$u_\omega = -|\omega - \sum_{i=0}^{\infty} \gamma^i r_i| \quad (9)$$

This could of course also be achieved via changing the reward definition within the MDP. However that approach would not be amenable to multi-policy learning whereas a multi-policy UBRL algorithm could simultaneously learn optimal policies for various

³We note that the utility definition in Equation 8 differs from both forms of utility previously used in MORL (Equations 4 and 5), in that it is the per-time-step reward, rather than the summed return, that is being transformed.

values of the threshold ω , allowing insight into the appropriate value of ω which will produce the desired safe behaviour while still performing to a satisfactory level in terms of the actual task.

5 IMPLICATIONS OF NON-LINEAR UTILITY

The benefits of utility-based learning described earlier do not come without some additional considerations. Research in utility-based MORL has identified several issues which need to be taken into account when designing and applying UBRL methods, particularly when using non-linear forms of u , such as that in Equation 9.⁴

One fundamental issue relates to the selection of the optimisation criteria. For tasks where we care about the outcome on a per-episode basis then the Expected Scalarised Return (ESR) defined in Equation 5 is most appropriate, whereas for problems where we care about the average outcome over multiple episodes then Scalarised Expected Return (SER, Equation 4) is the correct criteria⁵. Several researchers have reported that different algorithms may be required to address the ESR and SER settings [20, 23, 30].

A non-linear utility function has further implications for value-based UBRL algorithms, as it means the returns are no longer additive which is not compatible with the Bellman equation [24]. To address this, UBRL algorithms based on temporal difference approaches may need to use an *augmented state* which concatenates the environment state with a history of prior received rewards [30].

One further issue to consider is the implementation of reward shaping. In SORL, shaping rewards are simply added to the MDP rewards prior to presenting the reward to the agent. As long as certain conditions are met, the optimal policy is not altered by this process [18]. However this may not be the case if a non-linear utility function is applied to these combined reward values. To allow the agent to appropriately evaluate the true utility, it may prove necessary to treat the shaping reward as a separate objective [4].

6 CONCLUSION

The utility-based RL framework presented here unifies the previously disparate areas of single-objective RL and multi-objective RL. The utility-based approach increases the flexibility of SORL agents by facilitating multi-policy learning. This is the main novel capability of UBRL compared to standard SORL, and increases the control human decision-makers have over the RL agent, by allowing them to make an informed selection of their preferred agent after seeing which alternative behaviours exist, rather than making *a priori* decisions regarding the reward design and optimisation criteria and hoping these achieve the desired outcome. UBRL may also have benefits in terms of simplifying the task of reward design. We have presented several motivating cases where multi-policy UBRL approaches have advantages over standard SORL – uncertain objectives, risk-aware RL, discounted returns, and satisficing agents. Finally we believe that wide-spread adoption of UBRL as a unified framework would facilitate faster and easier transfer of novel ideas, algorithms and software implementations between the previously somewhat disconnected fields of SORL and MORL.

⁴MORL may use either linear or non-linear utility functions, but single-objective UBRL will almost exclusively use non-linear forms as a linear function would merely scale the range of the scalar reward without changing the optimal policy.

⁵The distinction between ESR and SER does not arise in SORL or in UBRL with linear U as the values of Equations 4 and 5 are equal under those settings.

REFERENCES

- [1] Axel Abels, Diederik Roijers, Tom Lenaerts, Ann Nowé, and Denis Steckelmacher. 2019. Dynamic weights in multi-objective deep reinforcement learning. In *International conference on machine learning*. PMLR, 11–20.
- [2] Lucas Nunes Alegre, Ana Bazzan, and Bruno C Da Silva. 2022. Optimistic linear support and successor features as a basis for optimal policy transfer. In *International Conference on Machine Learning*. PMLR, 394–413.
- [3] Marc G Bellemare, Will Dabney, and Mark Rowland. 2023. *Distributional reinforcement learning*. MIT Press.
- [4] Tim Brys, Anna Harutyunyan, Peter Vrancx, Matthew E Taylor, Daniel Kudenko, and Ann Nowé. 2014. Multi-objectivization of reinforcement learning problems by reward shaping. In *2014 international joint conference on neural networks (IJCNN)*. IEEE, 2315–2322.
- [5] Xin-Qiang Cai, Pushi Zhang, Li Zhao, Jiang Bian, Masashi Sugiyama, and Ashley Juan Llorens. 2023. Distributional Pareto-Optimal Multi-Objective Reinforcement Learning. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- [6] Xi Chen, Ali Ghadirzadeh, Márten Björkman, and Patric Jensfelt. 2019. Meta-learning for multi-objective reinforcement learning. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 977–983.
- [7] Ziteng Cheng, Sebastian Jaimungal, and Nick Martin. 2023. Distributional Method for Risk Averse Reinforcement Learning. *arXiv preprint arXiv:2302.14109* (2023).
- [8] Daniel Dewey. 2014. Reinforcement learning and the reward engineering principle. In *2014 AAAI Spring Symposium Series*.
- [9] Alhussein Fawzi, Matej Balog, Aja Huang, Thomas Hubert, Bernardino Romera-Paredes, Mohammadamin Barekatain, Alexander Novikov, Francisco J R Ruiz, Julian Schrittwieser, Grzegorz Swirszcz, et al. 2022. Discovering faster matrix multiplication algorithms with reinforcement learning. *Nature* 610, 7930 (2022), 47–53.
- [10] Ather Gattami, Qinbo Bai, and Vaneet Aggarwal. 2021. Reinforcement learning for constrained markov decision processes. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 2656–2664.
- [11] Ido Greenberg, Yinlam Chow, Mohammad Ghavamzadeh, and Shie Mannor. 2022. Efficient risk-averse reinforcement learning. *Advances in Neural Information Processing Systems* 35 (2022), 32639–32652.
- [12] Hossein Hajiabohassan and Ronald Ortner. 2023. Online Regret Bounds for Satisficing in MDPs. In *Sixteenth European Workshop on Reinforcement Learning*.
- [13] Conor F Hayes, Roxana Rădulescu, Eugenio Bargiacchi, Johan Källström, Matthew Macfarlane, Mathieu Reymond, Timothy Verstraeten, Luisa M Zintgraf, Richard Dazeley, Fredrik Heintz, et al. 2022. A practical guide to multi-objective reinforcement learning and planning. *Autonomous Agents and Multi-Agent Systems* 36, 1 (2022), 26.
- [14] Conor F Hayes, Mathieu Reymond, Diederik M Roijers, Enda Howley, and Patrick Mannion. 2023. Monte Carlo tree search algorithms for risk-aware and multi-objective reinforcement learning. *Autonomous Agents and Multi-Agent Systems* 37, 2 (2023), 26.
- [15] Conor F Hayes, Timothy Verstraeten, Diederik M Roijers, Enda Howley, and Patrick Mannion. 2022. Expected scalarised returns dominance: a new solution concept for multi-objective decision making. *Neural Computing and Applications* (2022), 1–21.
- [16] W Bradley Knox, Alessandro Allievi, Holger Banzhaf, Felix Schmitt, and Peter Stone. 2023. Reward (mis) design for autonomous driving. *Artificial Intelligence* 316 (2023), 103829.
- [17] John Martin, Michal Lyskawinski, Xiaohu Li, and Brendan Englot. 2020. Stochastically dominant distributional reinforcement learning. In *International conference on machine learning*. PMLR, 6745–6754.
- [18] Andrew Y Ng, Daishi Harada, and Stuart Russell. 1999. Policy invariance under reward transformations: Theory and application to reward shaping. In *ICML*, Vol. 99. Citeseer, 278–287.
- [19] Thanh Nguyen-Tang, Sunil Gupta, and Svetha Venkatesh. 2021. Distributional reinforcement learning via moment matching. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 9144–9152.
- [20] Roxana Rădulescu. 2020. A Utility-Based Perspective on Multi-Objective Multi-Agent Decision Making. In *Proceedings of the 19th International Conference on Autonomous Agents and Multi-Agent Systems*. Auckland, New Zealand: AAMAS, 2207–2208.
- [21] Roxana Rădulescu, Patrick Mannion, Diederik M Roijers, and Ann Nowé. 2020. Multi-objective multi-agent decision making: a utility-based analysis and survey. *Autonomous Agents and Multi-Agent Systems* 34, 1 (2020), 10.
- [22] Mathieu Reymond, Conor F Hayes, Denis Steckelmacher, Diederik M Roijers, and Ann Nowé. 2023. Actor-critic multi-objective reinforcement learning for non-linear utility functions. *Autonomous Agents and Multi-Agent Systems* 37, 2 (2023), 23.
- [23] Diederik M Roijers, Denis Steckelmacher, and Ann Nowé. 2018. Multi-objective reinforcement learning for the expected utility of the return. In *Proceedings of the Adaptive and Learning Agents workshop at FAIM*, Vol. 2018.
- [24] 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.
- [25] David Silver, Satinder Singh, Doina Precup, and Richard S Sutton. 2021. Reward Is Enough. *Artificial Intelligence* (2021), 103535.
- [26] Rahul Singh, Qinsheng Zhang, and Yongxin Chen. 2020. Improving robustness via risk averse distributional reinforcement learning. In *Learning for Dynamics and Control*. PMLR, 958–968.
- [27] Benjamin J Smith, Robert Klassert, and Roland Pihlakas. 2023. Using soft maximin for risk averse multi-objective decision-making. *Autonomous Agents and Multi-Agent Systems* 37, 1 (2023), 11.
- [28] Jessica Taylor. 2016. Quantilizers: A Safer Alternative to Maximizers for Limited Optimization. In *AAAI Workshop: AI, Ethics, and Society*.
- [29] 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 (2018), 27–40.
- [30] Peter Vamplew, Cameron Foale, and Richard Dazeley. 2022. The impact of environmental stochasticity on value-based multiobjective reinforcement learning. *Neural Computing and Applications* (2022), 1–17.
- [31] Peter Vamplew, Cameron Foale, Richard Dazeley, and Adam Bignold. 2021. Potential-based multiobjective reinforcement learning approaches to low-impact agents for AI safety. *Engineering Applications of Artificial Intelligence* 100 (2021), 104186.
- [32] Peter Vamplew, Benjamin J Smith, Johan Källström, Gabriel Ramos, Roxana Rădulescu, Diederik M Roijers, Conor F Hayes, Fredrik Heintz, Patrick Mannion, Pieter JK Libin, et al. 2022. Scalar reward is not enough: A response to Silver, Singh, Precup and Sutton (2021). *Autonomous Agents and Multi-Agent Systems* 36, 2 (2022), 41.
- [33] Qisong Yang, Thiago D Simão, Simon H Tindemans, and Matthijs TJ Spaan. 2023. Safety-constrained reinforcement learning with a distributional safety critic. *Machine Learning* 112, 3 (2023), 859–887.
- [34] Runzhe Yang, Xingyuan Sun, and Karthik Narasimhan. 2019. A generalized algorithm for multi-objective reinforcement learning and policy adaptation. *Advances in neural information processing systems* 32 (2019).