

# Fairness in Reinforcement Learning: A Survey

Anka Reuel<sup>1</sup>, Devin Ma<sup>2</sup>

<sup>1</sup>Stanford University

<sup>2</sup>University of Pennsylvania

anka.reuel@stanford.edu, devinma@upenn.edu

## Abstract

While our understanding of fairness in machine learning has significantly progressed, our understanding of fairness in reinforcement learning (RL) remains nascent. Most of the attention has been on fairness in one-shot classification tasks; however, real-world, RL-enabled systems (e.g., autonomous vehicles) are much more complicated in that agents operate in dynamic environments over a long period of time. To ensure the responsible development and deployment of these systems, we must better understand fairness in RL. In this paper, we survey the literature to provide the most up-to-date snapshot of the frontiers of fairness in RL. We start by reviewing where fairness considerations can arise in RL, then discuss the various definitions of fairness in RL that have been put forth thus far. We continue to highlight the methodologies researchers used to implement fairness in single- and multi-agent RL systems before showcasing the distinct application domains that fair RL has been investigated in. Finally, we critically examine gaps in the literature, such as understanding fairness in the context of RLHF, that still need to be addressed in future work to truly operationalize fair RL in real-world systems.

## Introduction

As artificial intelligence (AI) has been thrust into the public spotlight by the recent proliferation of generative AI applications, so has the issue of potential harm that AI can do to users and society in general (Bommasani et al. 2021). One particular issue is fairness. Before the advent of machine learning (ML), fairness was largely studied as a philosophical and societal concept (Rawls 1971; Wolff 1998; Rabin 1993). Nowadays, it is an active field of research (Corbett-Davies and Goel 2018; Mehrabi et al. 2021), studied with mathematical rigor by theorists and practitioners alike. Such efforts have enabled industry to start operationalizing both the implementation of fairness-aware algorithms and auditing fairness in products and services. However, one important area of ML where the study of fairness is still nascent is reinforcement learning (RL), in which autonomous agents interact in an environment to maximize a pre-specified reward (Sutton and Barto 2018). Whereas most efforts so far in fairness focus on

one-shot classification and decision-making tasks (Mehrabi et al. 2021), real-world systems, especially in robotics, often involve autonomous agents (e.g., the robot) interacting with the environment over a longer period of time. In such environments, the agent’s decision now will impact its future rewards. Often, these agents also impact humans during their interaction with the environment. Because RL is at the heart of many modern autonomous systems, including robotics, autonomous driving, and healthcare, the study of fairness in RL is particularly relevant today.

Typically, RL maximizes the long-term discounted reward of an agent where the reward is tailored to one specific objective (Sutton and Barto 2018). This reward can be, for example, overall output, points in a game, or stock market gains. While this ensures that the reward with respect to the defined objective is maximized, it does not necessarily involve considerations of group or individual fairness (Weng 2019). Only focusing on performance rewards but ignoring fairness notions comes with various negative ramifications in certain RL systems. First, using RL algorithms without fairness constraints can result in potentially discriminatory outcomes that are unethical and harmful. Depending on the application the algorithm is used in, unfair outcomes further have the potential to propagate existing societal biases (Vlasceanu and Amodio 2022). Second, in many developed countries, explicit mandates exist to guarantee fairness (Hacker 2018). Developers using RL-based systems need to ensure, by law, that their algorithms do not discriminate on protected attributes like gender, race, or religion. Third, if certain groups are underrepresented, user-centric systems might be abandoned altogether by the unfairly treated groups and other users. For example, if an RL algorithm pushes more white content creators on a platform like YouTube than minority creators, the latter might leave the platform due to inadequate exposure. Subsequently, the platform’s content may become more homogeneous, leading to less satisfied users (Liu et al. 2020). This outcome is undesirable to both the users and the company that markets the product. Since Jabbari et al. (2017) initiated the study of fairness in RL in 2017, there has been a growing interest in studying the theory and applications of fair RL. Despite these efforts, our understanding of fundamental questions about fairness in RL remains nascent. This issue

necessitates a comprehensive survey of the field to provide a snapshot of the frontiers of fairness in RL, offer a review of the definitions and methods to incorporate fairness in RL, and consider the implications of incorporating fairness into RL-enabled socio-technical systems.

Earlier work by Zhang and Liu (2021) presented a survey on fairness in learning-based sequential decision-making without considering bandit settings or a critical comparative analysis of definitions, methods, or application areas of fair RL. To our knowledge, only Gajane et al. (2022) has attempted to survey the field as we do. While they provide a foundational understanding, it predates several pivotal developments that have emerged over the last years, such as using RL in the context of large language models (LLMs), e.g., as part of RL from human feedback (RLHF). We further provide a more comprehensive analysis of the reviewed literature across multiple dimensions, namely, used fairness definitions, single- and multi-agent approaches, application domains, and trade-offs, and highlight papers not covered by previous work. Our paper offers an extensive survey of the most recent advancements in fairness in RL and uncovers gaps and potential avenues for future research that have yet to be explored.

The rest of the paper is organized as follows. The Section ‘Preliminaries’ will provide a brief theoretical overview of RL. In the Section ‘Definitions of Fairness in RL’, we introduce fairness definitions used in RL systems. We then show potential application domains for fair RL systems in the Section ‘Application Domains’. The Section ‘Methodology’ follows with a description of the methods used to introduce fairness notions in single- and multi-agent RL systems before we discuss trade-offs in the Section ‘Trade-Offs’. We conclude our work with an overview of potential future research directions in the Section ‘Future Research Directions’.

## Preliminaries

We start by reviewing the fundamentals of RL. For a more detailed account of RL, we refer the reader to Sutton and Barto (2018)’s work. RL is a process where an agent learns to make decisions by trying different actions to maximize rewards without being explicitly told which actions to take (Sutton and Barto 2018). The most significant aspect of this approach is the balance of exploration and exploitation. RL differs from other methods by starting with a fully interactive agent that has specific goals, can perceive its environment, and take actions. A key challenge is dealing with uncertainty in the environment (Sutton and Barto 2018). Key elements in RL include actions  $a$  an agent can take, states  $s$  in the environment the agent is interacting in, a policy  $\pi$ , which guides the agent’s behavior, and a reward signal  $r$  that the agent receives at some point during its interaction with the environment. The agent aims to maximize the total reward over time. The rewards can be immediate or delayed, depending on the problem formulation (see Section ‘Preliminaries’).

RL problems can further be formulated as single-objective, where the goal is to maximize a single metric or achieve a specific outcome, or multi-objective setting, where the agent has to balance multiple competing objectives, requiring a more nuanced decision-making process to optimize across various goals. The latter formulation adds a layer of complexity, as the agent must navigate and prioritize these objectives in its decision-making process. Another dimension to categorize RL methods is whether they are using an explicit model of the environment and interaction dynamics. Methods can be model-based, allowing for planning and prediction, or model-free, which relies on trial-and-error learning (Sutton and Barto 2018). Finally, problems can be formulated in RL based on the number of states: bandit problems are RL problems with only a single state, while Markov decision processes (MDPs) are a more general problem formulation that allows for multiple states.

## Multi-Armed Bandits

More straightforward RL problems are situations with only one state. These problems do not require learning different actions for different scenarios (Sutton and Barto 2018) and can be formulated as a version of the multi-armed bandit (MAB) problem: The agent is repeatedly choosing from  $k$  bandits, each yielding a numerical reward based on a fixed probability distribution. The goal is to maximize the total reward over a set number of choices or time steps (Sutton and Barto 2018). This approach, known as nonassociative learning, has been the foundation of most evaluative feedback research. On the contrary, imagine several  $k$ -armed bandit tasks, each randomly presented but with a unique clue, like a color change in the display, hinting at its identity (Sutton and Barto 2018). This clue does not reveal the (expected) rewards of each bandit but helps to associate each task with its best action; this setting is called an associative or contextual bandit task (Sutton and Barto 2018). It combines trial-and-error learning to find the best actions and the ability to associate these actions with the specific contexts where they are most effective (Sutton and Barto 2018). For an in-depth review of bandit algorithms, we refer the reader to Lattimore and Szepesvári (2020) and Sutton and Barto (2018).

## Markov Decision Processes

A more general formulation of RL problems is through MDPs. MDPs are a fundamental framework for sequential decision-making settings, where actions impact not only immediate rewards but also the reachability of future states and, consequently, future rewards (Sutton and Barto 2018). This setting introduces the complexity of balancing immediate and delayed rewards. Unlike bandit problems, where the value of each action is estimated independently, in MDPs, the value is determined for each action within each state. This state-action value, denoted as  $Q_\pi(s, a)$ , is the expected return of an agent, starting from state  $s$ , taking action  $a$ , and then following policy  $\pi$  (Sutton and Barto 2018). Relatedly, the value of a state assuming optimal action selection,  $V_\pi(s)$ , captures the expected value of an agent starting in state  $x$  and following policy  $\pi$ . These notions help attribute long-term strategies’ outcomes to specific actions (Sutton

and Barto 2018). MDPs frame the challenge of learning through interaction with the environment over multiple time steps to achieve a goal. This interaction is continuous, with the agent choosing actions and the environment responding, presenting new scenarios, and offering rewards (Sutton and Barto 2018). The traditional MDP model can be extended to multi-objective sequential decision-making, where the scalar reward is replaced by a vector whose components represent objectives, also called multi-objective MDPs (MOMDPs) (Pirotta, Parisi, and Restelli 2015). The agent aims to maximize these rewards over time through strategic action choices. The distinction between MDPs and bandit problems lies in their complexity and scope. While bandit problems focus on selecting the best action from a set without considering different states or the long-term consequences of actions, MDPs consider the entire sequence of actions and states, focusing on how each decision influences future scenarios and rewards (Sutton and Barto 2018).

### Definitions of Fairness in RL

One of the difficulties with studying fairness in RL comes from the numerous definitions of fairness one can adopt. Researchers debated notions of fairness in fields such as philosophy and economics long before fairness in RL – and ML more broadly – were discussed (Rawls 1971; Wolff 1998; Rabin 1993), underscoring the complexity and multifaceted nature of the concept. The notions used by ML researchers can be categorized based on where they fit into the ML development cycle (e.g., pre-processing, in-processing, or post-processing fairness) or by their intended aim (focusing on fairness with respect to individuals or groups). We recommend further reading for those interested in comprehensively examining fairness in ML (Mehrabi et al. 2021). This chapter, however, narrows its focus to the definitions of fairness specifically within the field of RL.

**Welfare-Based Definitions.** Siddique, Weng, and Zimmer (2020) and Zimmer et al. (2021) consider fairness in RL and encode fairness requirements as social welfare functions in the objective function. In both papers, the definition of fairness builds on three components: impartiality, equity, and efficiency. First, impartiality means that all agents are identical. Secondly, equity is based on the *Pigou-Dalton principle*, “which states that transferring utility from a better-off user to a worse-off user results in a fairer solution” (Yu, Siddique, and Weng 2023), given that all agents’ total reward/utility remains the same (Yu, Siddique, and Weng 2023). Thirdly, efficiency is the idea that between two potential solutions, if one solution is preferred by all agents, weakly or strictly, then that should be chosen over the other (Zimmer et al. 2021). This requirement is necessary because it prevents a situation where giving no reward to all users is treated the same as giving non-zero rewards to all users, even though both satisfy the impartiality and equity requirements. These principles are encoded in social welfare functions (SWF), which measure how good a utility vector is concerning social good. Concretely, a SWF is a function  $\phi : \mathbb{R}^D \rightarrow \mathbb{R}$ , where  $D$  is the number of objectives in Siddique, Weng, and Zimmer (2020) and the number of

users<sup>1</sup> in Zimmer et al. (2021).

Despite building on the same three components, it is important to realize the differences in the definitions Siddique, Weng, and Zimmer (2020) and Zimmer et al. (2021) proposed. Siddique, Weng, and Zimmer (2020) considers learning fair policies in *single-agent* deep reinforcement learning. Because of the single-agent setting, the notion of “fairness” is not across multiple users or agents but across multiple objectives or criteria. Informally, fairness in this single-agent context refers to how the solution “balances” the different objectives. The critical assumption in the multi-agent setting in Zimmer et al. (2021)’s work is that all users are identical, implying that with regard to the impartiality requirement, users should be treated at least similarly. Another important difference is that Siddique, Weng, and Zimmer (2020) operationalize these fairness components using the generalized *Gini* social welfare function (GGF), where the GGF is defined as

$$\text{GGF}_w(v) = \sum_{i=1}^D w_i v_i^\uparrow \quad (1)$$

where  $v \in \mathbb{R}^D$  is the utility vector,  $w \in \mathbb{R}^D$  is the fixed positive weight vector, and  $v^\uparrow$  is the vector with the components of vector  $v$  sorted in ascending order (Zimmer et al. 2021). Depending on the choice of  $w$ , this approach allows for some flexibility with respect to the exact fairness notion adopted. For example, if  $w_1 = 1$  and all other  $w_i = 0$  where  $i \neq 1$ , then the implementation corresponds to the maxmin egalitarian notion of fairness (Rawls 1971).

**Weighted Proportional Fairness Definition.** Kushner and Whiting (2004) originally defined proportional fairness in the scheduling context to balance two conflicting objectives, namely maximizing the total throughput of a network and guaranteeing a minimum level of service to all users. Liu et al. (2020) extended this definition to use weighted proportional fairness as their target fairness metric in the context of interactive recommender systems (RS). The authors first define an allocation vector  $x_t^i$  representing the allocation proportion of a group  $i$  up to time  $t$ :

$$x_t^i = \frac{\sum_{k=1}^t y_{a_k} \mathbb{I}_A c_i(a_k)}{\sum_{i=1}^l \sum_{k=1}^t y_{a_k} \mathbb{I}_A c_i(a_k)} \quad (2)$$

where  $A$  denotes a group of items with an attribute value  $c$  and  $i$  refers to the  $i$ th group.  $\mathbb{I}_A(x)$  is 1 if  $x \in A$  and 0 otherwise. They further define  $y_{a_k}$  as a user’s feedback on a recommended item  $a_k$ . The authors then define weighted proportional fairness as a “generalized Nash solution for multiple groups” (Liu et al. 2020), where a Nash solution describes “agent decisions that are self-enforceable, e.g. decisions such that no agent has an incentive to deviate from” (Krilašević and Grammatico 2021). The weighted proportionally fair allocation is, hence, the solution to the following

<sup>1</sup>Zimmer et al. (2021) formulated the problem in terms of “users”, not agents, because a user can represent an individual or a group of individuals, leading to a more general representation.

optimization problem:

$$\begin{aligned} & \max x_t \sum_{i=1}^l w_i \log(x_t^i) & (3) \\ \text{s.t. } & \sum_{i=1}^l x_t^i = 1, x_t^i \geq 0, i = 1, \dots, l \end{aligned}$$

where  $w_i$  is a parameter used to weigh the importance of each group. Introducing a weighted component allows for a more differentiated, nuanced notion of fairness that can be adapted to a specific context where there is a need for a nuanced consideration of protected groups.

**Coefficient of Variation.** In multi-agent systems where we want to ensure a fair distribution of resources or rewards among agents, the coefficient of variation ( $cv$ ) is often used to measure fairness among agents (Jiang and Lu 2019; Elmalaki 2021). The metric is defined as follows, capturing the sum of individual differences from the mean:

$$cv = \sqrt{\frac{1}{n-1} \sum_{i=1}^n \frac{(u_i - \bar{u})^2}{\bar{u}^2}} \quad (4)$$

where  $u_i$  is the utility of agent  $i$ ,  $n$  is the number of agents, and  $\bar{u}$  is the average utility of all agents. The lower  $cv$ , the fairer the system. Jiang and Lu (2019) note that in multi-agent sequential decision-making, optimizing the coefficient of variation of individual agents is often hard because its value depends on the joint policy of all agents. In these cases, each agent  $i$ 's contribution to  $cv$  is usually approximated by  $(u^i - \bar{u})^2 / \bar{u}^2$ . The coefficient of variation is then often built into the reward function of each agent so that each agent is punished for getting too much or too little resources or utility (Jiang and Lu 2019; Elmalaki 2021).

**Q-Value Based Definitions.** A natural idea of fairness in RL is based on Q-values. In the MDP setting, the Q-function  $Q(s, a)$  is the expected utility of taking a given action  $a$  in a given state. In this context,  $Q^*(s, a)$  is the expected utility of taking action  $a$  from state  $s$  and then acting optimally afterwards. Based on this notion, Jabbari et al. (2017) introduce the idea of exact fairness and its two relaxations: approximate-choice fairness and approximate-action fairness. Exact fairness requires that “in state  $s$ , the algorithm never chooses an available action  $a$  with probability higher than another action  $a'$  unless  $Q^*(s, a) > Q^*(s, a')$ , i.e., the long term [(discounted)] reward of  $a$  is higher than that of  $a'$ ” (Jabbari et al. 2017). Approximate-choice fairness “requires that an algorithm never chooses a worse action with probability substantially higher than better actions” (Jabbari et al. 2017). On the other hand, approximate-action fairness “requires that an algorithm never favors an action of substantially lower quality than that of a better action” (Jabbari et al. 2017). Both exact fairness and approximate-choice fairness require exponential time learning algorithms to approach optimality (Jabbari et al. 2017). Further relaxation to the probabilistic

requirement results in a weaker definition of fairness but ensures a polynomial-time learning algorithm.

**$\alpha$ -Fair Utility.** In computer networking, fairness is often defined as fairly allocating network resources (i.e., bandwidth) to different data flows. In this context, Chen, Wang, and Lan (2021) use an  $\alpha$ -fair utility function to capture fairness notions. For  $\alpha \geq 0$ , the  $\alpha$ -fair utility is defined as:

$$U(x) = \begin{cases} x^{1-\alpha}/(1-\alpha) & \text{for } \alpha \neq 1 \\ \log(x) & \text{for } \alpha = 1 \end{cases} \quad (5)$$

Note that this notion of fairness is different from the previous notions of fairness in that this measure can be “toggled” or controlled to achieve different levels of fairness: Setting  $\alpha = 0$ , for example, leads to throughput maximization,  $\alpha = 1$  to proportional fairness, and  $\alpha \rightarrow \infty$  to max-min fairness (Chen, Wang, and Lan 2021). The developer can set the desired fairness parameter  $\alpha$ . In contrast, definitions like the coefficient of variation are fairness measures embedded into the RL problem itself via the reward function. In these cases, the reward function can account for fairness, even though the programmer cannot precisely control how fair the outcomes are. Hao et al. (2023) also use  $\alpha$ -fair utility in the context of computing offloading in edge computing similarly.

**Issue-based Fairness.** Zhu et al. (2023) studies fairness in the context of robot assistant services. Compared to the previously discussed notions of fairness, the authors identify issues in human-robot interactions and define fairness as different possibilities that these issues present given groups of people with and without sensitive identities. They identified

- *willingness issues*, i.e., the issue that a robot is more likely to respond to groups without sensitive identities
- *quality issues*, i.e., the issue that a robot is more likely to be at an uncomfortable distance from groups with sensitive identities
- *priority issues*, i.e., the issue that a robot is more likely to prioritize groups without sensitive identities
- *risk issues*, i.e., the issue that a robot is more likely to be at a risky distance from groups with sensitive identities

The authors calculate the issue-specific score as  $d(g_c) = |\Pr(R = r | P \in g_c) - \Pr(R = r | P \notin g_c)|$ . They then calculate an *overall issue score* as the average of the scores of all issue-specific scores; this average issue score is subsequently included in their approach as a proxy for fairness (see Section ‘Methodology’ for more details).

**Calibrated Fairness.** Liu et al. (2017) introduce smooth and calibrated fairness as two notions of fairness in sequential decision-making problems. Smooth fairness ensures that if two groups (or arms in the context of multi-armed bandit problems) have similar quality distributions, they should be selected with similar probabilities, addressing the issue of minor quality differences unfairly influencing selection. Calibrated fairness, on the other hand, demands that the probability of selecting a group matches the probability that this group has the best quality realization. This

concept ensures meritocracy by requiring that dissimilar groups be treated in accordance with their differences in quality, thus addressing both the need for equality among similarly qualified groups and the requirement to recognize and reward merit when differences are significant.

**Regularized Maximin Fairness.** Zhang and Shah (2014) introduce a new fairness criterion specifically designed for multi-agent MDPs, where agents may have conflicting interests. The core idea behind this fairness criterion is to enhance equity among agents by focusing on the least advantaged positions. It seeks to maximize the minimum (or worst) performance across all agents, thereby ensuring that the most disadvantaged agent’s outcome is improved as much as possible. This maximization is not done in isolation but is regularized with an overall performance consideration, aiming for a balanced approach that uplifts the worst-off agent without excessively compromising the collective outcome. This approach is reflective of a Pareto-efficient solution, where improvements can be made to at least one agent’s situation without making any other agent worse off (Zhang and Shah 2014).

**Observations on Fairness Definitions.** In general, we notice that definitions of fairness in RL are highly inconsistent, some of them are even mutually exclusive, which has been previously discussed in the literature (Berk et al. 2017; Friedler, Scheidegger, and Venkatasubramanian 2016). This inconsistency arises due to several reasons. First, the majority of the existing literature is focused on incorporating fairness in RL in one specific application domain. For example, the work of Elmalaki (2021) focuses on incorporating fairness in RL for IoT devices, while Chen, Wang, and Lan (2021) studies fairness in RL for network utility optimization. This approach results in many fairness approaches being highly specialized to certain contexts that may not work well in more general, context-agnostic RL settings. Second, many notions of fairness originated from fields outside of ML. For example, the notion of  $\alpha$ -fairness was originally studied in the context of networks while welfare economics definitions have their origin in social welfare and game theory. This causes a discrepancy in definitions adopted in RL, since there has neither been an agreement on a general fairness definition in other fields nor in ML or RL – and there likely never will be, given that the most suitable fairness definition strongly depends on compromises and context of application. However, a structured approach to decide on a fairness definition to be used in a specific RL application is not available, either, leading to often arbitrary choices of fairness definitions, without sufficient justification. While Mandal and Gan (2022) outline four axioms that a fairness measure should satisfy, there’s more work necessary on how to decide which fairness measure is appropriate given an RL problem.

### Application Domains

Fairness in RL is most relevant in cases where multiple stakeholders use a system, or its decisions will impact multiple users. Ensuring fairness, or at least the attempt

to ensure fairness under generally accepted definitions, is crucial for the users’ fair treatment and acceptance of these systems. This chapter presents specific domains in which fair RL has been studied.

**Recommendation Systems.** Liu et al. (2020) developed a model to balance fairness and accuracy in interactive RS. These systems are being used to recommend items of interest (e.g., news, movies, or articles (Steck, van Zwol, and Johnson 2015)) to individual users. The recommendations are updated in an online setting based on user feedback, often expressed as taking a desired action or not, such as buying an article after seeing a recommendation. The latter is also known as conversion rate. Only using the conversion rate as the single objective can lead to an imbalance across different demographic groups, which can lead to minorities being ignored in recommendations. Ge et al. (2022) identified the same issue in RS and used a fairness-aware MOMDP approach to optimally balance these two objectives (see Section ‘Methodology’). Other work by Huang et al. (2022) focused on fair personalized RS utilizing contextual bandits while Huang, Zhang, and Wu (2022) used causal bandits for fair recommendations. Similar work on fair RL for recommendations was done by Singh et al. (2020) and Wang et al. (2021a). For a broader overview on fairness in RS, we refer the reader to Li et al. (2023) and Zhang and Wang (2021).

**Work Distribution.** Claire et al. (2019) looked at fair RL in the context of resource distribution in human-robot teams. They specifically focused on a fair candidate selection for a work task because in an unconstrained RL setting, the agent would learn the individual worker’s performance first and then assign as many tasks as possible to the highest-performing candidate. In real life, this can lead to a two-fold issue: Firstly, when a human worker is allocated significantly more resources to process than his team members, he might burn out quickly, which is a factor that is not being taken into account by the agent by default. Secondly, favoring one employee might hurt all team members’ motivation. Perceived inequalities have been shown to motivate people to act against their self-interest to eliminate the inequality (Camerer 2003), sometimes including actions to retaliate (Skarlicki and Folger 1997). It further undermines their trust in the system due to a perceived unfairness. All these aspects can negatively impact overall performance if fairness is not explicitly considered in the system.

**Scheduling and Resource Distribution.** Chen, Wang, and Lan (2021) look at fair RL solutions in the context of “wireless network scheduling and Quality-of-Experience (QoE) optimization in video streaming”. The former is concerned with scheduling the transmission and reception of sequences of network packages across a network of users. In this case, fairness considerations are necessary to prevent the discrimination of certain parties and, in the worst case, the restriction of access to the network. Yamazaki and Yamamoto (2021) also look into fair RL for wireless network scheduling. Fairness regarding QoE, on

the other hand, means that all users have a similar, e.g., video streaming experience based on metrics to capture the (perceived) video quality by users. Fair RL in the context of QoE has also been addressed by Arani, Hu, and Zhu (2021); Tong et al. (2021) and Petrangeli et al. (2014). Similarly, Hao et al. (2023) addressed the challenge of computing offloading in edge computing, where balancing service provision at network edges with limited resources is crucial for user experience. Other work that pertains to fair RL in scheduling and resource distribution was done by Agarwal and Aggarwal (2019), Comşa et al. (2019), and Arianpoo and Leung (2016). Another area in this context is Unmanned Aerial Vehicles (UAV) Control. UAVs serve as dynamic flying communication platforms, enabling quick and flexible deployment with the ability to establish direct Line-of-Sight links, yet face significant challenges such as limited battery life and coverage area, alongside high costs Qi et al. (2020). RL can be employed to optimize UAV control for efficient coverage and communication, ensuring fairness in resource allocation across the targeted regions despite these constraints (Qi et al. 2020). This setting has further been explored by Liu et al. (2018).

**Infrastructure Control.** Siddique, Weng, and Zimmer (2020) and Zimmer et al. (2021) consider applications to traffic light control and data center control. In the data center control setting, the goal is to minimize the queue length of each network switch, which must be fair. In the traffic light control problem, without fairness constraints, the problem seeks to minimize the total wait time across all lanes. By incorporating fairness considerations, the agent is then tasked to learn how to optimize the expected wait time per road. Athalye and Nayak (2021); Raeis and Leon-Garcia (2021); Valkanis et al. (2022) and Li et al. (2020) also looked into fair RL-based ground traffic control. de Arruda et al. (2010) use multi-agent RL for fairness-aware air traffic flow control.

**Internet-of-Things (IoT).** In the area of IoT, fairness considerations could enable IoT devices to mitigate better the challenges brought by intra-human, inter-human, and multi-human variability. Take smart thermometers, for example. The same user’s temperature preference might change over time (intra-human variability), each user has different temperature preferences (inter-human variability), and multiple people in the same room can have a wide range of temperature preferences (multi-human variability). Elmalaki (2021) show that using their developed RL-based “FaiR-IoT fairness-aware human-in-the-loop framework” will improve the user experience and improve fairness (measured by the coefficient of variation, discussed in Section ‘Definitions of Fairness in RL’).

**Robotics.** A key area for the use of RL is to train robots to interact in the real world. For example, robots can be used to provide services in airports or restaurants to humans; fair RL in this context was studied by Zhu et al. (2023) to ensure that humans impacted by the robot are treated fairly. On the other hand, Zhu and Oh (2018) researched fairness

pertaining to RL-based multi-robot coordination.

**Other Areas.** For completeness, we also want to mention less prominent application areas for fair RL here. For example, Atwood et al. (2019) studied the so-called precision disease control problem, i.e., the optimal allocation of vaccines in a social network, using fair RL. Bao (2019) used a multi-agent RL approach for fair stock trading. Finally, Wang and Deng (2020) used RL to mitigate bias in facial recognition.

## Methodology

Current literature on fairness in RL can be categorized into research that focuses on single- or multi-agent setups. We will discuss these approaches in the following chapters. For a general overview on multi-agent RL, we refer the reader to Busoniu, Babuska, and De Schutter (2006).

### Single-agent RL

In the single-agent multi-objective setting studied by Siddique, Weng, and Zimmer (2020), the authors integrate the GGF (see Section ‘Definitions of Fairness in RL’) into their MOMDP formulation. The problem of solving for an optimal policy that ensures a fair distribution of rewards to  $D$  users then becomes

$$\operatorname{argmax}_{\pi} \operatorname{GGF}_w(J(\pi)) \quad (6)$$

where  $J(\pi)$  is the discounted reward (Siddique, Weng, and Zimmer 2020). Note that because the GGF is a non-linear function, this is a non-linear convex optimization problem (Siddique, Weng, and Zimmer 2020). To solve it, a modified Deep Q Network (DQN) or policy gradient methods could be used. The latter can be advantageous because they directly optimize for the desired objective function and can also learn stochastic policies instead of just deterministic ones.

Liu et al. (2020) propose the RL-based *FairRec* framework to ensure a balanced long-term trade-off between accuracy and fairness in a single-agent MDP setup in RS. The authors use the weighted proportional fairness notion described in Section ‘Definitions of Fairness in RL’. (Liu et al. 2020) build an actor-critic architecture where the actor network is responsible for dynamic recommendations depending on the fairness status and the user preferences. On the other hand, the critic network encourages or discourages a recommended item based on its value estimate of the actor’s output. To accomplish a focus on both accuracy and fairness, Liu et al. (2020) design a two-fold reward based on a personalized fairness-aware state representation. They consider whether a user performed a desired activity on a recommended item before they evaluate the fairness gain of the activity.

Ge et al. (2022) introduce a fairness-aware recommendation framework called MoFIR, which utilizes a MOMDP setting to find an optimal balance between conflicting objectives of utility and fairness. MoFIR is designed to learn a single parametric representation that can adapt to

different preferences over these objectives, accommodating the varying needs of real-world e-commerce platforms (Ge et al. 2022). The key innovation in MoFIR is the modification of the Deep Deterministic Policy Gradient (DDPG) algorithm by incorporating conditioned networks (CN) (Ge et al. 2022). These networks are conditioned directly on the preferences for utility and fairness, allowing the system to output Q-value vectors that represent the optimal policy for any given preference mix (Ge et al. 2022). Constrained optimization problem formulations as by, e.g., Nabi, Malinsky, and Shpitser (2019), aim to maintain a certain level of fairness while optimizing for the main performance objective. However, these approaches often compromise recommendation accuracy for fairness. Ge et al. (2022) overcome this limitation with their approach.

Claire et al. (2019), on the other hand, consider a stochastic MAB framework together with an unconstrained upper confidence bound (UCB) algorithm in their work. The aim for an agent in the context of an MAB is to maximize cumulative rewards by pulling bandits’ arms based on previous information the agent obtained. In the UCB algorithm, the agent does this by using the information on the number of times an arm was pulled and the average empirical rewards the agent received to estimate the expected reward of each arm. However, the authors find that an unconstrained UCB algorithm fails to ensure fairness since the agent will not use under-performing arms after a certain number of time steps. To prevent this, Claire et al. (2019) propose two adjusted UCB algorithms: a strict-rate-constrained and a stochastic-rate-constrained version. The former guarantees<sup>2</sup> that each lever has a minimum pull rate because arms are pre-scheduled to be pulled in fixed time slots. In all other time slots, the agent will follow the standard UCB algorithm, i.e., choose the best arm according to the benchmark strategy. The stochastic-rate-constrained UCB algorithm, on the other hand, only guarantees<sup>3</sup> that the expected pulling rate is at least the minimum pull rate at any time. Compared with the strict-rate-constrained UCB algorithm, the authors introduce randomness, thereby ensuring that the probability that the agent pulls an arm at time  $t$  is equivalent to the minimum pull rate. In this case, with a probability of  $1 - Kv$ , where  $K$  is the arm set and  $v$  is the minimum pull rate, the benchmark UCB policy is followed.

Looking at other fairness in MAB problem formulations, Chen et al. (2020) utilize contextual information about users and tasks, allowing their algorithm to dynamically adjust its allocation strategy to satisfy fairness constraints while still aiming to maximize a given performance metric. This algorithm uniquely does not assume a fixed distribution for the performance losses associated with each user, as did previous related work by Patil et al. (2021); Li, Liu, and Ji (2019) and Claire et al. (2019), making it adaptable to non-stationary and adversarial environments. Ghodsi and Mir-

fakhar (2021) also explore fairness in MAB settings. The authors employ the EXP\_3 algorithm to address the online cake-cutting problem with non-contiguous pieces. This approach allows for sub-linear fairness and revenue regret by balancing exploration and exploitation in an environment where agents’ valuations change. The authors’ methodology involves running two separate algorithms independently to achieve sub-linear regret for revenue and fairness, then merging these algorithms to tackle both objectives simultaneously. Finally, Gillen et al. (2018) present an algorithm for online learning in linear contextual bandits settings with individual fairness constraints governed by an unknown Mahalanobis similarity metric. The algorithm learns from weak feedback identifying fairness violations without quantifying their extent, akin to regulator interventions that can recognize unfairness but cannot define it quantitatively. The main result is an adversarial-context algorithm that limits the number of fairness violations logarithmically over time  $T$  while obtaining optimal regret bounds concerning the best fair policy.

## Multi-agent RL

In the context of multi-agent fair RL, Chen, Wang, and Lan (2021) take the idea by Liu et al. (2020) to the multi-agent setting and aim to optimize general fairness utility functions in actor-critic RL. They specifically developed a method to adjust the standard RL rewards by a multiplicative weight that considers the history of rewards and the shape of the fairness utility. The multiplicative adjustment is defined using a uniformly continuous function  $\Phi(h_{\pi,t})$  that is dependent on a statistic  $h_{\pi,t}$  capturing past rewards. The adjusted rewards are calculated using  $\hat{r}_{k,t} = r_{k,t} \cdot \phi(h_{\pi,t})$ .

The authors then employ a rewards-adjusted actor-critic architecture to address the issue that a non-linear  $\alpha$ -fair utility ( $\forall \alpha > 0$ ) does not satisfy the Markovian property, which would be necessary to formulate  $\alpha$ -fair network utility optimization as an MDP and guarantee convergence under the Policy Gradient Theorem (Chen, Wang, and Lan 2021). Instead, their approach guarantees that, given a proper choice of  $\Phi$  and  $h$ , it converges to “at least a stationary point of the  $\alpha$ -fair utility optimization” (Chen, Wang, and Lan 2021). Another advantage to this approach is that it builds on actor-critic architectures, meaning that the optimization converges quicker because of variance reduction and because it is not based on the Monte Carlo method (Chen, Wang, and Lan 2021), especially when optimizing in large state/action spaces.

Agarwal and Aggarwal (2019) propose both model-based and model-free algorithms to optimize a non-linear function of long-term average rewards in multi-agent contexts. Their model-based algorithm utilizes posterior sampling with a Dirichlet distribution, showing convergence to an optimal point under certain conditions and achieving a sub-linear regret bound in the number of time steps and objectives. The paper also introduces a model-free policy gradient algorithm that can be implemented efficiently using neural networks. These methods are designed to overcome the limitations of

<sup>2</sup>See Claire et al. (2019) for formalized theoretical guarantees.

<sup>3</sup>See Claire et al. (2019) for formalized theoretical guarantees.

traditional RL approaches, such as SARSA, Q-Learning, and DQN, which are sub-optimal for non-linear, multi-agent optimization problems.

Arianpoo and Leung (2016) use a distributed, cross-layer approach leveraging Q-learning to monitor network dynamics and adjust TCP parameters autonomously for fair resource allocation. Each node hosting a TCP source models the system’s state as an MDP. This approach allows each node to autonomously adjust TCP parameters without central control or the need for control message exchanges among nodes. The authors emphasize the suitability of Q-learning for this application due to its model-free nature, enabling it to find optimal strategy-selection policies for any given finite state MDP, particularly under the dynamic conditions of wireless networks.

In the multi-agent setup researched by Zimmer et al. (2021), the authors formulate the problem of cooperative multi-agent RL as

$$\max_{\theta} \phi(J(\theta)) \quad (7)$$

where  $\theta$  is the joint policy of all the agents and  $J_k(\theta) = E_{\theta}[\sum_k \gamma^t r_{k,t}]$  the expected sum of discounted rewards of user  $k$ , and  $J(\theta) = \sum_k J_k(\theta)$ . The algorithmic solution for this optimization problem includes a policy gradient approach implemented in an actor-critic architecture. Zimmer et al. (2021) propose “Self-Oriented Team-Oriented (SOTO) networks updated by dedicated policy gradients”. On a high level, the network includes “a self-oriented and a team-oriented policy” (Zimmer et al. 2021). The former optimizes for an individual policy, whereas the team-oriented policy optimizes for the SWF  $\phi(J(\theta))$ , i.e., the two sub-networks focus on efficiency and equity, respectively. An advantage of the approach by Zimmer et al. (2021) is that it is not domain-specific and allows for the adoption of a variety of fairness notions since it only demands a (sub-)differentiable welfare function (Zimmer et al. 2021).

Jiang and Lu (2019), on the other hand, introduce the Fair-Efficient Network (FEN), an RL-based model that makes use of a “fair-efficient reward” (Jiang and Lu 2019) to address multi-agent RL settings. Each agent learns this reward to optimize its own policy. Additionally, an average consensus among agents as part of the fair-efficient rewards allows coordination between agents’ policies. Specifically, the authors consider a setting with  $n$  agents and limited, commonly-accessible resources. Each agent’s fair-efficient reward at time  $t$  is:

$$\hat{r}_t^i = \frac{\bar{u}_t/c}{\epsilon + |u_t^i/\bar{u}_t - 1|} \quad (8)$$

where the utility of agent  $i$  at time step  $t$  is the average reward over time it received, i.e.,  $u_t^i = \frac{1}{t} \sum_{j=0}^t r_j^i$ , and  $c$  is the maximum environmental reward an agent can obtain at a time step. Hence,  $\bar{u}_t/c$  can be seen as a proxy for system efficiency since it represents the system-wide

resource allocation. On the other hand,  $|u_t^i/\bar{u}_t - 1|$  represents an agent’s utility deviation from the average, which is taken as a proxy for fairness based on the idea of the coefficient of variation (see Section ‘Definitions of Fairness in RL’). Each agent uses this reward to learn its own policy  $F_i = \mathbb{E}[\sum_{t=0}^{\infty} \gamma^t \hat{r}_t^i]$  with a discount  $\gamma$ . Potential multi-objective conflicts are circumvented by using a hierarchical RL model with a controller that maximizes the fair-efficient reward by changing between multiple sub-policies. These sub-policies are designed with respect to different goals: to maximize an environmental reward and to explore different fair behaviors. This approach aims to enable agents to learn efficiency and fairness simultaneously.

Similarly, Zhu et al. (2023) study a multi-agent setting and employ a self-reflecting approach that allows agents to “self-identify biased behavior during interactions with humans” (Zhu et al. 2023); they call this Fairness-Sensitive Policy Gradient Reinforcement Learning (FSPGRL). Based on the issue score  $I$  explained in Section ‘Definitions of Fairness in RL’, the authors define a reward penalty parameter  $\tau_{penalty}$  that captures the tolerance for the identified biases. They then define the reward as  $R_t = \tau_{penalty} \sum (1 - I(g_c))$ , where  $g_c$  is the group of people with sensitive attributes. They combine the REINFORCE algorithm, a policy gradient method, and Proximal Policy Optimization (PPO) to mitigate the bias and allow the agent to self-correct biased behavior. The agent itself is based on an actor-critic network. FSPGRL subsequently “identifies bias by examining the abnormal update along particular gradients and updates the policy network to support fair decision-making” (Zhu et al. 2023).

Grupe, Selman, and Lee (2022) introduce a different multi-agent angle to fair RL: team-based fairness, “a group-based fairness measure that requires equitable reward distributions across sensitive groups” within a team. The authors subsequently introduce Fairness through Equivariance, a strategy that transforms “the team’s joint policy into an equivariant map” (Grupe 2023) to enforce team fairness during policy optimization. They further introduce “Fairness through Equivariance Regularization as a soft-constraint version of [Fairness through Equivariance]” (Grupe, Selman, and Lee 2022), which allows for tuning fairness constraints over multi-agent policies through regularization. The latter aims to balance fairness and utility by adjusting the weight of equivariance regularization.

Finally, Hao et al. (2023) formulate the setting of computing offloading in edge computing as a “long-term average optimization problem to maximize an  $\alpha$ -fair utility function” under computing capacity and storage constraints and approach it as an MDP. To solve it, they propose a multi-update deep RL algorithm that optimizes for fairness and efficiency.

## Trade-Offs

Imposing fairness requirements to RL algorithms often results in worse running time compared to problems



without fairness constraints. Jabbari et al. (2017) show the trade-offs between efficiency and fairness. In particular, both fair and approximate-choice fair requirements impose an exponential time step  $\mathcal{T} = \Omega(k^n)$ . The authors’ defined approximate-action fairness (see Section ‘Methodology’) requires steps  $\mathcal{T} = \Omega(k^{1/(1-\gamma)})$ , where  $n$  is the size of the state space and  $\gamma$  is the discount factor. Note that without fairness constraints, standard RL algorithms learn an “ $\epsilon$ -optimal policy in a number of steps polynomial in  $n$ ,  $1/\epsilon$ ” (Jabbari et al. 2017), and all parameters of the MDP. This example shows that imposing fairness requirements potentially comes with a cost regarding run-time efficiency.

Imposing fairness requirements further decreases performance. Liu et al. (2020) discussed this trade-off in their paper. The authors find that their *FairRec* framework (see Section ‘Methodology’) does increase fairness in RS while maintaining a good recommendation quality but that, compared to non-fair baselines, there are still performance losses. Ge et al. (2022) also identified this trade-off as an issue; by proposing a Pareto-optimal approach, they aim to provide optimal fairness-performance combinations and give users the option of choosing whichever optimal performance-fairness combination suits their application context best. The trade-off between fairness and performance is a general issue in designing fair RL systems since any fair solution will introduce additional constraints and/or objectives that will, compared to an unconstrained, non-fair problem formulation, result in a decreased performance quality (Berk et al. 2017). This trade-off is an ongoing challenge in ML more broadly (Wang et al. 2021b). However, since the perceived fairness of a system can impact human performance or trust, not considering fairness considerations can lead to an overall sub-optimal performance (Elmalaki 2021). To better understand these trade-offs, Cimpean et al. (2024) develop a performance-fairness trade-off framework that allows users to explore “obtainable performance-fairness trade-offs, [that] can be used by stakeholders to select the best policy for the problem at hand.” (Cimpean et al. 2024),

## Future Research Directions

The study of fairness in RL has been and will continue to be an interdisciplinary field involving economics, mathematics, computer science, philosophy, and political science. As the deployment of RL-enabled systems accelerates, the different communities need to address the following gaps and challenges.

**Fairness in RLHF.** RLHF is a method for optimizing language models based on human feedback by integrating and preferences into the learning process (Kaufmann et al. 2023). Traditionally, this is being done by having a single reward model that is derived from the given preference data (Chakraborty et al. 2024). In implementing fairness in the context of RLHF, developers must consider three dimensions: 1) the representation of the human feedback providers to encompass a broad range of demographics, cultures, and

ethical viewpoints, 2) how these differing preferences can fairly be integrated to fine-tune the reward mechanisms of the agent, and 3) ensuring that the subsequent model not only maximizes its primary objectives but also takes into account a fairness objective. All three dimensions are understudied. On a conceptual level, Liu (2023) looked at how RLHF can influence society more broadly. Beyond that, Zhong et al. (2024) published theoretical work on explicitly modelling diverse perspectives and integrating fairness concerns in an offline RL setting. Future work is necessary to understand the empirical performance and scalability of Zhong et al. (2024)’s work along with research into incorporating fairness in a multi-party RLHF in an online setting, as well as extending the analysis to more general function approximation classes for modeling individual rewards. Finally, Chakraborty et al. (2024) propose MaxMin-RLHF, which learns a mixture of reward models to capture diverse human preferences and optimizes a max-min objective over user group utilities to achieve socially fair alignment of language models. Important extensions of this work would be evaluating Chakraborty et al. (2024)’s approach on real-world human preference data with organic diversity instead of just synthetic preference data, as well as exploring alternative social welfare objectives.

**Cross-Domain Fair RL Approach.** As discussed in Section ‘Definitions of Fairness in RL’, the diversity of fairness definitions in the field of fair RL are considerable. In most works we surveyed, the definitions were derived from domain-specific requirements. To advance fair RL algorithms, we propose investigating the merits and drawbacks of each of these definitions to find a more general definition that can be applied across domains or a structured approach to determining which fairness definition is appropriate for a specific context. To do so, one could, for example, expand the approach by (Zimmer et al. 2021), which accepts different (sub-)differentiable notions of fairness, and test the effects of different fairness definitions on the performance of subsequent models. Defining fairness across contexts is a more broad open-problem in ML, so making progress in either ML or RL could advance the field. Frameworks that can be used with multiple fairness definitions could also help to study the trade-offs between fairness definitions or achieve multiple fairness requirements at the same time<sup>4</sup>; e.g., Cimpean et al. (2023) outline a fairness-definition-agnostic framework but do not validate it or show its effectiveness. In the bandit context, Metevier et al. (2019)’s work is promising for accepting multiple fairness definitions in an offline contextual bandit algorithm. A related open challenge is fair RL methods that can be applied across domains. While developing domain-specific fair RL algorithms ensures a high degree of applicability to the respective context, we recommend investigating more general models that can be applied across domains. The approaches taken by (Zimmer et al.

<sup>4</sup>This cannot be achieved for all known fairness definitions since some of them are mutually exclusive as shown by Miconi (2017).

2021), (Liu et al. 2020), and (Chen, Wang, and Lan 2021) are promising in this direction but should be tested for effectiveness in a more diverse set of domains to verify their applicability and robustness. Finally, another interesting approach that has not extensively been explored in an RL context yet is personalized fairness (Li et al. 2021), where users can articulate their unique expectations of fairness. This approach prevents the system from solely relying on the designers' interpretations to navigate through conflicting definitions of fairness.

**Additional Application Domains.** While we advocate for the focus on general cross-domain fair RL algorithms, we further encourage research in application-specific settings to find optimal, context-specific models, especially in high-stakes contexts to ensure the highest possible performance. Potential domains to investigate are routing, traffic light control systems, and cloud computing (Jiang and Lu 2019).

**Fairness over Time.** In the existing fair-RL literature, fairness has often been treated as a final objective to be fulfilled or maximized at the final time step  $T$ . However, ensuring fairness at any point during the interaction should become a research focus, especially for continuous RL applications. For example, the research by Henzinger et al. (2023) could be expanded to the RL setting. We further recommend studying the differences between ensuring fairness at a given time and ensuring fairness in long-term decision-making settings. Specifically, approaches to guarantee fairness not only at the final time step  $T$  but throughout the runtime should be investigated.

**Adversarial Fairness.** As RL systems are increasingly used in potentially adversarial settings, the imperative to design learning policies that are both fair and robust becomes critical. Hence, future research should focus on developing approaches that can identify and counteract bias-inducing adversarial attacks while also ensuring equitable treatment of all entities within an RL setting.

## Conclusion

Given the increased adoption of RL in socio-technical systems, supporting – or even guaranteeing – fairness in these algorithms is more important than ever. In this paper, we discussed research that focuses on developing fair RL approaches. We compared different definitions of fairness, showcased the methodologies pursued in the different papers, and proposed future research directions. With our survey, we hope to build a discussion base and encourage further research in the field.

## Acknowledgements

We thank Prof. Michael Kearns for his valuable feedback on the initial idea for this paper.

## References

- Agarwal, M.; and Aggarwal, V. 2019. Reinforcement learning for joint optimization of multiple rewards. *arXiv preprint arXiv:1909.02940*.
- Arani, A. H.; Hu, P.; and Zhu, Y. 2021. Fairness-aware link optimization for space-terrestrial integrated networks: A reinforcement learning framework. *IEEE Access*, 9: 77624–77636.
- Arianpoo, N.; and Leung, V. C. 2016. How network monitoring and reinforcement learning can improve tcp fairness in wireless multi-hop networks. *EURASIP Journal on Wireless Communications and Networking*, 2016: 1–15.
- Athalye, A.; and Nayak, S. 2021. Fairness and robustness of mixed autonomous traffic control with reinforcement learning.
- Atwood, J.; Srinivasan, H.; Halpern, Y.; and Sculley, D. 2019. Fair treatment allocations in social networks. *arXiv preprint arXiv:1911.05489*.
- Bao, W. 2019. Fairness in multi-agent reinforcement learning for stock trading. *arXiv preprint arXiv:2001.00918*.
- Berk, R.; Heidari, H.; Jabbari, S.; Joseph, M.; Kearns, M.; Morgenstern, J.; Neel, S.; and Roth, A. 2017. A convex framework for fair regression. *arXiv preprint arXiv:1706.02409*.
- Bommasani, R.; Hudson, D. A.; Adeli, E.; Altman, R.; Arora, S.; von Arx, S.; Bernstein, M. S.; Bohg, J.; Bosselut, A.; Brunskill, E.; et al. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.
- Busoniu, L.; Babuska, R.; and De Schutter, B. 2006. Multi-agent reinforcement learning: A survey. In *2006 9th International Conference on Control, Automation, Robotics and Vision*, 1–6. IEEE.
- Camerer, C. F. 2003. Behavioral game theory: Plausible formal models that predict accurately. *Behavioral and Brain Sciences*, 26(2): 157–158.
- Chakraborty, S.; Qiu, J.; Yuan, H.; Koppel, A.; Huang, F.; Manocha, D.; Bedi, A. S.; and Wang, M. 2024. MaxMin-RLHF: Towards Equitable Alignment of Large Language Models with Diverse Human Preferences. *arXiv preprint arXiv:2402.08925*.
- Chen, J.; Wang, Y.; and Lan, T. 2021. Bringing Fairness to Actor-Critic Reinforcement Learning for Network Utility Optimization. In *IEEE INFOCOM 2021-IEEE Conference on Computer Communications*, 1–10. IEEE.
- Chen, Y.; Cuellar, A.; Luo, H.; Modi, J.; Nemlekar, H.; and Nikolaidis, S. 2020. The fair contextual multi-armed bandit. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*.
- Cimpean, A.; Jonker, C. M.; Libin, P. J. K.; and Nowe, A. 2024. A Group And Individual Aware Framework For Fair Reinforcement Learning. In *The Sixteenth Workshop on Adaptive and Learning Agents*.
- Cimpean, A.; Libin, P.; Coppens, Y.; Jonker, C.; and Nowé, A. 2023. Towards Fairness In Reinforcement Learning. In *Proceedings of the Adaptive and Learning Agents Workshop (ALA 2023)*, 1–5.

- Claire, H.; Chen, Y.; Modi, J.; Jung, M.; and Nikolaidis, S. 2019. Reinforcement learning with fairness constraints for resource distribution in human-robot teams. *arXiv preprint arXiv:1907.00313*.
- Comşa, I.-S.; Zhang, S.; Aydin, M.; Kuonen, P.; Trestian, R.; and Ghinea, G. 2019. A comparison of reinforcement learning algorithms in fairness-oriented OFDMA schedulers. *Information*, 10(10): 315.
- Corbett-Davies, S.; and Goel, S. 2018. The measure and mismeasure of fairness: A critical review of fair machine learning. *arXiv preprint arXiv:1808.00023*.
- de Arruda, A. C.; Leite, A. F.; de Almeida, C. R.; Crespo, A. M.; and Weigang, L. 2010. Fairness analysis with cost impact for Brasilia’s Flight Information Region using reinforcement learning approach. In *13th International IEEE Conference on Intelligent Transportation Systems*, 539–544. IEEE.
- Elmalaki, S. 2021. FaiR-IoT: Fairness-aware Human-in-the-Loop Reinforcement Learning for Harnessing Human Variability in Personalized IoT. In *Proceedings of the International Conference on Internet-of-Things Design and Implementation*, 119–132.
- Friedler, S. A.; Scheidegger, C.; and Venkatasubramanian, S. 2016. On the (im) possibility of fairness. *arXiv preprint arXiv:1609.07236*.
- Gajane, P.; Saxena, A.; Tavakol, M.; Fletcher, G.; and Pechenizkiy, M. 2022. Survey on fair reinforcement learning: Theory and practice. *arXiv preprint arXiv:2205.10032*.
- Ge, Y.; Zhao, X.; Yu, L.; Paul, S.; Hu, D.; Hsieh, C.-C.; and Zhang, Y. 2022. Toward Pareto efficient fairness-utility trade-off in recommendation through reinforcement learning. In *Proceedings of the fifteenth ACM international conference on web search and data mining*, 316–324.
- Ghods, M.; and Mirfakhar, A. 2021. Online Fair Revenue Maximizing Cake Division with Non-Contiguous Pieces in Adversarial Bandits. *arXiv preprint arXiv:2111.14387*.
- Gillen, S.; Jung, C.; Kearns, M.; and Roth, A. 2018. Online learning with an unknown fairness metric. *Advances in neural information processing systems*, 31.
- Gruppen, N. A. 2023. *Responsible Emergent Multi-Agent Behavior*. Ph.D. thesis, Cornell University.
- Gruppen, N. A.; Selman, B.; and Lee, D. D. 2022. Cooperative multi-agent fairness and equivariant policies. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36 (9), 9350–9359.
- Hacker, P. 2018. Teaching fairness to artificial intelligence: Existing and novel strategies against algorithmic discrimination under EU law. *Common Market Law Review*, 55(4).
- Hao, H.; Xu, C.; Zhang, W.; Yang, S.; and Muntean, G.-M. 2023. Computing Offloading with Fairness Guarantee: A Deep Reinforcement Learning Method. *IEEE Transactions on Circuits and Systems for Video Technology*.
- Henzinger, T.; Karimi, M.; Kueffner, K.; and Mallik, K. 2023. Runtime Monitoring of Dynamic Fairness Properties. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 604–614.
- Huang, W.; Labille, K.; Wu, X.; Lee, D.; and Heffernan, N. 2022. Achieving user-side fairness in contextual bandits. *Human-Centric Intelligent Systems*, 2(3-4): 81–94.
- Huang, W.; Zhang, L.; and Wu, X. 2022. Achieving counterfactual fairness for causal bandit. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36(6), 6952–6959.
- Jabbari, S.; Joseph, M.; Kearns, M.; Morgenstern, J.; and Roth, A. 2017. Fairness in reinforcement learning. In *International conference on machine learning*, 1617–1626. PMLR.
- Jiang, J.; and Lu, Z. 2019. Learning fairness in multi-agent systems. *Advances in Neural Information Processing Systems*, 32.
- Kaufmann, T.; Weng, P.; Bengs, V.; and Hüllermeier, E. 2023. A Survey of Reinforcement Learning from Human Feedback. *arXiv preprint arXiv:2312.14925*.
- Krilašević, S.; and Grammatico, S. 2021. Learning generalized Nash equilibria in multi-agent dynamical systems via extremum seeking control. *Automatica*, 133: 109846.
- Kushner, H. J.; and Whiting, P. A. 2004. Convergence of proportional-fair sharing algorithms under general conditions. *IEEE transactions on wireless communications*, 3(4): 1250–1259.
- Lattimore, T.; and Szepesvári, C. 2020. *Bandit algorithms*. Cambridge University Press.
- Li, C.; Ma, X.; Xia, L.; Zhao, Q.; and Yang, J. 2020. Fairness control of traffic light via deep reinforcement learning. In *2020 IEEE 16th International Conference on Automation Science and Engineering (CASE)*, 652–658. IEEE.
- Li, F.; Liu, J.; and Ji, B. 2019. Combinatorial sleeping bandits with fairness constraints. *IEEE Transactions on Network Science and Engineering*, 7(3): 1799–1813.
- Li, Y.; Chen, H.; Xu, S.; Ge, Y.; Tan, J.; Liu, S.; and Zhang, Y. 2023. Fairness in Recommendation: Foundations, Methods, and Applications. *ACM Transactions on Intelligent Systems and Technology*, 14(5): 1–48.
- Li, Y.; Chen, H.; Xu, S.; Ge, Y.; and Zhang, Y. 2021. Towards personalized fairness based on causal notion. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1054–1063.
- Liu, C. H.; Chen, Z.; Tang, J.; Xu, J.; and Piao, C. 2018. Energy-efficient UAV control for effective and fair communication coverage: A deep reinforcement learning approach. *IEEE Journal on Selected Areas in Communications*, 36(9): 2059–2070.
- Liu, G. K.-M. 2023. Transforming Human Interactions with AI via Reinforcement Learning with Human Feedback (RLHF). *Massachusetts Institute of Technology*.
- Liu, W.; Liu, F.; Tang, R.; Liao, B.; Chen, G.; and Heng, P. A. 2020. Balancing between accuracy and fairness for interactive recommendation with reinforcement learning. *Advances in Knowledge Discovery and Data Mining*, 12084: 155.

- Liu, Y.; Radanovic, G.; Dimitrakakis, C.; Mandal, D.; and Parkes, D. C. 2017. Calibrated fairness in bandits. *arXiv preprint arXiv:1707.01875*.
- Mandal, D.; and Gan, J. 2022. Socially fair reinforcement learning. *arXiv preprint arXiv:2208.12584*.
- Mehrabi, N.; Morstatter, F.; Saxena, N.; Lerman, K.; and Galstyan, A. 2021. A survey on bias and fairness in machine learning. *ACM computing surveys (CSUR)*, 54(6): 1–35.
- Metevier, B.; Giguere, S.; Brockman, S.; Kobren, A.; Brun, Y.; Brunskill, E.; and Thomas, P. S. 2019. Offline contextual bandits with high probability fairness guarantees. *Advances in neural information processing systems*, 32.
- Miconi, T. 2017. The impossibility of “fairness”: a generalized impossibility result for decisions. *arXiv preprint arXiv:1707.01195*.
- Nabi, R.; Malinsky, D.; and Shpitser, I. 2019. Learning optimal fair policies. In *International Conference on Machine Learning*, 4674–4682. PMLR.
- Patil, V.; Ghalme, G.; Nair, V.; and Narahari, Y. 2021. Achieving fairness in the stochastic multi-armed bandit problem. *The Journal of Machine Learning Research*, 22(1): 7885–7915.
- Petrangeli, S.; Claeys, M.; Latré, S.; Famaey, J.; and De Turck, F. 2014. A multi-agent Q-Learning-based framework for achieving fairness in HTTP Adaptive Streaming. In *2014 IEEE Network Operations and Management Symposium (NOMS)*, 1–9. IEEE.
- Pirotta, M.; Parisi, S.; and Restelli, M. 2015. Multi-objective reinforcement learning with continuous pareto frontier approximation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 29 (1).
- Qi, H.; Hu, Z.; Huang, H.; Wen, X.; and Lu, Z. 2020. Energy Efficient 3-D UAV Control for Persistent Communication Service and Fairness: A Deep Reinforcement Learning Approach. *IEEE Access*, 8: 53172–53184.
- Rabin, M. 1993. Incorporating fairness into game theory and economics. *The American economic review*, 1281–1302.
- Raeis, M.; and Leon-Garcia, A. 2021. A Deep Reinforcement Learning Approach for Fair Traffic Signal Control. In *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, 2512–2518.
- Rawls, J. 1971. A theory of justice. *Cambridge (Mass.)*.
- Siddique, U.; Weng, P.; and Zimmer, M. 2020. Learning Fair Policies in Multi-Objective (Deep) Reinforcement Learning with Average and Discounted Rewards. In *International Conference on Machine Learning*, 8905–8915. PMLR.
- Singh, A.; Halpern, Y.; Thain, N.; Christakopoulou, K.; Chi, E.; Chen, J.; and Beutel, A. 2020. Building healthy recommendation sequences for everyone: A safe reinforcement learning approach. *FACCTRec Workshop*.
- Skarlicki, D. P.; and Folger, R. 1997. Retaliation in the workplace: The roles of distributive, procedural, and interactional justice. *Journal of applied Psychology*, 82(3): 434.
- Steck, H.; van Zwol, R.; and Johnson, C. 2015. Interactive recommender systems: Tutorial. In *Proceedings of the 9th ACM Conference on Recommender Systems*, 359–360.
- Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement learning: An introduction*. MIT press.
- Tong, L.; Chen, Y.; Zhou, X.; and Sun, Y. 2021. Qoe-fairness tradeoff scheme for dynamic spectrum allocation based on deep reinforcement learning. In *Proceedings of the 5th International Conference on Computer Science and Application Engineering*, 1–7.
- Valkanis, A.; Papadimitriou, G.; Beletsoti, G.; Varvarigos, E.; and Nicopolitidis, P. 2022. Efficiency and fairness improvement for elastic optical networks using reinforcement learning-based traffic prediction. *Journal of Optical Communications and Networking*, 14(3): 25–42.
- Vlasceanu, M.; and Amodio, D. M. 2022. Propagation of societal gender inequality by internet search algorithms. *Proceedings of the National Academy of Sciences*, 119(29): e2204529119.
- Wang, L.; Bai, Y.; Sun, W.; and Joachims, T. 2021a. Fairness of Exposure in Stochastic Bandits. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, 10686–10696. PMLR.
- Wang, M.; and Deng, W. 2020. Mitigating Bias in Face Recognition Using Skewness-Aware Reinforcement Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Wang, Y.; Wang, X.; Beutel, A.; Prost, F.; Chen, J.; and Chi, E. H. 2021b. Understanding and Improving Fairness-Accuracy Trade-offs in Multi-Task Learning. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, KDD ’21*, 1748–1757. New York, NY, USA: Association for Computing Machinery. ISBN 9781450383325.
- Weng, P. 2019. Fairness in reinforcement learning. *arXiv preprint arXiv:1907.10323*.
- Wolff, J. 1998. Fairness, respect, and the egalitarian ethos. *Philosophy & public affairs*, 27(2): 97–122.
- Yamazaki, M.; and Yamamoto, M. 2021. Fairness Improvement of Congestion Control with Reinforcement Learning. *Journal of Information Processing* 29, 592–595.
- Yu, G.; Siddique, U.; and Weng, P. 2023. Fair Deep Reinforcement Learning with Generalized Gini Welfare Functions. In *International Conference on Autonomous Agents and Multiagent Systems*, 3–29. Springer.
- Zhang, C.; and Shah, J. A. 2014. Fairness in Multi-Agent Sequential Decision-Making. In Ghahramani, Z.; Welling, M.; Cortes, C.; Lawrence, N.; and Weinberger, K., eds., *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc.
- Zhang, D.; and Wang, J. 2021. Recommendation Fairness: From Static to Dynamic. *CoRR*, abs/2109.03150.
- Zhang, X.; and Liu, M. 2021. Fairness in learning-based sequential decision algorithms: A survey. *Handbook of Reinforcement Learning and Control*. Cham: Springer International Publishing, 525–555.

Zhong, H.; Deng, Z.; Su, W. J.; Wu, Z. S.; and Zhang, L. 2024. Provable Multi-Party Reinforcement Learning with Diverse Human Feedback. *arXiv preprint arXiv:2403.05006*.

Zhu, J.; Hu, M.; Liang, X.; Zhang, A.; Jin, R.; and Liu, R. 2023. Fairness-Sensitive Policy-Gradient Reinforcement Learning for Reducing Bias in Robotic Assistance. *arXiv preprint arXiv:2306.04167*.

Zhu, Q.; and Oh, J. 2018. Deep Reinforcement Learning for Fairness in Distributed Robotic Multi-type Resource Allocation. In *2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA)*, 460–466.

Zimmer, M.; Glanois, C.; Siddique, U.; and Weng, P. 2021. Learning fair policies in decentralized cooperative multi-agent reinforcement learning. In *International Conference on Machine Learning*, 12967–12978. PMLR.