

Welfare and Fairness in Multi-objective Reinforcement Learning

Zimeng Fan*
Duke University
Durham, NC, USA
zf59@duke.edu

Muhang Tian*
Duke University
Durham, NC, USA
muhang.tian@duke.edu

Nianli Peng*
Duke University
Durham, NC, USA
nianli.peng@duke.edu

Brandon Fain
Duke University
Durham, NC, USA
brandon.fain@duke.edu

ABSTRACT

We study fair multi-objective reinforcement learning in which an agent must learn a policy that simultaneously achieves high reward on multiple dimensions of a vector-valued reward. Motivated by the fair resource allocation literature, we model this as an expected welfare maximization problem, for some non-linear fair welfare function of the vector of long-term cumulative rewards. One canonical example of such a function is the Nash Social Welfare, or geometric mean, the log transform of which is also known as the Proportional Fairness objective. We show that even approximately optimal optimization of the expected Nash Social Welfare is computationally intractable even in the tabular case. Nevertheless, we provide a novel adaptation of Q-learning that combines non-linear scalarized learning updates and non-stationary action selection to learn effective policies for optimizing nonlinear welfare functions. We show that our algorithm is provably convergent, and we demonstrate experimentally that our approach outperforms techniques based on linear scalarization, mixtures of optimal linear scalarizations, or stationary action selection for the Nash Social Welfare Objective.

CCS CONCEPTS

• **Computing methodologies** → **Reinforcement learning.**

KEYWORDS

Multi-objective Reinforcement Learning; Algorithmic Fairness

ACM Reference Format:

Zimeng Fan*, Nianli Peng*, Muhang Tian*, and Brandon Fain. 2023. Welfare and Fairness in Multi-objective Reinforcement Learning. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 9 pages.

1 INTRODUCTION

Suppose a logistics company announces they will deploy a reinforcement learning agent that optimizes for completed deliveries. It works and the number of deliveries increase. Some days later, the company begins receiving complaints that delivery service to

*These authors contributed equally to this work, ordered alphabetically by last name.

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

some locations has actually gotten much worse than before the AI deployment. The company assures its customers that the AI is learning and things will improve. But in another week, the situation is the same. Desperate, the company’s engineers boost the reward weights associated with deliveries to those locations, only to find that now other locations are being neglected.

This is an example where data-driven algorithmic systems may be generally quite performant but nonetheless fail on structured subsets of input. This is well-known in reinforcement learning, where extensive “reward shaping” is sometimes necessary to achieve desired behavior. In the opening example, the desired behavior is a policy that achieves high delivery service rates at *all* customer locations. However, standard reinforcement learning, in which the reward signal is a scalar value and the goal is to maximize total discounted reward, might naturally learn a policy that prioritizes “easy” to optimize regions (perhaps clusters of many tightly packed locations with many deliveries) at the expense of more difficult ways to achieve reward. Furthermore, because standard techniques rely on learning a *stationary* policy, the policy continues to prioritize the same customers day after day. Addressing this problem may require problem-specific fine tuning of rewards, and even then “fixing” the original problem can introduce new ones.

In this paper, we take a different approach. We study nonlinear welfare optimization in the context of Multi-objective Reinforcement Learning (MORL). A Multi-Objective Markov Decision Process (MOMDP) is a Markov Decision Process where rewards are vectors instead of scalars. The components of this vector can be viewed either as different criteria like cost and time or, as we interpret them, as individual utilities of “users” to whom the learning agent should be *fair*. In the opening example, the customers are the users and the vector reward tracks how well the learning agent is doing at optimizing deliveries for each user separately.

The solution to a MOMDP is a policy that seeks to maximize some function of the cumulative reward vector. Due to the linearity of expectation, linear functions maximizing some weighted arithmetic mean of the cumulative reward vector are the simplest to use. However, for any particular selection of weights, the resulting policies may be undesirable from a fairness perspective as any linear function may ignore the utility of some users. For instance, for equal weights, a policy that gives user 1 a utility of 10 and user 2 a utility of 0 is preferred over another policy that yields a utility of 4 for both. We therefore study a more general class of *welfare* functions with a particular emphasis on *nonlinear* welfare functions that optimize for fairness and efficiency.

Optimizing a nonlinear welfare function in a MOMDP is a substantial algorithmic challenge. The Bellman optimality principles [6][27] no longer hold, and stationary policies, in which the action selection depends only on the current state and not history, are no longer necessarily optimal. This is quite intuitive for fairness. For example, if an AI personal assistant is tasked with grocery shopping for a household with competing preferences over desserts, it is relevant to the current decision whether one member of the household got their most or least favorite dessert in previous weeks, as the agent may wish to be fair to its users across time.

Though these examples are toys, there are many real-world decision problems in which a learning agent may need to simultaneously prioritize more than a single utility or goal in a balanced and fair way. In telecommunications and wireless networking, one may want to allocate bandwidth in a way that balances the quality of service in many different locations. In autonomous driving, one may want to balance vehicle speed and passenger comfort [14].

1.1 Contributions and Outline

Sections 2 and 3 introduce related work and preliminaries for MOMDPs. Our results are as follows:

- (1) In Section 4 we introduce and characterize the problem of optimizing expected welfare for fairness in a MOMDP. We specifically focus on nonlinear welfare functions, with the Nash Social Welfare (NSW) as our canonical example of a fair welfare function.
- (2) In Section 5 we give a reduction in Theorem 5.2 to show that optimizing expected NSW is computationally intractable, even in the tabular setting. We further show that stationary policies cannot, in general, guarantee high approximation to optimality as the number of dimensions of reward grows.
- (3) On the positive side, also in Section 5 we define Algorithm 1 *Welfare Q-Learning* that adapts model-free Q-learning in two important ways to optimize non-linear welfare functions: (1) nonlinear learning updates, and (2) non-stationary action selection. We show in Theorem 5.3 that our algorithm is provably convergent.
- (4) In Section 6 we deploy our algorithm in two simulated environments to optimize expected NSW. Our algorithm substantially outperforms the following baselines: (1) optimal [for NSW] linear scalarization, (2) optimal [for NSW] mixtures of optimal policies in each dimension, and (3) stationary action selection on our algorithm’s learned Q-table.

2 RELATED WORK

Multi-objective reinforcement learning (MORL) algorithms include *single-policy* and *multi-policy* methods [18]. Single-policy methods use a scalarization function to reduce the problem to scalar optimization for a single policy. The simplest form is linear scalarization, applying a weighted sum on the Q vector [20].

Multi-policy methods search for a set of policies that approximate the Pareto frontier of the problem. For instance, the *convex hull value-iteration algorithm* [5] computes the deterministic stationary policies on the convex hull of the Pareto front. *Pareto Q learning* [20] integrates temporal difference algorithms with Pareto dominance

relations to learn a set of Pareto dominating policies. *Stochastic mixture policy* [28] combines multiple deterministic base policies with a convex combination, choosing a base policy with a given probability at the start of each episode. We focus on single-policy methods with nonlinear scalarization, as the size of the Pareto frontier may grow exponentially with the dimensionality of the problem, and because the Pareto frontier may not be well-approximated by its convex hull for nonlinear welfare functions.

Fairness in Reinforcement Learning has been recently considered, beginning with [15] in a scalar setting. More directly related to our work, [25] investigated the (Deep) MORL problem of learning a fair policy to optimize the *Generalized Gini Social Welfare function* using nonlinear scalarization. [1] studied a similar problem and considered maximizing concave welfare functions generally and Nash welfare specifically, showing an optimal approach to optimizing the welfare of expected rewards in the tabular setting and an extension to the function approximation setting.

Our work differs from these in two major ways: (1) we seek to optimize the expected welfare, rather than welfare of expected rewards (see Section 4) which is fundamentally more challenging computationally (see Section 5), and (2) we learn a non-stationary policy, as stationary policies may be far from optimal for optimizing expected (nonlinear) welfare (see Section 5 and [23]).

We formulate our consideration of welfare functions in Section 4 based on consideration from the resource allocation literature [21]. Our canonical example of a fair welfare function, the Nash Social Welfare (NSW) derives from Nash’s solution to the bargaining game [22] and its n-player extension [19]. Its log transform is commonly known as the proportional fairness objective. More recent studies have shown NSW maximization provides outstanding fairness guarantees when allocating both divisible and indivisible goods [8].

3 PRELIMINARIES

Multi-objective Markov Decision Process. A Multi-objective Markov Decision Process (MOMDP) consists of a finite set \mathcal{S} of states, a starting state $s_1 \in \mathcal{S}$,¹ a finite set \mathcal{A} of actions (we let $\mathcal{A}(s)$ denote the subset of actions available in state s), and probabilities $\mathcal{P}_{a,s,s'} \in [0, 1]$ that determine the probability of transitioning to state s' from state s after taking action a . Probabilities are normalized so that $\sum_{a \in \mathcal{A}(s), s'} \mathcal{P}_{a,s,s'} = 1$ for all s . We also have a reward function $R(s, a) : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^n$ for taking action a in state s .² Some states may be *terminal*, meaning they transition only to themselves and yield $\mathbf{0}$ reward.

Each of the n dimensions of the reward vector correspond to one of the multiple objectives that are to be maximized. At each time step t , the agent observes state $s_t \in \mathcal{S}$, takes action $a_t \in \mathcal{A}(s_t)$, and receives a reward vector $\mathbf{r}_t = R(s_t, a_t) \in \mathbb{R}^n$. The environment, in turn, transitions into s_{t+1} with probability $\mathcal{P}_{a_t, s_t, s_{t+1}}$. Where clear from context, we will often omit the subscript and simply write the immediate reward vector as \mathbf{r} .

A *trajectory* is a sequence of state, action, reward tuples $\tau = (s_1, a_1, \mathbf{r}_1), (s_2, a_2, \mathbf{r}_2), \dots, (s_T, a_T, \mathbf{r}_T)$. A trajectory that begins in the starting state s_1 and ends in a terminal state defines an *episode*. For

¹In general we may have a distribution over starting states; we assume a single starting state for ease of exposition.

²For simplicity of exposition, we assume rewards are deterministic.

a discounting factor $\gamma \in [0, 1)$, the *discounted cumulative return* of a trajectory is the vector

$$G(\tau) = \sum_{t=1}^{\infty} \gamma^{t-1} \mathbf{r}_t.$$

A *stationary policy* is function $\pi(a | s) : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ that forms a probability distribution such that $\sum_{a \in \mathcal{A}(s)} \pi(a | s) = 1$ for all s . Such a policy is *stationary* since the probability with which an action is selected depends only on the current state. More generally, a policy (not necessarily stationary) is a function $\pi(a | \tau, s)$ that may additionally depend on a given trajectory (intuitively, the history prior to reaching state s).

An *action value function* is defined as the expected total reward starting from s , taking action a , and following policy π thereafter:

$$q_{\pi}(s, a) := \mathbb{E}_{\tau \sim \pi} \left[\sum_{k=0}^{\infty} \gamma^k \mathbf{r}_{t+k} \mid s_t = s, a_t = a \right].$$

The value function of a policy π from state s is defined by:

$$v_{\pi}(s) := \mathbb{E}_{\tau \sim \pi} \left[\sum_{k=0}^{\infty} \gamma^k \mathbf{r}_{t+k} \mid s_t = s \right].$$

Our algorithms will aim to solve the learning problem by finding an estimate of $q_{\pi}(s, a)$ and $v_{\pi}(s)$. We denote such estimates as $Q_{\pi}(s, a)$ and $V_{\pi}(s)$, respectively.

4 PROBLEM FORMULATION

Our goal is to learn a policy π that maximizes $\mathbb{E}_{\tau \sim \pi} [G(\tau)]$ in all dimensions. To make this optimization objective concrete, we must specify a *scalarization* function $f : \mathbb{R}^n \rightarrow \mathbb{R}$. In fair reinforcement learning, we think of each of the n dimensions of the reward vector \mathbf{r} as corresponding to a distinct user to whom the learning agent wishes to be fair. The scalarization function can thus be thought of as a *welfare* function W over the users, and the learning agent is a welfare maximizer. For a given welfare function W , our goal is then to compute a policy π^* that maximizes expected welfare:

$$\pi^* = \arg \max_{\pi} \mathbb{E}_{\tau \sim \pi} [W(G(\tau))]. \quad (1)$$

4.1 Welfare Axioms

Here we describe desirable properties of welfare functions in terms of general outcomes as vectors \mathbf{v} . You can think of these vectors as possible discounted cumulative reward vectors corresponding to different policies. In the fair division literature [21], the most basic requirement of a welfare function is *monotonicity*.

Definition 4.1. W satisfies **monotonicity** if and only if for all \mathbf{v} and \mathbf{v}' with $v_j \geq v'_j$ for all j and $v_i > v'_i$ for some i , $W(\mathbf{v}) > W(\mathbf{v}')$.

Intuitively, *monotonicity* specifies that given all else equal, one prefers to increase a user's utility. Related is *Pareto optimality*, that an outcome should be efficient in the sense of not being dominated by any other outcome.

Definition 4.2. An outcome \mathbf{v} satisfies **Pareto optimality** if there is no other outcome \mathbf{v}' such that $v'_i \geq v_i$ for all users i , and at least one inequality is strict. W satisfies Pareto optimality if any \mathbf{v} (within some feasible space) that maximizes W is Pareto optimal.

Welfare functions should also satisfy *symmetry*, or indifference towards permutations of the input [21]. This is the most basic form of a fairness guarantee, that different users are treated similarly.

Definition 4.3. W satisfies **Symmetry** if for every permutation σ of its inputs \mathbf{v} , $W(\sigma(\mathbf{v})) = W(\mathbf{v})$.

A family of welfare functions satisfying the above properties are *generalized mean p -welfare functions* [4, 21], where $W_p(\mathbf{v}) = 1/n (\sum_{i=1}^n (v_i)^p)^{1/p}$. For instance, when $p = 1$, we have the utilitarian welfare function [7], the arithmetic mean of utilities.

Note that the utilitarian social welfare that may not be suitable for ensuring fairness on outcomes. Monotonicity, Pareto optimality, and Symmetry are merely minimal requirements for a welfare function. One may wish to introduce a stronger axiom such as the *Pigou-Dalton Principle* [9]. This principle states that one-to-one transfer of utility (or rewards in the MOMDP) from a better-off user to a worse-off user should increase the overall welfare.

Definition 4.4. W satisfies the **Pigou-Dalton Principle** if for all \mathbf{v}, \mathbf{v}' equal except for $v_i = v'_i + \delta$ and $v_j = v'_j - \delta$ where $v'_j - v'_i > \delta > 0$, $W(\mathbf{v}) > W(\mathbf{v}')$.

In other words, more equal distributions of utility are preferred. Functions that satisfy this formulation of fairness are often concave, capturing the diminishing marginal returns of increasing the utility of a user who already enjoys high utility relative to other users. Among the generalized mean p -welfare functions, the Pigou-Dalton Principle is satisfied for all $p < 1$. Our algorithm is designed to maximize welfare functions in this class.

4.2 Nash Social Welfare Function

The extreme case of a fair welfare function is the generalized mean p -welfare function where $p \rightarrow -\infty$, which corresponds to the *egalitarian* welfare function [24] that maximizes the minimum utility (and subject to that, optimizes the next smallest, and so forth). In-between the extremes of the utilitarian and egalitarian social welfare functions, we specifically focus on the **Nash Social Welfare** (NSW) function as our canonical example of a fair welfare function that also balances efficiency with fairness [8, 12, 16, 22].

$$\text{NSW}(\mathbf{v}) = \left(\prod_{i=1}^n v_i \right)^{\frac{1}{n}} \quad (2)$$

NSW is also simply the geometric mean of utilities, and in its log transform is also known as the proportional fairness objective. Note that NSW is a generalized mean p -welfare function where $p \rightarrow 0$. In addition to the previous desirable properties, NSW also enjoys the property of being *scale invariant*, meaning that the arg max of NSW is invariant under scaling of a given dimension of reward. From a practical perspective, this means that the relative scales of utility or reward for each dimension are not significant and do not need to be tuned during reward shaping.

Though we focus on NSW as our canonical example, we note that other reasonable welfare functions exist. For example, in Multi-Objective optimization, several works have studied Ordered Weighted Average (OWA) operators, a family of operators that contains many types of means [31].

4.3 Expected Welfare

In contrast to some prior work [1, 25] we focus on optimizing the expected welfare $\mathbb{E}_{\tau \sim \pi} [W(G(\tau))]$ rather than the welfare of the expectation $W(\mathbb{E}_{\tau \sim \pi} [G(\tau)])$. Note that for any concave welfare function, including NSW, *Jensen’s inequality* [11] implies

$$\mathbb{E}_{\tau \sim \pi} [W(G(\tau))] \leq W(\mathbb{E}_{\tau \sim \pi} [G(\tau)]). \quad (3)$$

We optimize for the lower bound (which turns out to be a more computationally challenging objective, see Section 5) in order to avoid treating policies as “fair” that are unfair in every particular episode and satisfy fairness only across several episodes on average.

Example (Expected Welfare). Consider the example diagrammed in Figure 1 with $n = 2$ users. Suppose we want to learn a policy that maximizes NSW. There is a stochastic policy π_1 that yields discounted cumulative reward of $(1, 0)$ with probability 0.5 and $(0, 1)$ with probability 0.5. There is also a deterministic policy π_2 that yields $(0.5 - \epsilon, 0.5 - \epsilon)$ (where $\epsilon > 0$ is small). The NSW expected reward under π_1 is 0.5, even though with probability 1, the NSW of every trajectory generated by π_1 is 0. By contrast, the NSW of π_2 is always $0.5 - \epsilon$. Maximizing the expected welfare, our optimization problem would prefer π_2 .




Figure 1: Example MOMDP. Dotted lines represent trajectories generated by π_1 , solid line for π_2

This example shows the intuition for why we choose to maximize $\mathbb{E}_{\tau \sim \pi} [W(G(\tau))]$. We seek to find a policy that generates trajectories with high expected welfare, a stronger property than generating high welfare of expected rewards. As we see in the next section, the problem is also computationally more challenging.

5 OPTIMIZING WELFARE

In general, one cannot provably and efficiently optimize all fair welfare functions. We first demonstrate that finding a policy that maximizes the NSW is APX-hard, implying one cannot get an arbitrarily close approximation efficiently, even in the tabular setting. We note that the same is not true for optimizing the NSW of expected rewards, for which the optimal stochastic policy can be computed efficiently [1].

Our argument follows via a reduction from the problem of allocating indivisible goods, in which m items must be partitioned among n users, where user a_i has utility $u_{i,j} \geq 0$ for good j and their utility for multiple goods is the sum of their utilities for the individual goods.

LEMMA 5.1. [17] *It is APX-hard to compute an indivisible allocation of goods optimizing the NSW.*

From this we can show the following impossibility.

THEOREM 5.2. *Computing the policy that maximizes $\text{NSW}(G(\tau))$ is APX-hard, even in a deterministic environment.*

PROOF. We reduce the problem of allocating a set of indivisible items to users with additive utilities. Given such an allocation problem, consider the following MOMDP. Find an arbitrary enumeration $\{s_1, s_2, \dots, s_m\}$ of the m items. These are the states of the MOMDP. At each time step, the environment transitions from s_j to s_{j+1} , beginning at s_1 and with s_m as a terminal state. In each state, s_j , the agent has n available actions $\{a_1, \dots, a_n\}$ where taking action a_i in state s_j corresponds to allocating item s_j to user a_i and receives a reward vector with $u_{i,j}$ in dimension i and 0 elsewhere.

Consider a policy π that is μ -approximately optimal on the NSW objective. The indivisible allocation corresponding to each user receiving the set of goods for which π chooses a_i is also μ -approximately optimal on the indivisible allocation problem. By Lemma 1, there exists a constant $\mu > 1$ such that it is NP-hard to approximate the NSW optimal allocation of indivisible items to agents with additive utilities within a factor μ . It must also be NP-hard to approximate the NSW optimal policy in a MOMDP within a factor μ . □

5.1 Non-stationary Welfare Q-Learning

We now present our algorithm, *Welfare Q-Learning*, which implements a variant on *Q-learning* [30], a model-free temporal-difference learning algorithm [26]. Our algorithm differs in two major ways from standard Q-learning.

- (1) Q table updates are chosen to maximize the (potentially) nonlinear W , and each value $Q(s, a)$ is a vector in \mathbb{R}^n corresponding to an estimate of the future reward vector possible that is welfare maximal.
- (2) Action selection is non-stationary. We keep track of the discounted cumulative reward vector within a trajectory so far and select the action that maximizes total estimated welfare including that already accumulated and future estimates.

We show experimentally in Section 6 that both of these changes are crucial to achieving high expected welfare for fair welfare functions such as NSW. We show in Theorem 5.3 that the algorithm still provably converges even with the nonlinear learning updates on a vector-valued Q table. To see the intuition for the significance of non-stationary action selection for optimizing the expected welfare of a nonlinear welfare function, we present the following example.

5.1.1 Non-stationarity. Consider the MOMDP diagrammed in Figure 2. There are n users. $|\mathcal{A}(s_1)| = 1$, and at $t = 1$, the environment transitions from s_1 to $s_{2,i}$ with probability $1/n$ and reward $\mathbf{0} \in \mathbb{R}^n$. At $t = 2$, the environment transitions from $s_{2,i}$ to s_3 with probability 1 and a reward vector that has 1 at the i^{th} component and 0 elsewhere. At $t = 3$, the agent gets to choose from n actions, a_1, a_2, \dots, a_n that yield associated with rewards $(1, \dots, 1, 0), (1, \dots, 0, 1), \dots, (0, \dots, 1, 1)$, respectively.

A stationary policy cannot achieve expected NSW greater than $1/n$ on this MOMDP, even though a non-stationary policy can achieve expected NSW of 1. Without loss of generality, assume a stationary stochastic policy π chooses the action a_i with probability p_i for $i = 1, \dots, n$ such that $\sum_i p_i = 1$. Then the expected Nash social

Algorithm 1 Welfare Q-Learning

```

1: Parameters: Learning rate  $\alpha \in (0, 1]$ , Discount factor  $\gamma \in$ 
    $[0, 1)$ , exploration rate  $\epsilon > 0$ , welfare function  $W$ 
2: Require: Initialize  $Q(s, a)$  for all  $s \in \mathcal{S}, a \in \mathcal{A}(s)$  arbitrarily
   except  $Q(\bar{s}, \cdot) \leftarrow 0$  for terminal states  $\bar{s}$ 
3: for each episode do
4:   Initialize  $s \leftarrow s_1, r_{acc} \leftarrow \mathbf{0}, c \leftarrow 0$ 
5:   repeat ▷ each step in an episode
6:      $a \leftarrow \begin{cases} \text{a uniform random action} & \text{with Pr}(\epsilon) \\ \arg \max_{a'} W(r_{acc} + \gamma^c Q(s, a')) & \text{otherwise} \end{cases}$ 
7:     Take action  $a$ , observe  $r, s'$ 
8:      $a^* \leftarrow \arg \max_a W[\gamma Q(s', a)]$ 
9:      $Q(s, a) \leftarrow Q(s, a) + \alpha[r + \gamma Q(s', a^*) - Q(s, a)]$ 
10:     $s \leftarrow s'$ 
11:     $r_{acc} \leftarrow r_{acc} + \gamma^c r$ 
12:     $c \leftarrow c + 1$ 
13:  until  $s$  is terminal
14: end for

```

welfare resulting from such a policy is

$$\mathbb{E}_{\tau \sim \pi} \left[\text{NSW} \left[\sum_t r_t \right] \right] = \sum_{i=1}^n \frac{1}{n} p_i = \frac{1}{n}$$

On the other hand, an optimal non-stationary policy π^* that keeps track of accumulated rewards is able to choose the correct complementary action at s_3 depending on the random transition at $t = 1$, so that $\mathbb{E}_{\tau \sim \pi^*} [\text{NSW} [\sum_t r_t]] = 1$.




Figure 2: Example of a stochastic MOMDP in which the expected Nash social welfare of stationary policies shrinks to 0 as the dimension of rewards n increases.

This example shows why, when optimizing for fairness, one should keep track of the discounted cumulative reward r_{acc} allocated to thus far. A stationary policy does not distinguish which users have received higher or lower rewards on a given trajectory thus far. Our approach seeks to be greedy on the sum of this discounted cumulative reward and future estimates of reward stored in the Q table. In this manner, users that have not received much reward so far within an episode are prioritized in action selection.

It is worth noting that in order to incorporate both information from the past and future, we need a consistent accounting of discounting for both terms to ensure our agent has correct information

from both past and future when deciding a fair policy:

$$r_{acc_t} = r_1 + \gamma r_2 + \gamma^2 r_3 + \dots + \gamma^{t-1} r_t$$

$$\gamma^t Q_\pi(s, a) = \mathbb{E}_\pi [\gamma^t r_{t+1} + \gamma^{t+1} r_{t+2} + \dots \mid S_t = s, A_t = a].$$

It is also important that each value in our learned Q table stores the vector of estimated future rewards rather than simply a scalar estimate of the welfare achievable, as the true greedy objective is the welfare of the sum of these vectors, not sum of the welfare of the two. Because of this, our Q table has size $|\mathcal{S}| \times |\mathcal{A}| \times n$ where n is the dimension of the reward. Indeed, as we show in Section 6, this does result in decreased convergence rates of the algorithm for larger n .

5.1.2 Convergence. We now argue that the Q values of our algorithm converge to an interpretable set of values.

THEOREM 5.3. For discount factor $\gamma \in [0, 1)$, the Q values of Welfare Q-Learning converge.

PROOF. The proof follows from *Banach's Fixed-Point Theorem* [3], which guarantees the existence and uniqueness of fixed-point of a contraction map on a complete metric space. The update step of Algorithm 1 can be seen as applying, in expectation, an operator on the Q-table. To apply the fixed-point theorem, we define a metric and show that this operator is a contraction. The Generalized Banach Fixed-Point Theorem therefore implies that Algorithm 1 will converge toward a unique fixed-point of this operator. Proofs for the two technical lemmas are provided in the full version [13].

Definition 5.4. Define d on the space of Q-tables \mathcal{Q} by

$$d(Q, Q') := \max_{\substack{s \in \mathcal{S}, a \in \mathcal{A} \\ i \in \{1, \dots, d\}}} |Q_i(s, a) - Q'_i(s, a)|.$$

LEMMA 5.5. $\langle \mathcal{Q}, d \rangle$ is a complete metric space.

Next, we define the *optimality filter* \mathcal{H} .

Definition 5.6. The optimality filter \mathcal{H} is an operator defined as

$$(\mathcal{H}Q)(s) = \arg \max_{a' \in \mathcal{A}} W(Q(s, a')),$$

where $\arg \max$ takes the multi-objective value corresponding to the maximum, i.e., $Q(S, a'')$ such that $a'' \in \arg \max_{a \in \mathcal{A}} W(Q(S, a))$.

We can define an *optimality operator* \mathcal{T} in terms of the optimal filter.

Definition 5.7. The optimality operator \mathcal{T} is defined as

$$(\mathcal{T}Q)(s, a) = r(s, a) + \gamma \mathbb{E}_{s' \sim \mathcal{P}(\cdot | s, a)} (\mathcal{H}Q)(s').$$

Note that in the algorithm, at each iteration, we sample from $\mathcal{P}(\cdot | s, a)$ to make an update. If the learning rate α satisfies the usual Robbins-Monro type conditions, namely $\sum \alpha = \infty$ and $\sum \alpha^2 < \infty$, the update at each iteration is, in expectation, applying the optimality operator \mathcal{T} . Thus, to show convergence, it suffices to show that iteratively applying \mathcal{T} on any Q leads to a unique Q-table.

LEMMA 5.8 (THE OPTIMALITY OPERATOR IS A CONTRACTION). Let Q, Q' be any two multi-objective Q-value functions, then $d(\mathcal{T}Q, \mathcal{T}Q') \leq \gamma d(Q, Q')$, where $\gamma \in [0, 1)$ is the discount factor of the underlying MOMDP.

Finally, since in our design, the distance d is a well-defined metric, to prove convergence to a unique fixed point, we will use the Generalized Banach Fixed Point Theorem as in [32].

LEMMA 5.9 (GENERALIZED BANACH FIXED-POINT [32]). *Given that \mathcal{T} is a contraction mapping with Lipschitz coefficient γ on the complete pseudo-metric space (\mathcal{Q}, d) , then there exists Q^* such that*

$$\lim_{n \rightarrow \infty} d(\mathcal{T}^n Q, Q^*) = 0$$

for any $Q \in \mathcal{Q}$.

It follows from the Lemmas that there exists Q^* such that

$$\lim_{n \rightarrow \infty} d(\mathcal{T}^n Q, Q^*) = 0$$

for any $Q \in \mathcal{Q}$.

Note that if the distance d between two tables is 0, the tables are equal. In other words, iteratively applying the optimality operator \mathcal{T} to a multi-objective Q -table will converge toward a unique Q -table. Since the update step in Algorithm 1 is applying \mathcal{T} in expectation, the algorithm also converges toward a unique Q -table. This concludes the proof of Theorem 5.3. \square

Note that the convergence result is not dependant on any particular welfare function W , but applies generally.

Next, we provide an interpretation of the unique fixed point Q^* of our algorithm. Note that the Bellman optimality conditions are not satisfied for nonlinear welfare functions, precluding the typical interpretation of the optimal Q -table. We nevertheless show that a very similar interpretation can be given in which $Q^*(s, a)$ provides a lower bound estimate on the discounted cumulative reward vector that is achievable after taking action a in state s and then optimizing for W .

Definition 5.10. Define a policy π_Q^* given a Q -table as follows: at a given state s , let $\pi_Q^*(s_t) = a_t^* = \arg \max_a W(Q(s, a))$ and let $\mathbf{r}(s_t, a)$ be the immediate reward of performing action a in state s . Then the value function corresponding to $\pi_Q^*(s)$ is

$$V_{\pi_Q^*}(s) = \mathbb{E}_{\tau \sim (\mathcal{P}, \pi_Q^*)} \sum_{t=1}^{\infty} \gamma^{t-1} \mathbf{r}(s_t, a_t^*).$$

LEMMA 5.11 (INTERPRETING THE FIXED-POINT). *Let Q^* be the unique fixed point of the algorithm, i.e. $Q^* = \mathcal{T}Q^*$, then*

$$Q^*(s, a) = \mathbf{r}(s, a) + \gamma \mathbb{E}_{s' \sim \mathcal{P}(\cdot | s, a)} V_{\pi_{Q^*}^*}(s').$$

PROOF. Since Q^* is the fixed point of the algorithm, $Q^* = \mathcal{T}Q^*$. Expanding this using the definition of \mathcal{T} , we get

$$Q^*(s, a) = \mathbf{r}(s, a) + \gamma \mathbb{E}_{s' \sim \mathcal{P}(\cdot | s, a)} \arg \max_{a' \in \mathcal{A}(s')} W(Q^*(s', a')).$$

So it suffices to show that $\arg \max_{a' \in \mathcal{A}} W(Q^*(s', a')) = V_{\pi_{Q^*}^*}(s')$. But expanding the LHS and RHS recursively and recalling $a_0^* = \arg \max_{a'} W(Q^*(s', a'))$ (see Definition 5.10) as well as the definition of $\arg \max$ in Definition 5.6, we get

$$\begin{aligned} \arg \max_{a' \in \mathcal{A}} W(Q^*(s', a')) &= Q^*(s', a_0^*) = \mathcal{T}[Q^*(s', a_0^*)] \\ &= \mathbf{r}(s', a_0^*) + \gamma \mathbb{E}_{s'' \sim \mathcal{P}(\cdot | s', a_0^*)} \arg \max_{a'' \in \mathcal{A}} W(Q^*(s'', a'')) \end{aligned}$$

and

$$\begin{aligned} V_{\pi_{Q^*}^*}(s') &= \mathbb{E}_{\tau \sim (\mathcal{P}, \pi_{Q^*}^*)} \sum_{t=1}^{\infty} \gamma^{t-1} \mathbf{r}(s_t, a_t^*) \\ &= \mathbf{r}(s', a_0^*) + \gamma \mathbb{E}_{s'' \sim \mathcal{P}(\cdot | s', a_0^*)} \left[\mathbb{E}_{\tau \sim (\mathcal{P}, \pi_{Q^*}^*)} \sum_{t=1}^{\infty} \gamma^{t-1} \mathbf{r}(s_t, a_t^*) \right] \\ &= \mathbf{r}(s', a_0^*) + \gamma \mathbb{E}_{s'' \sim \mathcal{P}(\cdot | s', a_0^*)} \left[V_{\pi_{Q^*}^*}(s'') \right]. \end{aligned}$$

So in turn, it suffices to show that $\arg \max_{a'' \in \mathcal{A}} W(Q^*(s'', a''))$ is equal to $V_{\pi_{Q^*}^*}(s'')$. Repeat this argument until the agent is in one of the terminal states \bar{s} . Note that $Q(\bar{s}, \cdot) = 0$, i.e. Q -values for the terminal states are zero for all actions. So

$$\arg \max_{a \in \mathcal{A}} W(Q^*(\bar{s}, a)) = 0 = V_{\pi_{Q^*}^*}(\bar{s}).$$

This completes the proof of Lemma 5.11. \square

Lemma 5.11 implies that each entry $Q^*(s, a)$ represents what the agent could actually expect to receive as total discounted future reward in expectation if one performs action a in initial state s , then follows greedy stationary action selection using the Q -table. This is essentially the same interpretation as in traditional scalar Q -learning, except that there is no optimality guarantee of the result for non-linear welfare optimization.

Intuitively, these entries only serve as an estimate of the lower bound of total discounted rewards that are achievable in the future. Algorithm 1 combines these estimates with keeping track of the discounted cumulative rewards up to a given point in order to greedily optimize for the welfare of the sum of the two vectors. This is the essential intuition behind Algorithm 1.

6 EXPERIMENTS

We run experiments under two simulated environments diagrammed in Figure 3.³ Our results demonstrate that (a) *Welfare Q-Learning* is effective in finding policies with high expected welfare compared with other baselines, (b) the rate of convergence depends on n , the dimensionality of the reward space, and (c) linear scalarization and mixture policies are generally inadequate for optimizing fair welfare functions. All results for all algorithms are obtained using an average of NSW and utilitarian welfare of \mathbf{r}_{acc} for each episode over 50 runs. For all the experiments on both environments, each unit on the x -axis corresponds to an episode, which equals to 10000 timesteps or action selections in the environment. The duration of a timestep is the same for all methods.

6.1 Metrics, Methods, and Baseline Algorithms

All of our experiments attempt to optimize NSW (results for other welfare functions are provided in the full version [13]). We measure the NSW function on \mathbf{r}_{acc} earned thus far. As elaborated in Section 1, NSW satisfies all of our basic desiderata plus scale invariance, and is an intermediate welfare function between the extremes of egalitarian and utilitarian social welfare. For comparison, we also show utilitarian social welfare (the arithmetic mean of reward vectors) alongside NSW.

³The implementation is available at <https://github.com/MuhangTian/Fair-MORL-AAMAS>

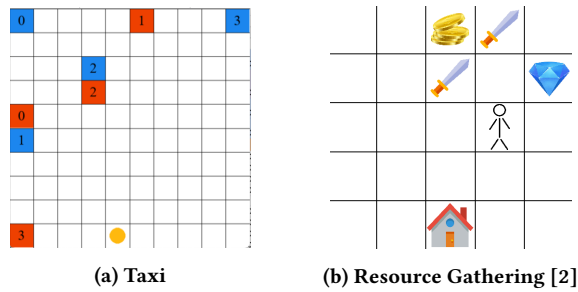


Figure 3: Simulated Environments

The geometric mean can be numerically unstable. In experiments we work with the log transform of NSW. That is, instead of maximizing $NSW(R) = (\prod_i^n R_i)^{1/n}$, we equivalently maximize $\sum_i^n \ln(R_i + \lambda)$, where $\lambda > 0$ is included as a smoothing factor in case $R_i = 0$. Due to the nature of the NSW function, the NSW of rewards with negative elements is undefined, or alternatively can be defined as $-\infty$. Thus, our scope of exploration is restricted to policies that yield all non-negative accumulated rewards.

6.1.1 Baseline Algorithms. We compare our algorithm against three baselines.

- (1) *Optimal Linear Scalarization.* A simple MORL technique is to apply linear scalarization on the Q-table [29]. Given weights $\mathbf{w} \in \mathbb{R}^n$, where $\sum_{i=1}^n w_i = 1$ for n objectives, let $SQ(s, a) = \sum_{i=1}^n w_i Q(s, a)_i$, where $Q(s, a)_i$ is the Q-value for i^{th} objective. For each time step, $SQ(s, a)$ is treated by the algorithm as the objective to perform both action selection with ϵ -greedy and learning updates of $Q(s, a)$. We chose the weights \mathbf{w} that performed best on the NSW objective as determined by a grid search through combinations of \mathbf{w} . We include this algorithm as the baseline to demonstrate the limitations of linear scalarization on non-linear objectives and show the importance of our non-linear learning updates in Algorithm 1.
- (2) *Stationary Policy.* Algorithm 1, *Welfare Q-Learning*, learns a particular Q-table corresponding to a (potentially) non-linear welfare function, then performs non-stationary action selection. By contrast, we also show the results if one stationary ϵ -greedy action selection on the same learned Q-table. That is, the stationary policy algorithm does not consider r_{acc} in its action selection. We include this algorithm to show the importance of non-stationary action selection.
- (3) *Optimal Mixture Policy.* [28] proposed the idea of combining multiple Pareto Optimal base policies into a single mixture policy. We chose our base policies as those that optimize each dimension of the reward vector independently. The algorithm then uses one of these policies for I time steps, then switches to the next. To determine the optimal value of I for optimizing NSW, we used a grid search. We use the resulting optimal value I^* . This baseline examines the effectiveness of intuitive approaches (combining optimal policies for each user) for optimizing fairness.

6.2 Taxi Environment

Description. Inspired by the taxi toy example problem for single-objective RL [10], we designed a multi-objective taxi environment⁴. In this grid world, our agent is a taxi driver whose goal is to deliver passengers from their origins to their destinations. There are n origin-destination pairs, one for each dimension of reward, and the agent earns reward in that dimension when dropping off a passenger from that origin-destination pair. There are an unlimited number of passengers for each origin-destination pair, but the taxi can only take one passenger at a time. This constraint enforces objectives to be conflicting, thus our agent’s fairness performance becomes more important—it should provide its delivery service to each origin successfully and fairly over time, without ignoring origins that are more difficult to deliver (such as number 3 origin/destination pair in Figure 3a).

Results. Results are shown in Figure 4. *Welfare Q-Learning* achieves the maximum average NSW score among all the algorithms, and still manages to achieve the second highest utilitarian score. We observe that our non-stationary policy outperforms the stationary policy on the same Q-table for both NSW and utilitarian score. Note that a stationary policy that optimizes NSW on this environment must essentially make a large loop always taking each origin-destination pair in turn, whereas a non-stationary policy can selectively optimize a single origin-destination pair for several time steps before switching to another pair.

Linear scalarization has the lowest average NSW since there simply does not exist a set of weights that would produce accumulated rewards in all dimensions. It achieves highest utilitarian score since it favors to complete delivery for closest origin/destination pairs (such as index 2 pair in Figure 3a). The mixture policy performs generally well but slightly lower than that of *Welfare Q-Learning*, this is because an optimal fair policy in this environment does have the structure of alternating between optimizing on different dimensions at a time. This is not, however, true in general, as is seen in the next environment. Although the mixture policy converges quickly (each dimension independently is very easy to optimize), such performance is also subject to finding the optimal interval for the taxi environment I^* (227 timesteps) via a search through the parameter space, which involves a computational cost not reflected by the figures.

For *Welfare Q-Learning*, we observe that there is an inverse correlation between the dimensionality of the reward space n and the rate of convergence, as shown in Figure 4c. A possible explanation is that the increase in dimensionality increases the size of the Q-table, which is of size $|S| \times |\mathcal{A}| \times n$, thus more updates are needed to converge.

6.3 Resource Gathering (RG) Environment

Description. Building on the resource-collection domain [5], we modified the environment to include more complexity and randomness, and we also restructured it as a multi-objective setting.⁵ In this grid world, our agent collects three types of resources (gold, gem, and sword) spawned randomly at different locations which

⁴See detailed description of the environment in the full version [13]

⁵See detailed description of the environment in the full version [13]

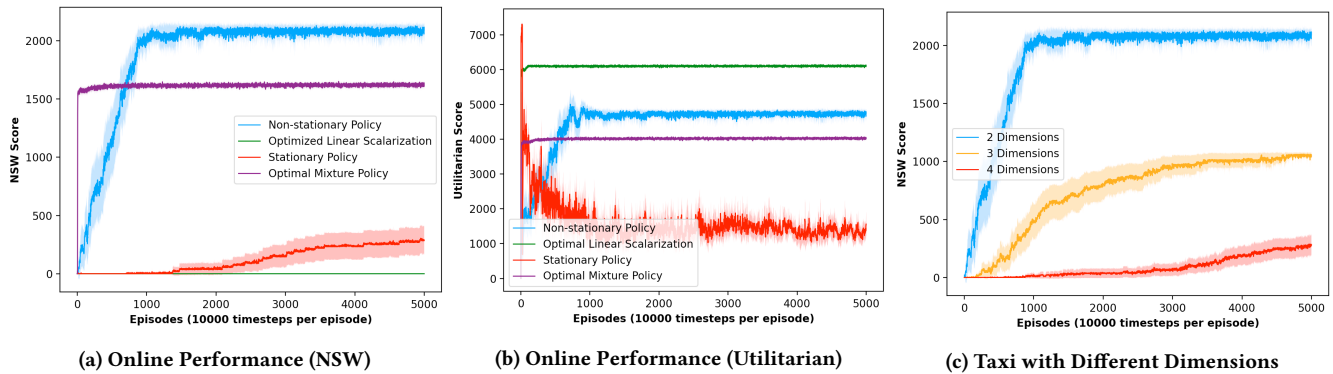


Figure 4: Experiment Results for Taxi Environment. Non-stationary Policy is Welfare Q-Learning




Figure 5: Experiment Results for RG Environment. Non-stationary Policy is Welfare Q-Learning

also disappear with a given probability. The dimensions of reward correspond to the resources. That is, gathering a given resource earns reward in that dimension and 0 the others. The goal of our agent is to collect as many resources as possible while maintaining a balance between difference types of resources.

Results. *Welfare Q-Learning* achieves the maximum average NSW score among all the algorithms. Although linear scalarization achieves a reasonable utilitarian score, it fails to satisfy fairness, as it largely ignores one of the dimensions/resource types. The diminishing performance of the linear scalarization on NSW is presumably due to the algorithm increasingly coming to optimize for some dimensions and ignoring others. The mixture policy achieved the most comparable utilitarian score to the non-stationary policy, yet substandard NSW score. Unlike in the taxi environment, the optimal policy in RG is not characterized by optimizing each objective sequentially, which results in ignoring certain type of nearby resources.

An additional desirable property of optimizing NSW is *scale invariance*, that the optimizing policy is invariant with respect to changes in the scale of a dimension of reward. We demonstrate this empirically by comparing the distribution over resources gathered using Nash Welfare Q-Learning (*Welfare Q-Learning* with NSW as the welfare function) versus linear scalarization in Figure 5c. When all resources are worth the same amount of reward in their

respective dimension (Figure 5c Top) the NSW agent achieves a balanced distribution between the three types of resources, while the linear scalarized agent tends to neglect swords in the third dimension. To rectify this one might try to rescale the rewards (Figure 5c Bottom) so that the value of swords is scaled to 50, but then the scalarized policy changes drastically to gather swords and ignore gems. However, the NSW agent is immune to scaling of this kind, retaining the same general distribution of reward in both cases. In other words, fair welfare (and especially NSW) optimization in a MOMDP may be more robust to specifications of the environment as compared to techniques based on linear scalarization.

7 FUTURE WORK

While we know that exact optimization of, say, NSW is intractable, we do not know what provable approximation factors might be achievable in tabular MORL. Additionally, we observe a dependence between the dimensionality n and the convergence rate of our algorithm, but we do not know whether this “curse of dimensionality” is fundamental to nonlinear welfare optimization in MORL. Finally, though even the tabular setting is challenging for nonlinear welfare optimization in MORL, we believe the intuition of non-stationary action selection coupled with nonlinear learning updates can be extended to the function approximation setting and combined with deep neural network representations.

REFERENCES

- [1] Mridul Agarwal, Vaneet Aggarwal, and Tian Lan. 2022. Multi-Objective Reinforcement Learning with Non-Linear Scalarization. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (Virtual Event, New Zealand) (AAMAS '22)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 9–17.
- [2] Lucas N. Alegre, Florian Felten, El-Ghazali Talbi, Grégoire Danoy, Ann Nowé, Ana L. C. Bazzan, and Bruno C. da Silva. 2022. MO-Gym: A Library of Multi-Objective Reinforcement Learning Environments. In *Proceedings of the 34th Benelux Conference on Artificial Intelligence BNAIC/Benelearn 2022*.
- [3] Stefan Banach. 1922. Sur les opérations dans les ensembles abstraits et leur application aux équations intégrales. *Fund. math* 3, 1 (1922), 133–181.
- [4] Siddharth Barman, Arindam Khan, and Arnab Maiti. 2022. Universal and Tight Online Algorithms for Generalized-Mean Welfare. *Proceedings of the AAAI Conference on Artificial Intelligence* 36, 5 (Jun. 2022), 4793–4800. <https://doi.org/10.1609/aaai.v36i5.20406>
- [5] Leon Barrett and Srinu Narayanan. 2008. Learning all optimal policies with multiple criteria. In *Proceedings of the 25th international conference on Machine learning*, 41–47.
- [6] Richard Bellman. 1966. Dynamic programming. *Science* 153, 3731 (1966), 34–37.
- [7] Ken Binmore. 1998. Egalitarianism versus utilitarianism. *Utilitas* 10, 3 (1998), 353–367.
- [8] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D Procaccia, Nisarg Shah, and Junxing Wang. 2019. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation (TEAC)* 7, 3 (2019), 1–32.
- [9] Erio Castagnoli and Pietro Muliere. 1990. A note on inequality measures and the Pigou-Dalton principle of transfers. In *Income and wealth distribution, inequality and poverty*. Springer, 171–182.
- [10] Thomas G Dietterich. 2000. Hierarchical reinforcement learning with the MAXQ value function decomposition. *Journal of artificial intelligence research* 13 (2000), 227–303.
- [11] Rick Durrett. 2019. *Probability: theory and examples*. Vol. 49. Cambridge university press.
- [12] Brandon Fain, Kamesh Munagala, and Nisarg Shah. 2018. Fair allocation of indivisible public goods. In *Proceedings of the 2018 ACM Conference on Economics and Computation*. 575–592.
- [13] Zimeng Fan, Nianli Peng, Muhang Tian, and Brandon Fain. 2022. Welfare and Fairness in Multi-objective Reinforcement Learning. <https://doi.org/10.48550/ARXIV.2212.01382>
- [14] Tianmeng Hu, Biao Luo, and Chunhua Yang. 2021. Multi-objective optimization for autonomous driving strategy based on Deep Q Network. *Discover Artificial Intelligence* 1, 1 (2021), 1–11. <https://doi.org/10.1007/s44163-021-00011-3>
- [15] Shahin Jabbari, Matthew Joseph, Michael Kearns, Jamie Morgenstern, and Aaron Roth. 2017. Fairness in reinforcement learning. In *International conference on machine learning*. PMLR, 1617–1626.
- [16] Mamoru Kaneko and Kenjiro Nakamura. 1979. The Nash social welfare function. *Econometrica: Journal of the Econometric Society* (1979), 423–435.
- [17] Euiwoong Lee. 2017. APX-hardness of maximizing Nash social welfare with indivisible items. *Inform. Process. Lett.* 122 (2017), 17–20. <https://doi.org/10.1016/j.ipl.2017.01.012>
- [18] Chunming Liu, Xin Xu, and Dewen Hu. 2014. Multiobjective reinforcement learning: A comprehensive overview. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 45, 3 (2014), 385–398.
- [19] R Duncan Luce and Howard Raiffa. 1989. *Games and decisions: Introduction and critical survey*. Courier Corporation.
- [20] 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.
- [21] Hervé Moulin. 2004. *Fair division and collective welfare*. MIT press.
- [22] John F Nash Jr. 1950. The bargaining problem. *Econometrica: Journal of the econometric society* (1950), 155–162.
- [23] Włodzimierz Ogryczak, Hanan Luss, Michał Pióro, Dritan Nace, and Artur Tomaszewski. 2014. Fair optimization and networks: A survey. *Journal of Applied Mathematics* 2014 (2014).
- [24] Amartya Sen. 2018. *Collective choice and social welfare*. Harvard University Press.
- [25] Umer Siddique, Paul Weng, and Matthieu Zimmer. 2020. Learning Fair Policies in Multi-Objective (Deep) Reinforcement Learning with Average and Discounted Rewards. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119)*, Hal Daumé III and Aarti Singh (Eds.). PMLR, 8905–8915. <https://proceedings.mlr.press/v119/siddique20a.html>
- [26] Richard S Sutton. 1988. Learning to predict by the methods of temporal differences. *Machine learning* 3, 1 (1988), 9–44.
- [27] Richard S Sutton, Andrew G Barto, et al. 1998. *Introduction to reinforcement learning*. (1998).
- [28] 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.
- [29] 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 (AD-PRL)*. IEEE, 191–199.
- [30] Christopher John Cornish Hellaby Watkins. 1989. Learning from delayed rewards. (1989).
- [31] Ronald R. Yager. 1993. Families of OWA operators. *Fuzzy Sets and Systems* 59, 2 (1993), 125–148. [https://doi.org/10.1016/0165-0114\(93\)90194-M](https://doi.org/10.1016/0165-0114(93)90194-M)
- [32] Runzhe Yang, Xingyuan Sun, and Karthik Narasimhan. 2019. *A Generalized Algorithm for Multi-Objective Reinforcement Learning and Policy Adaptation*. Curran Associates Inc., Red Hook, NY, USA.