Survey on Fair Reinforcement Learning: Theory and Practice

PRATIK GAJANE, Eindhoven University of Technology, The Netherlands

AKRATI SAXENA, Eindhoven University of Technology, The Netherlands

MARYAM TAVAKOL, Eindhoven University of Technology, The Netherlands

GEORGE FLETCHER, Eindhoven University of Technology, The Netherlands

MYKOLA PECHENIZKIY, Eindhoven University of Technology, The Netherlands

Fairness-aware learning aims at satisfying various fairness constraints in addition to the usual performance criteria via data-driven machine learning techniques. Most of the research in fairness-aware learning employs the setting of fair-supervised learning. However, many dynamic real-world applications can be better modeled using sequential decision-making problems and fair reinforcement learning provides a more suitable alternative for addressing these problems. In this article, we provide an extensive overview of fairness approaches that have been implemented via a reinforcement learning (RL) framework. We discuss various practical applications in which RL methods have been applied to achieve a fair solution with high accuracy. We further include various facets of the theory of fair reinforcement learning, organizing them into single-agent RL, multi-agent RL, long-term fairness via RL, and offline learning. Moreover, we highlight a few major issues to explore in order to advance the field of fair-RL, namely – i) correcting societal biases, ii) feasibility of group fairness or individual fairness, and iii) explainability in RL. Our work is beneficial for both researchers and practitioners as we discuss articles providing mathematical guarantees as well as articles with empirical studies on real-world problems.

1 INTRODUCTION

Recently, the machine learning (ML) and artificial intelligence (AI) community have taken major steps in advancing the research towards fairness-aware learning. We point the readers to the following non-exhaustive list of references: Barocas et al. [12], Gajane and Pechenizkiy [43], Mehrabi et al. [89]. Most of the research has been focused on addressing fairness concerns in supervised learning. Let us consider a real-world application of hiring, which we shall use as a running example throughout the article. Consider a firm aiming to hire employees for a number of positions. Common applications of ML in hiring include identifying the candidates, processing incoming applications, and selection (see Krishnakumar [69, Table 4.1] and references therein). Typically, ML tools for such scenarios use supervised learning, i.e., they model the relationship between applicants and outcomes from a training dataset, and then apply their model

to predict outcomes for subsequent applicants [75]. It has been shown that supervised learning techniques can effectively identify high-quality candidates as well as reduce the duration and cost of hiring [54]. Nevertheless, supervised learning methods are prone to discrimination or biased outcomes if they are trained on inaccurate [68], biased [13], or unrepresentative input data [119]. A number of pre-processing and post-processing methods have been proposed to mitigate these issues (e.g., [67],[44],[50]). However, these methods can be expensive, cumbersome, and may also reduce the accuracy of the assessment [107]. Thus, it is reasonable to look for valid alternatives for fairness-aware supervised learning techniques.

In addition, it has been shown that imposing fairness constraints as a static singular decision (as standard supervised learning methods do) while ignoring subsequent dynamic feedback can harm sub-populations [32, 34, 82]. Hence, it is imperative for researchers and practitioners to account for the indirect and delayed effects of the decisions being made by the ML system. It is customary to model such scenarios as sequential decision-making problems as they allow the consequences of a decision to last for an arbitrarily long time [104]. Indeed, regarding our running example of hiring, Zhao and Gordon [147] note: "Hiring is rarely a single decision point, but rather a cumulative series of small decisions."

A number of classical optimization methods can be used for sequential decision-making problems; however, they require as input a complete characterization of the learning environment [120]. In the case of hiring, this translates to a comprehensive knowledge of how hiring decisions affect the goal of a fair system (i.e., identify high-quality candidates while complying with fairness constraints) which, in practice, is not always available. Reinforcement learning (RL) methods provide a suitable solution in such scenarios as they can learn from interacting with an unknown environment. Furthermore, as explained in Sutton and Barto [120], the ability to take into account indirect and delayed consequences of actions is a key feature of reinforcement learning methods. This is because an RL agent's actions are permitted to affect the future state of the environment, thereby affecting the options available to the agent at later times.

For the task of hiring, supervised learning may also tend to select from groups with proven track records (i.e., *exploitation*), rather than selecting from non-traditional applicants (i.e., *exploration*), raising valid concerns about access to opportunity [75]. Even the goal of selecting high-quality candidates requires some exploration in order to possibly discover a better pool of applicants and consequently make better hiring decisions. However, the hiring firm must also exploit its knowledge about the well-understood parts of applicants' pool to make rewarding decisions. Therefore, neither exploration nor exploitation can be pursued exclusively without failing at the task. RL methods are designed to provide a meticulous and goal-directed trade-off between exploration and exploitation, unlike supervised learning methods. Hence, reinforcement learning can be effectively used for real-world applications (such as hiring, recommendation systems, resource allocation, etc.), posed as a sequential decision-making problem.

Li et al. [75] provide comparative results of using reinforcement learning and supervised learning in hiring. They focus on the decision to grant first-round interviews for high-skilled positions in consulting, financial analysis, and data science. Using data from professional services recruiting within a Fortune 500 firm, they demonstrate that the approach based on RL improves the quality (as measured by eventual hiring rates) of candidates selected for an interview while also increasing demographic diversity. On the other hand, supervised learning-based algorithms improve hiring rates but select far fewer Black and Hispanic applicants. Thus, fair reinforcement learning (fair-RL) methods provide a viable solution for fairness-aware decision-making problems which might not be optimally solvable by supervised learning.

Fair-RL poses additional challenges that traditional fairness-oblivious RL methods are unable to solve. A frequently used performance measure in RL algorithms is the *regret*. Regret of an algorithm measures the difference between the best possible performance and the performance of the said algorithm. Therefore, optimizing the performance is equivalent to minimizing the regret, which is a commonly used goal in fairness-oblivious RL. However, minimizing

the regret might not be sufficient for satisfying fairness constraints. It is easy to construct RL solutions that achieve optimum regret while performing unacceptably on fairness requirements. At the same time, dismal regret cannot be tolerated for the sake of a higher level of compliance with the fairness constraints. Hence, fair-RL presents a dual goal of optimizing the usual performance criteria such as regret and adhering to fairness considerations.

The remainder of this article is structured as follows. In Section 2, we begin by providing the preliminaries of fair-RL. Then, in Section 3, we furnish the real-world usage of fair-RL methods, followed by an extensive overview of the theory of fair-RL methods in Section 4. The kind of problems demonstrated by our running example of hiring mainly fall under single-agent RL, and these methods are enlisted in Section 4.1. Fair-RL approaches to handle the extensions of our running example to multiple agents, long-term fairness, and offline learning are described in Section 4.2, 4.3, and 4.4 respectively. In Section 5, we discuss a few of the major challenges from fair-RL research, and in the end, we provide brief concluding remarks in Section 6.

To the best of our knowledge, the closest work to ours is from Zhang and Liu [146] that reviews fairness approaches in sequential decision-making problems. While there is a small overlap limited to Zhang and Liu [146, Section 3.1 and 4.2.5], following are the salient additions in our work – i) we include long-term fairness as well as multi-agent RL problems, ii) we include application of fair-RL approaches to practical problems, and iii) we discuss challenges in fair-RL research at length.

2 FUNDAMENTALS OF (FAIR-)REINFORCEMENT LEARNING

In this section, we briefly introduce the fundamental concepts of reinforcement learning. However, providing a thorough review of RL is beyond the scope of this article. We point the readers to Sutton and Barto [120], Szepesvári [121] for a thorough understanding of the subject matter. Reinforcement learning problems are usually modeled as either multi-armed bandits (MAB), Markov decision processes (MDP), or their variants. Below we briefly introduce some of these formulations.

2.1 Multi-Armed Bandit (MAB)

A multi-armed bandit problem can be symbolized as a game from time period $t = 1, \ldots, T$ between a learner and its environment where T is called the *horizon*. The learner has a set of arms (or actions, used interchangeably in the rest of the article) $\mathcal{A}$ available to it. At every time period $t = 1, \ldots, T$, each arm a is associated with a numerical value i.e., a reward [1]. The learner's task is to pull an arm from $\mathcal{A}$, whereupon it receives the reward associated with that arm. Such a feedback is called *bandit feedback*. The learner's goal is to optimize the value of the arms it chooses, and hence its task is to find a policy that selects the best arm for a given time period.

In the stochastic version of this problem, a stationary reward distribution is associated with each arm, and the rewards are drawn from the respective distributions. In the adversarial setting, the rewards are generated by an adversary, and the reward probabilities may not be stationary. Generally, in many real-life applications, some information or *context* is available that can be used to make a better decision when choosing amongst all actions. These scenarios are modeled by *contextual bandits*, in which at time step t, the learner first observes the context C_t before choosing an action. In the problems described so far, the learner receives absolute feedback about its choices; however, relative feedback is naturally suited to many practical applications where people are expected to provide feedback, e.g., user-perceived product preferences or information retrieval systems. These scenarios are modeled by *duelling bandits* in

---

[1]In some scenarios, considering a loss instead of a reward might be more suitable, however, mathematically, the problem settings as well as the general solution strategies are equivalent. In this article, for the sake of uniformity we consider rewards by default and not losses.

which the learner is to select two actions at every time step. As feedback, the learner receives information about which action gave a better reward, i.e., won the duel. Additionally, in order to produce scalable results when the number of actions is large, the setting of *infinite bandits* is used in the literature. Here, the learner has to choose from a convex set of arms contained in a ball of the given radius. Alternatively, *causal bandits*[71] are used to study the problem of using causal models to improve the rate at which good interventions[2] can be learned online. In a causal bandit problem, interventions are treated as arms, but their influence on the reward — along with any other observations — is assumed to conform to a known causal graph.

Algorithms for MAB problems are typically divided into two groups [65] – i) upper confidence bound (UCB) algorithms (frequentist algorithms), and ii) Thompson sampling based algorithms (Bayesian algorithms).

2.2 Markov Decision Process (MDP)

A Markov decision process is characterized by its state set $\mathcal{S}$, the action set $\mathcal{A}$, and one-step dynamics of the environment, which specify, for any state and action, the probability of each possible pair of next state and reward. The learner's task is to choose an action from $\mathcal{A}$ at each time step. As with MAB problems, the learner's goal is to optimize the received rewards. In some cases, a reward in the future is deemed not worth quite as much as a reward now. This is expressed by using a discount factor $0 < \gamma < 1$ and a reward that occurs n steps in the future is multiplied by γ^n. In episodic RL problems, an agent interacts with the environment in episodes of fixed length. A partially observable MDP (POMDP) is a generalization of an MDP to model planning under uncertainty. In a POMDP, the learner cannot directly observe the system state, and it uses its observations to form a belief about the current state. Q-learning is a popular RL algorithm [120, Chapter 6] and deep Q-learning is its variant using deep convolutional neural network [94].

2.3 Multi-Agent RL (MARL)

Multi-agent reinforcement learning studies how multiple agents can collectively learn in the same environment. To see how the single-agent RL concepts discussed above can be readily extended to MARL, we point the readers to Busoniu et al. [17].

2.4 Fairness in Reinforcement Learning

In the literature, fairness in reinforcement learning methods is aimed at either of the following:

- Alleviate *societal bias* from the decisions made by learning algorithms (societal fairness) – Societal bias or discrimination in this context refers to unfavorable treatment of people due to their membership to certain demographic groups that are distinguished by the *protected* or *sensitive* attributes, such as race, gender, age, etc. Discrimination, based on protected attributes, is prohibited by international legislation [2]. RL techniques have been applied to several real-life problems from different areas to achieve fair solutions, such as hiring [75], disease contagion control problem [8], face recognition [135], recommendation systems [14, 16, 115], and so on.
- Adhere to defined fairness constraints in decision-making or resource allocation (non-societal fairness) – On the other hand, some works frame fairness as a way of adhering to defined constraints in allocation tasks or real-time/sequential decision-making problems. This is particularly common in designing fair solutions for resources allocation schemes in computer networks. Other well known applications include HTTP adaptive streaming [4], traffic management in autonomous vehicles [7], etc.

---

[2]We borrow the term from Pearl [98].

3 REAL-WORLD IMPLEMENTATIONS OF FAIR REINFORCEMENT LEARNING

Fair-RL methods have been applied to achieve fair solutions for various real-world problems. In this section, we discuss some of these works categorized based on societal and non-societal fairness.

3.1 Implementations of Fair-RL for Societal Fairness

Li et al. [75] model the decision to extend interview opportunities in the recruitment process as contextual bandits and employ UCB-GLM algorithm [79] as a solution. They test the proposed solution on the administrative data of 88,666 job applications made to recruitment services within a Fortune 500 firm. Their results show that the used fair-RL method more than doubles the share of Black and Hispanic representation in selected applicants. At the same time, the quality of the selected applicants (as measured by their hiring potential) is also substantially increased.

Atwood et al. [8] study the *precision disease contagion control* problem, which aims to provide a fair allocation of vaccines in society. In such problems, uniform allocation might not be a fair solution as some people might be at more risk due to their characteristics, such as age, or due to their position in the network. A fairness-aware vaccine allocation policy aims at finding allocation strategies that equalize outcomes in the population. They train a Deep Q-Network [95] using dopamine's experiment runner and assume that the agents providing the vaccines can observe the contact network structure of individuals and their disease state. The reward function at each step is the negative number of newly sick nodes. The results show that the proposed solution achieves fairness as compared to baselines. These promising results confirm that RL can be further used to achieve fair solutions for such social problems having complex network structures.

Wang and Deng [135] propose an RL-based race balance network (RL-RBN) that reduces the skewness of features between different races for mitigating the bias in face recognition. An MDP-based setting is used to find the optimal margins for non-Caucasians, and then a deep Q-learning method is applied to learn policies for an agent to select an appropriate margin. Fairness is evaluated using standard deviation and skewed error ratio (SER), and the results show the effectiveness of RL-RBN.

3.2 Implementations of Fair-RL for Non-societal Fairness

Another well-known application of fairness-aware solutions using RL is *Enhancing transmission control protocol* (TCP) over wireless mesh networks (WMN). Generally, TCP is not fair in resource allocation over WMN [42]. Arianpoo and Leung [6] propose a distributed mechanism that observes the unfairness in resource allocation and then tunes the TCP parameters accordingly. In their solution, each TCP source models the state of the multi-hop network as an MDP, and then they use the Q-learning algorithm to monitor and learn the transition matrix of the proposed MDP. The agent learns the network usage behavior of each node using the local fairness index and the aggressive index. The reward function is used for guiding the learning agent to choose the correct actions that will eventually provide a fair solution in a distributed manner. The proposed solution improves the fairness of the flows traversing a larger number of hops with a negligible impact on the smaller size flows by tuning TCP parameters. The overall performance is enhanced by 10-20% without using feedback messaging and no extra overhead to the system. Yamazaki and Yamamoto [140] present an RL-based method, called QTCP-AIMD, that uses Q-learning based TCP with Additive Increase Multiplicative Decrease (AIMD) to improve fairness in congestion window control mechanism. The simulation analysis shows that the proposed method improves fairness without degrading both throughput as well as low latency characteristics of

QTCP. Tong et al. [126] propose a deep RL-based method for dynamic spectrum allocation in the case of resources shortage.

HTTP adaptive streaming (HAS) is a popular service for adaptive video streaming, in which each video is temporally segmented and stored in varying levels of quality. The quality selection heuristics used by the player dynamically request for the most appropriate quality level based on the network conditions. Most of these heuristic methods are fixed and deterministic, and are not able to provide good performance if the network conditions are highly dynamic. The quality of experience for users might be poor due to freezing or frequent quality switches while playing the videos. In real scenarios, multiple clients share a single medium and request videos from the HAS server. The mutual synchronization among clients gives rise to unfairness among them as one client might have a negative impact on others [4]. Petrangeli et al. [100] propose a multi-agent Q-learning based HAS client that will learn and dynamically adapt its behavior depending on the network conditions to achieve fairness and provide a high quality of experience which is measured using average mean opinion score [37]. Some other works on fairness-aware resource optimization include Chen et al. [20] and Valkanis et al. [128].

RL-based solution for autonomous vehicles by adding fairness terms in the reward function significantly improves the learned behavior of the agents to avoid bottleneck situation [7]. The results show that an RL method provides high outflow (vehicles per hour) and throughput efficiency [130]. The authors study a single-agent method, and further investigation in multi-agent setting would be interesting. Applying fairness penalties in multi-agent settings for other non-bottleneck situations, such as changing lanes or speed unnecessarily, are also open research questions.

RL has provided promising results to achieve fairness and quality of service in *radio resource management* by selecting scheduling rules to allocate users' data packets in the frequency domain. In Comsa et al. [29], Q-learning is used to learn different policies for scheduling rules at each transmission time interval. The proposed solution achieves different levels of throughput-fairness trade-off by offering optimal solutions according to the channel quality indicator for different classes of users. Comşa et al. [26] propose an RL-based technique that considers the traffic load and channel conditions to learn the parameters for a Generalized Proportional Fair scheduling rule that respects the fairness criteria. The authors further propose a more complex RL framework to achieve higher NGMN (Next Generation of Mobile Networks) fairness [27, 28]. Yuan et al. [141] model the *resource block group* allocation problem as a stochastic game framework and provide a multi-agent RL-based solution to achieve fairness by optimizing 5-percentile user data rate.

In the past few years, researchers have focused on designing fairness-aware recommendation systems [14, 16, 115]. Apart from recommendation systems based on static settings, RL-based systems have provided good performance for interactive systems where users' evolving preferences must be taken into account while determining if the system is fair [144]. Liu et al. [84] introduce an RL-based framework, called FairRec, that maximizes the cumulative reward function based on both fairness and accuracy, and maintains a dynamic balance between both in the long run for an interactive recommendation. The fairness is evaluated using Unit Fairness Gain, which considers both a higher conversion rate and a higher weighted proportional fairness.

Furthermore, human-in-the-loop IoT systems are popular means to provide a personalized experience. These systems continuously learn human behavior and evolve to adapt to the human and environment state and take actions autonomously or by way of recommendation. A major challenge in designing personalized IoT applications arises due to human variability as different people interact with IoT applications in different ways or change their behavior over time. Elmalaki [40] propose Fair-IoT using a reinforcement learning-based framework that continuously monitors the human state and changes in the environment to adapt its behavior accordingly. The proposed adaptive and fairness-aware solution is further used for two human-in-the-loop IoT applications – i) automotive advanced driver assistance

systems, and ii) smart house. The analysis shows that the proposed solution improves systems' fairness by 1.5 times and enhances the human experience by 40% to 60% compared to non-personalized systems.

RL-based techniques have been employed to find fair solutions in other domains as well, such as reducing the air traffic congestion [36], traffic light control to optimize the waiting time of all the drivers [74, 106], resources management [5, 80, 105], electric taxi charging [132], resource allocation in robot systems (human-robot systems, multi-robot systems) [24, 148], and same-day delivery services [21].

3.3 Unexplored Issues in the Implementations of Fair-RL

In this section, we summarized the fair-RL methods that have been deployed in the real-world. Following are a few of the major unexplored issues concerning them.

(1) Interpretability/Explainability: Many of the above works use deep RL methods and it remains difficult to diagnose what aspects of the input affect the decisions of these methods. Explainability in deep RL methods is still not widely studied [139]. The lack of comprehensible explainability of automated methods is a concern in many real-world domains such as legislation, law enforcement, healthcare, etc. As explained in Corbett-Davies and Goel [30], an opaque system may engender mistrust from policymakers and stakeholders thus hindering implementation. In particular, the need for explainability for algorithmic hiring is asserted by Schumann et al. [110]. This point is further discussed in detail in Section 5.3.

(2) Focus on societal-fairness: As we shall see in Section 4, most of the research in the theory of fair-RL is motivated by societal fairness. As witnessed in the current section, practical use of fair-RL lags behind the theory in this aspect. In addition to being a legal requirement in many parts of the world, satisfying a valid societal fairness constraint is beneficial from a business perspective too. Studies demonstrate that if a lack of fairness or equality is perceived by people, it might impact their job satisfaction [88], which results in retaliatory behavior from the affected parties [116]. We have already noted how fair-RL methods fare better than traditionally used fair-supervised learning methods for tasks like hiring [75]. We believe business needs and evidence of the viability of fair-RL methods will naturally motivate extensive use of fair-RL techniques for societal fairness in the real world.

(3) Cold start: RL methods sometimes tend to suffer from "cold start" i.e., poor performance at the beginning. This might hinder implementation of fair-RL methods in critical domains. Some techniques (e.g., [145]) have been suggested to overcome this problem and it needs to be studied if they can be incorporated in fair-RL methods too.

4 THEORY OF FAIR REINFORCEMENT LEARNING

We divide this section into multiple subsections based on the type of the used RL method. We elucidate some articles in detail while only briefly mentioning others and deferring their details to appendix B due to space constraints.

4.1 Fairness in Single-Agent RL

An extensive number of fair-RL approaches have been developed in single-agent settings. We categorize them according to the used fairness notion for ease of exposition. Table 1 provides a comparative view of some of the noteworthy theoretical results. Lack of space prevents us from including a definition of each symbol in the table. Please refer to the description in the relevant text or in the list of symbols given in Appendix A.

Table 1. Performance bounds

Reference	Formulation	Fairness Notion	Performance Measure	Bound
Joseph et al. [62, Theorem 2]	Stochastic bandits	Individual fairness	Cumulative regret	$\tilde{O}\left(\sqrt{\|\mathcal{A}\|^3 T}\right)$
Joseph et al. [63, Theorem 1]	Contextual bandits	Individual fairness	Cumulative regret	$\tilde{O}(d\|\mathcal{A}\|^2\sqrt{T})$
Joseph et al. [63, Theorem 2]	Infinite bandits	Individual fairness	Cumulative regret	$\tilde{O}\left(\frac{rad^6\Psi^2}{\kappa^2\lambda^2\Delta_{gap}^2}\right)$
Jabbari et al. [59, Theorem 6]	Discounted reward MDP	Individual fairness	Sample complexity to achieve (near-)optimality	$\tilde{O}\left(\frac{f\|\mathcal{S}\|^5\|\mathcal{A}\|^{\frac{1}{1-\gamma}+5}}{(1-\gamma)^{12}}\right)$
Liu et al. [85, Theorem 4.1]	Stochastic bandits	Individual fairness	Expected cumulative deviation from the fairness constraint	$\tilde{O}\left((\|\mathcal{A}\|T)^{2/3}\right)$
Liu et al. [85, Theorem 5.2]	Dueling bandits	Individual fairness	Expected cumulative deviation from the fairness constraint	$\tilde{O}\left(\|\mathcal{A}\|^{4/3}T^{2/3}\right)$
Gillen et al. [46, Theorem 3]	Linear contextual bandits	Individual fairness	Cumulative regret	$\tilde{O}\left(\|\mathcal{A}\|^2 d^2 \log{(T)} + d\sqrt{T}\right)$
Huang et al. [56, Corollary 4.1]	Contextual bandits	Group fairness	Cumulative regret	$O\left(d\sqrt{T}\log{(TL)}\right)$
Wen et al. [137, Theorem 5.1]	Episodic discounted MDP	Group fairness	Cumulative regret	$O\left(\left(N^{2/3}+1/\epsilon^2\right)\log{(1/\delta)}\right)$
Claure et al. [25, Theorem 1 & 2]	Stochastic bandits	Minimum selection criteria	Cumulative regret	$O\left(\sqrt{(\|\mathcal{A}\|T\log{T})} + \|\mathcal{A}\|\log{T}\right)$
Li et al. [76, Theorem 2]	Comb. sleeping bandits	Minimum selection criteria	Time-average regret	$O\left(\frac{\|\mathcal{A}\|}{2\eta} + \frac{\sqrt{m\|\mathcal{A}\|T\log{T}}+\|\mathcal{A}\|}{T}\right)$
Patil et al. [97, Theorem 8]	Stochastic bandits	Minimum selection criteria	Cumulative regret	$O(\sqrt{T\log{T}})$
Chen et al. [22]	Contextual bandits	Minimum selection criteria	Cumulative regret	$O\left(\sqrt{T\|C\|\|\mathcal{A}\|\log(\|\mathcal{A}\|)}\right)$
Celis et al. [18, Theorem 1]	Stochastic Bandits	Selection within pre-specified range	Cumulative regret	$O\left(\frac{\|\mathcal{A}\|}{\Delta_{gap}^2}\log(T)\right)$
Wang et al. [134, Theorem 3.2.2, 3.3.2]	Stochastic bandits	Selection proportional to merit	Cumulative regret	$\tilde{O}\left(\sqrt{\|\mathcal{A}\|T}\right)$
Wang et al. [134, Theorem 4.2.2, 4.3.2]	Linear stochastic bandits	Selection proportional to merit	Cumulative regret	$\tilde{O}\left(d\sqrt{T}\right)$
Huang et al. [57, Theorem 3]	Causal bandits	Counterfactual individual fairness	Cumulative regret	$\tilde{O}\left(\frac{\sqrt{\|\mathcal{W}\|T}}{\tau-\Delta}\right)$

4.1.1 Individual fairness in single-agent RL. A commonly used fairness notion is individual fairness [39] which stipulates an RL system to make similar decisions for similar individuals.

Joseph et al. [62] consider a notion of individual fairness called *meritocratic fairness* in stochastic bandits and contextual bandits. This fairness notion stipulates that a fair algorithm should not select one arm over another if the chosen arm has a lower expected reward than the unchosen arm. They provide a lower bound, which proves that no algorithm has diminishing regret before $\Omega(|\mathcal{A}|^3)$ time steps. For stochastic bandits, the authors propose an algorithm called FAIR-BANDITS and prove a $\tilde{O}\left(\sqrt{|\mathcal{A}|^3 T}\right)$ upper bound on its regret which matches the lower bound when $T \approx |\mathcal{A}|^3$. For any contextual bandit problem, the authors demonstrate that the optimal learning rate of any fair algorithm is determined by the best KWIK ("Knows When It Knows") [78] bound for the problem. Joseph et al. [62] also provide a reduction from KWIK learning algorithm to a fair contextual bandit algorithm and vice versa.

Jabbari et al. [59] extend meritocratic fairness to MDPs with discounted rewards. Their fairness constraint called, approximate-action fairness, requires that an algorithm never favors an action of substantially lower quality over a better action. They design an algorithm, called *Fair-E^3*, satisfying approximate-action fairness which achieves (near-)optimality after at most $\tilde{O}\left(\frac{f|\mathcal{S}|^5|\mathcal{A}|^{\frac{1}{1-\gamma}+5}}{(1-\gamma)^{12}}\right)$ steps where f is a problem-dependent value and γ is the discount factor.

It has been argued in the literature (e.g., [85]) that meritocratic fairness suffers from the following limitations – i) it allows a subset of arms best in expectation by only a small margin to be selected all the time, even if any single sample from the subset may be worse than a single sample from another subset, ii) it does not constrain the learner in case one subset of arms is much better than other subsets, iii) it does not account for/correct past inequities or inaccurate/biased data, and iv) it assumes an accurate mapping from features/arms/actions to true quality is available for the task at hand.

Liu et al. [85] impose the following constraints on stochastic bandits and dueling bandits: i) *smooth fairness* constraint - two arms with similar distribution should be selected with similar probability, and ii) *calibrated fairness* constraint - sample each arm with probability equal to its reward being the greatest. They further elucidate how these two constraints address the first two critiques of meritocratic fairness presented earlier. Joseph et al. [63] extend upon Joseph et al. [62] by addressing a couple of shortcomings of the latter. Firstly, the problem setting by Joseph et al. [63] allows for multiple individuals per group per round, and the learner can choose multiple individuals per round, unlike that of Joseph et al. [62]. Gillen et al. [46] contend that similarity-metric based fairness definitions suffer from the problem that it may be actually difficult for anyone to precisely express a quantitative metric over individuals. To circumnavigate this problem, they assume that the algorithm has access to an oracle that can ascertain what it means to be fair, but cannot explicitly enunciate the fairness metric. After the learner makes its decisions, it receives feedback from the oracle about violations of the fairness constraint. For more details about Liu et al. [85], Joseph et al. [63], and Gillen et al. [46], please refer to Appendix B.

4.1.2 Group fairness in single-agent RL. Another common notion is group fairness which imposes statistical/demographic parity in the decisions made by an RL system.

Huang et al. [56] consider group fairness in contextual bandits focusing on the practical application of recommendation systems. Group fairness constraint here necessitates the expected mean reward of the protected group and that of the unprotected group to be equal (barring a small additive tolerance degree). The considered performance measure is expected cumulative regret compared to the best algorithm satisfying the fairness constraint. They further propose an algorithm called FAIR-LINUCB , and prove an upper bound of $O\left(d\sqrt{T}\log(TL)\right)$, where d is the dimension of the feature vector and L is an upper bound on the $L2$ norm of all the feature vectors. The experimental results on

simulated datasets demonstrate that Fair-LinUCB achieves competitive regret while complying with the considered fairness notion.

Schumann et al. [111] study contextual bandits with the following two definitions of group fairness – i) demographic parity, wherein the probability of choosing a group is the same across groups, and ii) proportional parity, wherein the probability of choosing an arm from a particular group is proportional to the size of that group. Wen et al. [137] develop fair decision-making policies in discounted MDPs, and their approach works with demographic parity and equal opportunity [50]. For more details about Schumann et al. [111] and Wen et al. [137], please refer to Appendix B.

4.1.3 Minimum selection criteria. Celis et al. [18] study stochastic bandits considering the practical application of preventing polarization in personalized online spaces [108, 118]. In this setting, the fairness constraint demands that the probability with which an algorithm selects a group of arms stays within a pre-specified interval. They propose an algorithm called Constrained-ϵ-Greedy and prove an upper bound of $O\left(\frac{|\mathcal{A}|}{\Delta_{gap}^2} \log(T)\right)$ on the expected cumulative regret where Δ_{gap} is the difference between the maximum and the second maximum expected rewards. Additionally, they provide empirical results showing competent performance of their algorithm on a dataset of news articles with the aim of diversifying across topics, and the MovieLens dataset [51] with the aim of diversifying recommendations.

Claure et al. [25] explore stochastic bandits with a fairness constraint that a minimum pulling rate for each arm is satisfied (in expectation or anytime throughout the time horizon). The performance of the proposed algorithms is measured in terms of regret with respect to the optimal benchmark strategies satisfying the fairness constraint. The authors provide UCB-based algorithms to satisfy both the flavors of the above fairness constraint and prove a regret bound of $O\left(\sqrt{(|\mathcal{A}|T \log T)} + |\mathcal{A}| \log T\right)$.

Li et al. [76] also employ minimum selection criteria for each individual arm in a setting called combinatorial sleeping bandits in which multiple arms (up to m) can be simultaneously played, and an arm could sometimes be "sleeping" (i.e., unavailable). The authors enlist real-time scheduling in wireless networks, ad placement in online advertising systems, and task assignment in crowd-sourcing platforms as practical applications and present experimental results on simulations showing the competent performance of their approach. Patil et al. [97] consider stochastic bandits where the minimum selection fraction of each arm is limited to $\left[0, \frac{1}{|\mathcal{A}|-1}\right)$. Chen et al. [22] investigate the problem of an AI system assigning tasks or distributing resources to multiple humans. They model this problem using the setup of contextual bandits, and their fairness criterion demands a minimum rate for assigning a task or a resource to an individual, which translates to a minimum probability $\in \left(0, \frac{1}{|\mathcal{A}|}\right)$ of each arm being pulled. For more details about Li et al. [76], Patil et al. [97], and Chen et al. [22], please refer to Appendix B.

4.1.4 Other notions. Wang et al. [134] aim to achieve merit-based fairness of exposure to the items while optimizing utility to the users in stochastic bandits and stochastic linear bandits. The fairness constraint requires that each arm receives an amount of exposure proportional to its merit, where merit is quantified through an application-dependent merit function. The experimental results on synthetic data and on real-world data from "Yahoo! Today" Module [77] show that their proposed algorithms provide better fairness in terms of exposure to the items. Huang et al. [57] study counterfactual individual fairness in causal bandits. They focus on the practical application of online recommendations and aim to provide user-side fairness for customers. Experimental results provided on the Email Campaign data [86] show that their proposed algorithm maintains good performance while satisfying counterfactual individual fairness in each round. Ghodsi and Mirfakhar [45] provide an RL solution for fair and efficient allocation of a divisible resource among a population with distinct preferences. They model this problem using adversarial bandits and propose an

algorithm based on EXP3 [10] which attains sub-linear bounds for both cumulative regret and fairness-regret if the adversary is restrained in a certain way. For more details about Wang et al. [134], Huang et al. [57], and Ghodsi and Mirfakhar [45], please refer to Appendix B.

Nabi et al. [96] study fairness in RL from the perspective of causal inference and constrained optimization. They present several strategies for learning optimal policies by modifying some of the existing RL algorithms, such as Q-learning, which also account for some fairness considerations, and further provide a theoretical guarantee that their proposed fair policy satisfies the specified fairness constraints.

Talebi and Proutiere [122] address the problem of providing *proportionally fair* allocations in a sequential resource allocation problem where a set of tasks are to be handled by some servers. In their work, a learner is considered fair if all the tasks are served at similar rates, avoiding starvation. In a proportionally fair allocation, the aim is to maximize the sum of the logarithm of the task expected service rates.

4.2 Fairness in Multi-Agent RL

Balancing fairness is essential in many multi-agent systems in which fairness constraints must be satisfied among several users simultaneously. In a modification of our running example, suppose multiple divisions/subsidiaries of the same employer are looking to hire for a number of positions. Here, the hiring process of each subsidiary can be thought of as a single agent. A fair-RL algorithm for this problem aims for a fair solution for both the employer – to have a desirable set of employees, and applicants – to receive decisions that are fair (non-discriminatory) with respect to the protected attributes.

Fairness in multi-agent systems has been studied in reinforcement learning scenarios [101]. In one of the early attempts, Zhang and Shah [143] model fairness in multi-agent sequential decision-making via a simple linear programming approach. The fairness optimization is then formulated in a game-theoretic framework to achieve a Nash equilibrium that contains a regularized max-min fairness policy. Zhu and Oh [148] propose a multi-type resource allocation method for multi-robot systems by leveraging a weighted combination of utility and inequality, which is measured using the Gini-index, as the overall objective function in their approach. Moreover, Agarwal and Aggarwal [3] explore multi-agent settings with non-Markovian reward functions and develop model-based and model-free algorithms for joint decision making of multiple agents. They present a tabular model-based algorithm using Dirichlet sampling to obtain a regret bound of $O\left(KD|\mathcal{S}||\mathcal{A}|\sqrt{\frac{|\mathcal{A}|}{T}}\right)$, where K is the number of agents and D is the diameter of the underlying Markov chain. In a more decentralized setting, Jiang and Lu [61] introduce a hierarchical approach for joint policy optimization. In this setting, each agent first learns its own policy to balance fairness versus efficiency, and at a higher level, a controller maximizes the multi-objective reward by switching between the local policies while interacting with the environment. In their approach, the fairness measure is decomposed between the agents such that they only focus on optimizing their corresponding sub-goals. Additionally, Zimmer et al. [149] study cooperative multi-agent scenarios and present a fair policy learning to optimize a welfare function that utilizes both fairness efficiency and fairness equity in a decentralized manner. Their framework consists of a self-oriented and a team-oriented network which are together optimized via a policy gradient algorithm with theoretical proof of convergence.

Furthermore, Wang et al. [133] specifically show that a model which only focuses on maximizing the team reward might lead to a conflict and potential unfair outcomes for individual members. Accordingly, they develop a modular architecture to improve cooperation in social dilemmas. The proposed architecture adapts a fast timescale learning of individual agents combined with a slow evolution mechanism that allows for a natural selection in a population.

Grupen et al. [49] define a group-based measure of fairness for fully cooperative multi-agent settings and present an approach in which the agents learn how to coordinate fairly toward their common goal. They benefit from equivariant policy learning to achieve provably fair outcomes for individual members of a team where the trade-off between fairness and utility is obtained dynamically. Multi-agent RL techniques are also employed in other applications, such as learning stock trading strategies to keep a balance between revenue and fairness for the involved clients [11].

4.3 Long-term Fairness via RL

Most approaches designed to deal with fairness-aware learning only consider the immediate implication of bias in a static context, whereas in many applications, the decisions made by RL systems have consequences in the long term. Recent work explores the long-term effects of RL and demonstrates that modeling the immediate effect of decisions for single-step prevention of bias does not guarantee fairness in later downstream tasks [64, 83, 93]. For example, Holzer [55] explore the long-term effects of affirmative action [38] in hiring. They cite several studies to demonstrate that there is little evidence for affirmative action leading to weaker performance. Moreover, it clearly generates positive benefits for the minority and low-income communities, and perhaps for employers as well. They also enlist positive indirect consequences of affirmative action in certain sectors – e.g., medical care, where minority physicians are more likely to provide medical care to minorities and low-income communities. Fairness methods in RL aiming to deal with such long term effects are presented below.

Kannan et al. [64] study two-step decision scenarios and illustrate that a fair decision of a classifier in the first step has an indirect unfair impact on the next Bayesian decision-making task. Milli et al. [93] introduce robust strategies to address a similar problem in the context of strategic classification, which aims at lowering the social burden and disadvantage in certain groups in the long-term. Liu et al. [83] propose a one-step feedback model that reveals how decisions change the underlying population over time, confirming that common fairness criteria, in general, do not lead to reshaping the population and promoting long-term improvements.

The consequences of bias can go beyond the immediate next step and impact the stakeholders further steps away. In such scenarios, reinforcement learning is a remedy to model the long-term implications of bias in the form of Markov decision processes. D'Amour et al. [35] discuss the growing need to understand the long-term behaviors of deployed ML-based decision systems and their potential consequences, and accordingly, propose a framework for systematically exploring these long-term effects. Tu et al. [127] study the dynamics of population qualification and algorithmic decisions under a partially observed MDP and show that the long-term effects are heavily shaped by the interplay between algorithmic decisions and individuals' reactions. They further illustrate that the same fairness constraint can have an opposite impact depending on the underlying problem scenarios, which highlights the importance of understanding real-world dynamics in decision-making systems. These approaches are solely based on simulations by reproducing the dynamics of certain applications with known parameters to highlight the need for modeling the long-term implications of bias. Subsequently, they have been employed in some applications such as recommendation systems to retain fairness for a longer period of time [114, 144].

Consequently, we require methods that can learn the dynamics of a population and the optimal policies which are unbiased in the long-term. Wen et al. [138] explore the temporal effects of fair/unfair decisions for every individual in the population. They first formulate fairness definitions in an MDP setting and then propose an algorithm to learn an optimal policy that satisfies the fairness constraints. The presented experimental results show that accounting for the dynamic effects of decisions improves the results of supervised learning under optimistic assumptions. Siddique et al. [113] formulate a fair optimization problem in a multi-objective sequential decision-making setting. They propose to

adjust the standard reward function by applying a social welfare function on the reward distribution, leading to a multi-objective reinforcement learning problem based on welfare optimization. All these works underline the importance of modeling long-term fairness, which is still in the beginning phase, and more research needs to be done in the domain of fairness in reinforcement learning to account for the dynamic effects of decisions in our society.

4.4 Offline Fair-RL

Consider another variant of our running example: hiring via predictive analysis of offline data using learning models [99]. In such scenarios, historical recruitment data is used to make predictions about future hiring activities and candidates. Such problems are solved using offline fair-RL methods.

Metevier et al. [91] present a theoretically grounded contextual bandit algorithm for offline data. Their approach can adapt to various definitions of fairness and returns "no solution found" if data is insufficient or there might be conflict. In addition to RL and standard bandit settings, fairness-aware learning has been employed in counterfactual learning frameworks for scenarios that can only be evaluated on the historical data, and where online learning/interaction is infeasible, costly, or dangerous, e.g., in healthcare applications. Tavakol [123] propose to model the task of fair classification as a bandit problem in which group-based fairness constraints can be satisfied via off-policy learning in a counterfactual bandit setting. Consequently, counterfactual risk minimization techniques are leveraged to learn fair policies from biased offline data. In another work, Coston et al. [31] present a counterfactual risk assessment method using doubly robust estimation. This method reduces the likelihood of an adverse event by identifying risky cases where the decision-making needs human intervention. They further demonstrate that modeling the counterfactual outcomes is necessary in order to effectively conduct fairness-adjustment procedures in such scenarios.

4.5 Unexplored Issues in the Theory of Fair-RL

In this section, we saw articles contributing mainly to the theory of fair-RL. Albeit, they advance the field by providing a mathematically-grounded understanding of the used methods, the following avenues are still to be explored.

(1) Fairness-performance trade-off: The fairness-performance trade-off is a fundamental question in fair-ML, and it has been discussed in several existing works [19, 90]. The necessity of this trade-off has been contended in other forms of learning [136, 147]. Indeed, a recent empirical study challenges the assumption that adhering to fairness constraints necessarily leads to a noticeable drop in accuracy [109]. However, this important matter is mostly not investigated in the theory of fair-RL yet (except a few articles, e.g., [18, 46, 97, 133]).

(2) Strictness of performance bounds: Almost all articles (barring a few, e.g., Joseph et al. [62], Wang et al. [134]) do not prove a corresponding lower bound on the used performance measures. This makes it hard to determine the strictness of the performance bounds, i.e., how far they are from the best possible bounds.

(3) Credible surveys from social sciences to determine the applicability of the fairness measures: Most of the research does not make explicit use of the vast amount of surveys available in social sciences assessing the applicability of fairness notions for particular domains (see Gajane and Pechenizkiy [43] and references therein). These surveys could help RL researchers and practitioners to determine the most suitable fairness notions for the target domain. Furthermore, the testimony of these surveys will also justify the usage of the RL approaches employing the said fairness notions. This point is further discussed in the challenges given in Section 5.1 and 5.2.

5 PRESENT CHALLENGES AND FUTURE DIRECTIONS

In this section, we describe major challenges in fair-RL research and discuss possible solutions to overcome them.

5.1 Correcting Societal Biases

In this article, we reviewed a number of computational approaches aiming to alleviate societal biases. However, it is important to recognize and take into consideration the fact that the presence of biases in automated decisions is due to issues such as unequal access to resources and social conditioning [58, 125]. This is true for all forms of learning; nevertheless, it assumes additional relevance for RL methods as they traditionally learn from their own experiences without a knowledgable external supervisor. The absence of such a supervisor precludes the possibility to externally correct biases during the decision-making process.

Since individuals should not be held responsible for the attributes they can not change or had no say in (i.e., sensitive attributes), the automated decisions, which affect the social benefits they receive and their prospects in life, should not depend upon those attributes. If RL algorithms are used to make decisions about the benefits whose allocation exhibits discrimination for some people owing to the attributes they had no say in, then it is reasonable that the algorithms should offset the existing discrimination due to such attributes. Of course, the obvious difficulty lies in determining which attributes that an individual has no say in. Education, at first glance, might seem like an attribute that an individual can choose. However, several studies show that the attributes that an individual has no say in (e.g., birth-place, race, caste) can impact the level of their education. For example, Jacobs [60] provide a comprehensive survey on education in the USA, which demonstrates that women are particularly disadvantaged with respect to the outcomes of education. This can be seen in academia in the USA as women make up only 26% of full professors, 23% of university presidents, and 14% of presidents of doctoral degree-granting institutions [102]. Similar gender disparity at higher echelons of positions is found in politics and business. Moreover, this is despite the fact that women fare relatively well in access to education in the USA. Indeed, 57% of all the college students in the USA are women [102]. However, easy access to education for women is not a universal phenomenon, and discrimination is present in many parts of the world as confirmed by the Department of Economic and Social Affairs, the United Nations [124, Chapter 3]. This leads to us a key insight that the level of discrimination varies with the domain (or even for different issues within a single domain as seen above) and the place of interest. Thus, while proposing fair-RL (and in general fair-ML) algorithms, credible surveys appraising the discrimination caused by the protected attributes in the target domain should be taken into consideration.

Which fairness approaches should be used for tasks corresponding to a particular social benefit should also depend upon whether the benefit in question can be considered to be a basic human right. For domains like affordable housing, essential health-care, and basic education, fairness approaches that actively try to remove disparity caused due to protected attributes and provide benefits to all the individuals should be considered. However, for other domains which require qualifications not evenly distributed in the population, a justification could be made for relaxing these stipulations. At the same time, independent efforts could be made to diffuse the ability to have such qualifications evenly in the population.

5.2 Group Fairness or Individual Fairness?

While individual fairness and group fairness remain the most popular approaches in fair-RL, their applicability has been called into question. It has been argued that group fairness measures may miss unfairness against people as a result

of their membership to multiple protected groups [33], or groups which are not (yet) defined in anti-discrimination laws but may need protection [131]. It has also been noted that many of the most promising group fairness constraints are mutually incompatible [23]. If group fairness measures are applied to protected groups, such as race and gender, separately, they might permit unfairness for structured combinations of those groups (e.g., black women), also known as *fairness gerrymandering* [66]. Moreover, if group fairness is applied across many different combinations of protected characteristics, it does not scale well with a large number of groups and may lead to overfitting [52].

On the other hand, Fleisher [41] present the following critiques of individual fairness – i) counterexamples show that similar treatment guaranteed by individual fairness is insufficient to guarantee fairness, ii) the used similarity metrics/arbiters may suffer from implicit systematic biases, iii) it cannot offer a substantive, non-circular definition of fairness because determining which features are task-relevant and thus apt for measuring similarity requires making moral judgments about what fairness constitutes, and iv) if incommensurable moral values are relevant for determining similarity for a task, similarity cannot be represented as a distance metric.

Furthermore, a recent work argues that the conflict between group fairness and individual fairness is more of an artifact of the blunt application of fairness measures, rather than a matter of conflicting principles [15]. Zemel et al. [142] also suggest pairing individual fairness with statistical parity conditions. While the stated aim in Lahoti et al. [70] is individual fairness, they claim that their approach indirectly improves group fairness because it reduces information on protected attributes.

How to disentangle the apparent conflict between group fairness and individual fairness and improve their applicability remains an important challenge for fair-RL and fair-ML researchers and practitioners.

5.3 Explainability in RL

Another important issue to consider is the opacity of many RL systems since, unlike other forms of learning, RL algorithms work using a trial-and-error method without human interaction. It has been shown that transparency not only increases users' trust [47] but also user's acceptance of a system [53]. The importance of explainability to RL systems is crucial from a legal perspective (see GDPR [48]), and it also enhances their usefulness [73].

The field of eXplainable AI (XAI) is aimed at verifying the fairness of ML systems by providing insights into how the machine learning model generates predictions. XAI has mainly focused on supervised learning, but recent efforts in eXplainable Reinforcement Learning (XRL) are specific to reinforcement learning. Verma et al. [129] introduce an inherently interpretable model framework, called Programmatically Interpretable Reinforcement Learning (PIRL), that represents policies using a high-level, human-readable programming language. This is achieved by limiting the set of target policies to conform to a policy sketch: a grammar of expressions over atoms and sequences of atoms, essentially similar to regularization. Liu et al. [81] introduce an XRL post-hoc explainability approach that mimics the model's Q-function using Linear Model U-Trees (LMUTs): simple and understandable models designed to approximate continuous functions. Shu et al. [112] propose a framework for multi-task reinforcement learning which is capable of composing hierarchical plans in an interpretable manner. Madumal et al. [87] present a method to generate explanations of the behavior of RL agents based on counterfactual analysis of the causal model. For more details on explainable reinforcement learning, we refer to the survey by Puiutta and Veith [103].

Notwithstanding these works, explainability for the intended target audience is often neglected in the development of RL methods [1] and remains a credible challenge for the fair-RL community. It is also important to cater the form and presentation of explanations to the needs of the target audience as it has been shown that human users favor some forms of explanations over others [92].

6 CONCLUDING REMARKS

Fairness in reinforcement learning is a rapidly-growing emerging field, where newer solutions are revealing newer challenges, and there are immense opportunities to take this discipline forward theoretically as well as empirically. At this juncture, we present an extensive survey covering the full breadth of fair-RL. The organization of this survey into various aspects of RL and applications of fair-RL is with the intention of aiding researchers and practitioners to gain a thorough understanding of the field. Towards the end, we also discuss a few of the most important challenges, which we hope will lead to productive conversations and advance the field of fair-RL.

REFERENCES

[1] Ashraf Abdul, Jo Vermeulen, Danding Wang, Brian Y. Lim, and Mohan Kankanhalli. 2018. *Trends and Trajectories for Explainable, Accountable and Intelligible Systems: An HCI Research Agenda.* Association for Computing Machinery, New York, NY, USA, 1–18. https://doi.org/10.1145/3173574.3174156

[2] Academy of European Law. 2018. *Handbook on European non-discrimination law.*

[3] Mridul Agarwal and Vaneet Aggarwal. 2019. Reinforcement Learning for Joint Optimization of Multiple Rewards. *arXiv preprint arXiv:1909.02940* (2019).

[4] Saamer Akhshabi, Lakshmi Anantakrishnan, Ali C Begen, and Constantine Dovrolis. 2012. What happens when HTTP adaptive streaming players compete for bandwidth?. In *Proceedings of the 22nd international workshop on Network and Operating System Support for Digital Audio and Video.* 9–14.

[5] Atefeh Hajijamali Arani, Peng Hu, and Yeying Zhu. 2021. Fairness-aware Link Optimization for Space-Terrestrial Integrated Networks: A Reinforcement Learning Framework. *IEEE Access* (2021).

[6] Nasim Arianpoo and Victor CM Leung. 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 (2016), 1–15.

[7] Ashay Athalye, Shannon Hwang, and Siddharth Nayak. [n. d.]. Fairness and Robustness of Mixed Autonomous Traffic Control with Reinforcement Learning. ([n. d.]).

[8] James Atwood, Hansa Srinivasan, Yoni Halpern, and David Sculley. 2019. Fair treatment allocations in social networks. *arXiv preprint arXiv:1911.05489* (2019).

[9] Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. 2002. Finite-Time Analysis of the Multiarmed Bandit Problem. *Mach. Learn.* 47, 2–3 (2002), 235–256.

[10] Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. 2003. The Nonstochastic Multiarmed Bandit Problem. *SIAM J. Comput.* 32, 1 (jan 2003), 48–77. https://doi.org/10.1137/S0097539701398375

[11] Wenhang Bao. 2019. Fairness in Multi-agent Reinforcement Learning for Stock Trading. *arXiv preprint arXiv:2001.00918* (2019).

[12] Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2019. *Fairness and Machine Learning.* fairmlbook.org. http://www.fairmlbook.org.

[13] Solon Barocas and Andrew D. Selbst. 2016. Big Data's Disparate Impact. *California Law Review* 104, 3 (2016), 671–732. http://www.jstor.org/stable/24758720

[14] Asia J Biega, Krishna P Gummadi, and Gerhard Weikum. 2018. Equity of attention: Amortizing individual fairness in rankings. In *The 41st international acm sigir conference on research & development in information retrieval.* 405–414.

[15] Reuben Binns. 2020. On the Apparent Conflict between Individual and Group Fairness. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Barcelona, Spain) *(FAT* '20).* Association for Computing Machinery, New York, NY, USA, 514–524. https://doi.org/10.1145/3351095.3372864

[16] Robin Burke, Nasim Sonboli, and Aldo Ordonez-Gauger. 2018. Balanced neighborhoods for multi-sided fairness in recommendation. In *Conference on Fairness, Accountability and Transparency.* PMLR, 202–214.

[17] Lucian Busoniu, Robert Babuska, and Bart De Schutter. 2006. Multi-Agent Reinforcement Learning: A Survey. In *2006 9th International Conference on Control, Automation, Robotics and Vision.* 1–6. https://doi.org/10.1109/ICARCV.2006.345353

[18] L. Elisa Celis, Sayash Kapoor, Farnood Salehi, and Nisheeth Vishnoi. 2019. Controlling Polarization in Personalization: An Algorithmic Framework. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (Atlanta, GA, USA) *(FAT* '19).* Association for Computing Machinery, New York, NY, USA, 160–169. https://doi.org/10.1145/3287560.3287601

[19] Irene Y. Chen, Fredrik D. Johansson, and David Sontag. 2018. Why is My Classifier Discriminatory?. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems* (Montréal, Canada) *(NIPS'18).* Curran Associates Inc., Red Hook, NY, USA, 3543–3554.

[20] Jingdi Chen, Yimeng Wang, and Tian Lan. 2021. Bringing Fairness to Actor-Critic Reinforcement Learning for Network Utility Optimization. In *IEEE INFOCOM 2021-IEEE Conference on Computer Communications.* IEEE, 1–10.

[21] Xinwei Chen, Tong Wang, Barrett W Thomas, and Marlin W Ulmer. 2020. Same-day delivery with fairness. *arXiv preprint arXiv:2007.09541* (2020).

[22] Yifang Chen, Alex Cuellar, Haipeng Luo, Jignesh Modi, Heramb Nemlekar, and Stefanos Nikolaidis. 2020. The Fair Contextual Multi-Armed Bandit. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS '20)*. 1810–1812.

[23] Alexandra Chouldechova. 2017. Fair Prediction with Disparate Impact: A Study of Bias in Recidivism Prediction Instruments. *Big Data* 5, 2 (2017), 153–163. https://doi.org/10.1089/big.2016.0047

[24] Houston Claure, Yifang Chen, Jignesh Modi, Malte Jung, and Stefanos Nikolaidis. 2019. Reinforcement learning with fairness constraints for resource distribution in human-robot teams. *arXiv preprint arXiv:1907.00313* (2019).

[25] Houston Claure, Yifang Chen, Jignesh Modi, Malte Jung, and Stefanos Nikolaidis. 2020. Multi-Armed Bandits with Fairness Constraints for Distributing Resources to Human Teammates. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*. 299–308.

[26] Ioan Sorin Comşa, Mehmet Aydin, Sijing Zhang, Pierre Kuonen, Jean-Frederic Wagen, and Yao Lu. 2014. Scheduling policies based on dynamic throughput and fairness tradeoff control in LTE-A networks. In *39th Annual IEEE conference on local computer networks*. IEEE, 418–421.

[27] Ioan Sorin Comşa, Sijing Zhang, Mehmet Aydin, Jianping Chen, Pierre Kuonen, and Jean-Frederic Wagen. 2014. Adaptive proportional fair parameterization based LTE scheduling using continuous actor-critic reinforcement learning. In *2014 IEEE global communications conference*. IEEE, 4387–4393.

[28] Ioan-Sorin Comşa, Sijing Zhang, Mehmet Aydin, Pierre Kuonen, Ramona Trestian, and Gheorghiță Ghinea. 2019. A comparison of reinforcement learning algorithms in fairness-oriented OFDMA schedulers. *Information* 10, 10 (2019), 315.

[29] Ioan Sorin Comsa, Sijing Zhang, Mehmet Aydin, Pierre Kuonen, and Jean-Frederic Wagen. 2012. A novel dynamic Q-learning-based scheduler technique for LTE-advanced technologies using neural networks. In *37th Annual IEEE Conference on Local Computer Networks*. IEEE, 332–335.

[30] Sam Corbett-Davies and Sharad Goel. 2018. The Measure and Mismeasure of Fairness: A Critical Review of Fair Machine Learning. arXiv:1808.00023 [cs.CY]

[31] Amanda Coston, Alan Mishler, Edward H Kennedy, and Alexandra Chouldechova. 2020. Counterfactual risk assessments, evaluation, and fairness. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 582–593.

[32] Elliot Creager, David Madras, Toniann Pitassi, and Richard Zemel. 2020. Causal Modeling for Fairness In Dynamical Systems. 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, 2185–2195. https://proceedings.mlr.press/v119/creager20a.html

[33] Kimberle Crenshaw. 1991. Mapping the Margins: Intersectionality, Identity Politics, and Violence against Women of Color. *Stanford Law Review* 43, 6 (1991), 1241–1299. http://www.jstor.org/stable/1229039

[34] Alexander D'Amour, Hansa Srinivasan, James Atwood, Pallavi Baljekar, D. Sculley, and Yoni Halpern. 2020. Fairness is Not Static: Deeper Understanding of Long Term Fairness via Simulation Studies. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Barcelona, Spain) *(FAT* '20)*. Association for Computing Machinery, New York, NY, USA, 525–534. https://doi.org/10.1145/3351095.3372878

[35] Alexander D'Amour, Hansa Srinivasan, James Atwood, Pallavi Baljekar, D Sculley, and Yoni Halpern. 2020. Fairness is not static: deeper understanding of long term fairness via simulation studies. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 525–534.

[36] Antônio C de Arruda, Alessandro F Leite, Cicero RF de Almeida, Antônio MF Crespo, and Li Weigang. 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*. IEEE, 539–544.

[37] Johan De Vriendt, Danny De Vleeschauwer, and David Robinson. 2013. Model for estimating QoE of video delivered using HTTP adaptive streaming. In *2013 IFIP/IEEE International Symposium on Integrated Network Management (IM 2013)*. IEEE, 1288–1293.

[38] Ashwini Deshpande. 2005. Affirmative action in India and the United States. (2005).

[39] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness Through Awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (ITCS '12)*.

[40] Salma Elmalaki. 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.

[41] Will Fleisher. 2021. What's Fair about Individual Fairness?. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society* (Virtual Event, USA) *(AIES '21)*. Association for Computing Machinery, New York, NY, USA, 480–490. https://doi.org/10.1145/3461702.3462621

[42] Mirko Franceschinis, Marco Mellia, Michela Meo, and Maurizio Munafo. 2005. Measuring TCP over WiFi: A real case. In *1st workshop on Wireless Network Measurements (Winmee), Riva Del Garda, Italy*. Citeseer.

[43] Pratik Gajane and Mykola Pechenizkiy. 2018. On formalizing fairness in prediction with machine learning. *5th Workshop on Fairness, Accountability, and Transparency in Machine Learning (FAT/ML) 2018* (2018).

[44] AmirEmad Ghassami, Sajad Khodadadian, and Negar Kiyavash. 2018. Fairness in Supervised Learning: An Information Theoretic Approach. In *2018 IEEE International Symposium on Information Theory (ISIT)*. 176–180. https://doi.org/10.1109/ISIT.2018.8437807

[45] Mohammad Ghodsi and Amirmahdi Mirfakhar. 2021. Online Fair Revenue Maximizing Cake Division with Non-Contiguous Pieces in Adversarial Bandits. *CoRR* abs/2111.14387 (2021). arXiv:2111.14387 https://arxiv.org/abs/2111.14387

[46] Stephen Gillen, Christopher Jung, Michael Kearns, and Aaron Roth. 2018. Online Learning with an Unknown Fairness Metric. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems* (Montréal, Canada) *(NIPS'18)*. Curran Associates Inc., Red Hook, NY, USA, 2605–2614.

[47] Alyssa Glass, Deborah L. McGuinness, and Michael Wolverton. 2008. Toward Establishing Trust in Adaptive Agents. In *Proceedings of the 13th International Conference on Intelligent User Interfaces* (Gran Canaria, Spain) *(IUI '08)*. Association for Computing Machinery, New York, NY, USA,

227–236. https://doi.org/10.1145/1378773.1378804

[48] Bryce Goodman and Seth Flaxman. 2017. European Union Regulations on Algorithmic Decision-Making and a "Right to Explanation". *AI Magazine* 38, 3 (2017).

[49] Niko A Grupen, Bart Selman, and Daniel D Lee. 2021. Fairness for Cooperative Multi-Agent Learning with Equivariant Policies. *arXiv preprint arXiv:2106.05727* (2021).

[50] Moritz Hardt, Eric Price, , and Nati Srebro. 2016. Equality of Opportunity in Supervised Learning. In *Advances in Neural Information Processing Systems 29*. 3315–3323.

[51] F. Maxwell Harper and Joseph A. Konstan. 2015. The MovieLens Datasets: History and Context. *ACM Trans. Interact. Intell. Syst.* 5, 4, Article 19 (Dec. 2015), 19 pages. https://doi.org/10.1145/2827872

[52] Ursula Hebert-Johnson, Michael Kim, Omer Reingold, and Guy Rothblum. 2018. Multicalibration: Calibration for the (Computationally-Identifiable) Masses. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer Dy and Andreas Krause (Eds.). PMLR, 1939–1948. https://proceedings.mlr.press/v80/hebert-johnson18a.html

[53] Jonathan L. Herlocker, Joseph A. Konstan, and John Riedl. 2000. Explaining Collaborative Filtering Recommendations. In *Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work* (Philadelphia, Pennsylvania, USA) *(CSCW '00)*. Association for Computing Machinery, New York, NY, USA, 241–250. https://doi.org/10.1145/358916.358995

[54] Mitchell Hoffman, Lisa B Kahn, and Danielle Li. 2015. *Discretion in Hiring*. Working Paper 21709. National Bureau of Economic Research. https://doi.org/10.3386/w21709

[55] H.J. Holzer. 2007. The economic impact of affirmative action in the US. *Swedish Economic Policy Review* 14 (03 2007), 41–71.

[56] Wen Huang, Kevin Labille, Xintao Wu, Dongwon Lee, and Neil Heffernan. 2020. Achieving User-Side Fairness in Contextual Bandits. *CoRR* abs/2010.12102 (2020). arXiv:2010.12102 https://arxiv.org/abs/2010.12102

[57] Wen Huang, Lu Zhang, and Xintao Wu. 2021. Achieving Counterfactual Fairness for Causal Bandit. arXiv:2109.10458 [cs.LG]

[58] P J Bickel, EA Hammel, and J W O'connell. 1975. Sex Bias in Graduate Admissions: Data from Berkeley. 187 (03 1975), 398–404.

[59] Shahin Jabbari, Matthew Joseph, Michael Kearns, Jamie Morgenstern, and Aaron Roth. 2017. Fairness in Reinforcement Learning. In *Proceedings of the 34th International Conference on Machine Learning*. 1617–1626.

[60] Jerry Jacobs. 2003. Gender Equality and Higher Education. 22 (11 2003), 153–185.

[61] Jiechuan Jiang and Zongqing Lu. 2019. Learning Fairness in Multi-Agent Systems. *Advances in Neural Information Processing Systems* 32 (2019), 13854–13865.

[62] Matthew Joseph, Michael Kearns, Jamie H Morgenstern, and Aaron Roth. 2016. Fairness in Learning: Classic and Contextual Bandits. In *Advances in Neural Information Processing Systems*.

[63] Matthew Joseph, Michael J. Kearns, Jamie Morgenstern, Seth Neel, and Aaron Roth. 2016. Fair Algorithms for Infinite and Contextual Bandits. *CoRR* abs/1610.09559 (2016). arXiv:1610.09559 http://arxiv.org/abs/1610.09559

[64] Sampath Kannan, Aaron Roth, and Juba Ziani. 2019. Downstream effects of affirmative action. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 240–248.

[65] Emilie Kaufmann. 2014. *Analysis of bayesian and frequentist strategies for sequential resource allocation*. Theses. Télécom ParisTech. https://pastel.archives-ouvertes.fr/tel-01413183

[66] Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. 2018. Preventing Fairness Gerrymandering: Auditing and Learning for Subgroup Fairness. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer Dy and Andreas Krause (Eds.). PMLR, 2564–2572. https://proceedings.mlr.press/v80/kearns18a.html

[67] Sajad Khodadadian, AmirEmad Ghassami, and Negar Kiyavash. 2021. Impact of Data Processing on Fairness in Supervised Learning. In *2021 IEEE International Symposium on Information Theory (ISIT)*. 2643–2648. https://doi.org/10.1109/ISIT45174.2021.9517766

[68] Pauline T. Kim. 2017. Data-Driven Discrimination at Work. *William & Mary Law Review* 58:857 (2017). Issue 3.

[69] Akhil Krishnakumar. 2019. *Assessing the Fairness of AI Recruitment systems*. Master's thesis.

[70] Preethi Lahoti, Krishna P. Gummadi, and Gerhard Weikum. 2019. iFair: Learning Individually Fair Data Representations for Algorithmic Decision Making. In *2019 IEEE 35th International Conference on Data Engineering (ICDE)*. 1334–1345. https://doi.org/10.1109/ICDE.2019.00121

[71] Finnian Lattimore, Tor Lattimore, and Mark D. Reid. 2016. Causal Bandits: Learning Good Interventions via Causal Inference. In *Proceedings of the 30th International Conference on Neural Information Processing Systems* (Barcelona, Spain) *(NIPS'16)*. Curran Associates Inc., Red Hook, NY, USA, 1189–1197.

[72] Tor Lattimore and Csaba Szepesvári. 2020. *Bandit Algorithms*. Cambridge University Press. https://doi.org/10.1017/9781108571401

[73] Jung Hoon Lee. 2019. Complementary reinforcement learning towards explainable agents. *CoRR* abs/1901.00188 (2019). arXiv:1901.00188 http://arxiv.org/abs/1901.00188

[74] Chenghao Li, Xiaoteng Ma, Li Xia, Qianchuan Zhao, and Jun Yang. 2020. Fairness Control of Traffic Light via Deep Reinforcement Learning. In *2020 IEEE 16th International Conference on Automation Science and Engineering (CASE)*. IEEE, 652–658.

[75] Danielle Li, Lindsey R Raymond, and Peter Bergman. 2020. *Hiring as Exploration*. Working Paper 27736. National Bureau of Economic Research. https://doi.org/10.3386/w27736

[76] Fengjiao Li, Jia Liu, and Bo Ji. 2019. Combinatorial Sleeping Bandits with Fairness Constraints. In *IEEE INFOCOM 2019 - IEEE Conference on Computer Communications*. 1702–1710. https://doi.org/10.1109/INFOCOM.2019.8737461

[77] Lihong Li, Wei Chu, John Langford, and Robert E. Schapire. 2010. A Contextual-Bandit Approach to Personalized News Article Recommendation. In *Proceedings of the 19th International Conference on World Wide Web* (Raleigh, North Carolina, USA) *(WWW '10)*. Association for Computing Machinery, New York, NY, USA, 661–670. https://doi.org/10.1145/1772690.1772758

[78] Lihong Li, Michael L. Littman, Thomas J. Walsh, and Alexander L. Strehl. 2011. Knows what it knows: aÂ framework forÂ self-aware learning. *Machine Learning* 82, 3 (Mar 2011), 399–443.

[79] Lihong Li, Yu Lu, and Dengyong Zhou. 2017. Provably Optimal Algorithms for Generalized Linear Contextual Bandits. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70* (Sydney, NSW, Australia) *(ICML'17)*. JMLR.org, 2071–2080.

[80] Chi Harold Liu, Zheyu Chen, Jian Tang, Jie Xu, and Chengzhe Piao. 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 (2018), 2059–2070.

[81] Guiliang Liu, Oliver Schulte, Wang Zhu, and Qingcan Li. 2018. Toward interpretable deep reinforcement learning with linear model u-trees. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 414–429.

[82] Lydia T. Liu, Sarah Dean, Esther Rolf, Max Simchowitz, and Moritz Hardt. 2018. Delayed Impact of Fair Machine Learning. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer Dy and Andreas Krause (Eds.). PMLR, 3150–3158. https://proceedings.mlr.press/v80/liu18c.html

[83] Lydia T Liu, Sarah Dean, Esther Rolf, Max Simchowitz, and Moritz Hardt. 2018. Delayed impact of fair machine learning. In *International Conference on Machine Learning*. PMLR, 3150–3158.

[84] Weiwen Liu, Feng Liu, Ruiming Tang, Ben Liao, Guangyong Chen, and Pheng Ann Heng. 2020. Balancing between accuracy and fairness for interactive recommendation with reinforcement learning. *Advances in Knowledge Discovery and Data Mining* 12084 (2020), 155.

[85] Yang Liu, Goran Radanovic, Christos Dimitrakakis, Debmalya Mandal, and David C. Parkes. 2017. Calibrated fairness in Bandits. In *Proceedings of the 4th Workshop on Fairness, Accountability, and Transparency in Machine Learning (Fat/ML 2017)*. https://arxiv.org/abs/1707.01875

[86] Yangyi Lu, Amirhossein Meisami, Ambuj Tewari, and Zhenyu Yan. 2020. Regret Analysis of Bandit Problems with Causal Background Knowledge. In *Proceedings of the 36th Annual Conference on Uncertainty in Artificial Intelligence*. http://www.auai.org/uai2020/proceedings/77_main_paper.pdf

[87] Prashan Madumal, Tim Miller, Liz Sonenberg, and Frank Vetere. 2019. Explainable Reinforcement Learning Through a Causal Lens. *CoRR* abs/1905.10958 (2019). arXiv:1905.10958 http://arxiv.org/abs/1905.10958

[88] Dean B. McFarlin and Paul D. Sweeney. 1992. Distributive and Procedural Justice as Predictors of Satisfaction with Personal and Organizational Outcomes. *The Academy of Management Journal* 35, 3 (1992), 626–637.

[89] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. A Survey on Bias and Fairness in Machine Learning. *ACM Comput. Surv.* 54, 6, Article 115 (July 2021), 35 pages. https://doi.org/10.1145/3457607

[90] Aditya Krishna Menon and Robert C Williamson. 2018. The cost of fairness in binary classification. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency (Proceedings of Machine Learning Research, Vol. 81)*, Sorelle A. Friedler and Christo Wilson (Eds.). PMLR, 107–118. https://proceedings.mlr.press/v81/menon18a.html

[91] Blossom Metevier, Stephen Giguere, Sarah Brockman, Ari Kobren, Yuriy Brun, Emma Brunskill, and Philip Thomas. 2019. Offline contextual bandits with high probability fairness guarantees. *Advances in neural information processing systems* 32 (2019).

[92] Tim Miller. 2019. Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence* 267 (2019), 1–38. https://doi.org/10.1016/j.artint.2018.07.007

[93] Smitha Milli, John Miller, Anca D Dragan, and Moritz Hardt. 2019. The social cost of strategic classification. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 230–239.

[94] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing Atari With Deep Reinforcement Learning. In *NIPS Deep Learning Workshop*.

[95] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *nature* 518, 7540 (2015), 529–533.

[96] Razieh Nabi, Daniel Malinsky, and Ilya Shpitser. 2019. Learning optimal fair policies. In *International Conference on Machine Learning*. PMLR, 4674–4682.

[97] Vishakha Patil, Ganesh Ghalme, Vineet Nair, and Y. Narahari. 2021. Achieving Fairness in the Stochastic Multi-Armed Bandit Problem. *Journal of Machine Learning Research* 22, 174 (2021), 1–31. http://jmlr.org/papers/v22/20-704.html

[98] Judea Pearl. 2000. *Causality: models, reasoning and inferenc*. MIT Press, Cambridge.

[99] Dana Pessach, Gonen Singer, Dan Avrahami, Hila Chalutz Ben-Gal, Erez Shmueli, and Irad Ben-Gal. 2020. Employees recruitment: A prescriptive analytics approach via machine learning and mathematical programming. *Decision Support Systems* 134 (2020), 113290. https://doi.org/10.1016/j.dss.2020.113290

[100] Stefano Petrangeli, Maxim Claeys, Steven Latré, Jeroen Famaey, and Filip De Turck. 2014. A multi-agent Q-Learning-based framework for achieving fairness in HTTP Adaptive Streaming. In *2014 IEEE Network Operations and Management Symposium (NOMS)*. IEEE, 1–9.

[101] Alexander Peysakhovich and Adam Lerer. 2018. Prosocial Learning Agents Solve Generalized Stag Hunts Better than Selfish Ones. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*. 2043–2044.

[102] White House Project. 2009. *The White House Project Report: Benchmarking Women's Leadership*. White House Project.

[103] Erika Puiutta and Eric MSP Veith. 2020. Explainable reinforcement learning: A survey. In *International Cross-Domain Conference for Machine Learning and Knowledge Extraction*. Springer, 77–95.

[104] Martin L. Puterman. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming* (1st ed.). John Wiley & Sons, Inc., USA.

[105] Hang Qi, Zhiqun Hu, Hao Huang, Xiangming Wen, and Zhaoming Lu. 2020. Energy efficient 3-D UAV control for persistent communication service and fairness: A deep reinforcement learning approach. *IEEE Access* 8 (2020), 53172–53184.

[106] Majid Raeis and Alberto Leon-Garcia. 2021. A deep reinforcement learning approach for fair traffic signal control. In *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2512–2518.

[107] Manish Raghavan and Solon Barocas. 2019. *Challenges for mitigating bias in algorithmic hiring*. Technical Report. The Brookings Institution's Artificial Intelligence and Emerging Technology (AIET).

[108] Ronald E. Robertson, David Lazer, and Christo Wilson. 2018. Auditing the Personalization and Composition of Politically-Related Search Engine Results Pages. In *Proceedings of the 2018 World Wide Web Conference* (Lyon, France) *(WWW '18)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 955–965. https://doi.org/10.1145/3178876.3186143

[109] Kit T. Rodolfa, Hemank Lamba, and Rayid Ghani. 2021. Empirical observation of negligible fairness–accuracy trade-offs in machine learning for public policy. *Nature Machine Intelligence* 3, 10 (01 Oct 2021), 896–904. https://doi.org/10.1038/s42256-021-00396-x

[110] Candice Schumann, Jeffrey S. Foster, Nicholas Mattei, and John P. Dickerson. 2020. We Need Fairness and Explainability in Algorithmic Hiring. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems* (Auckland, New Zealand) *(AAMAS '20)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1716–1720.

[111] Candice Schumann, Zhi Lang, Nicholas Mattei, and John P. Dickerson. 2019. Group Fairness in Bandit Arm Selection. *CoRR* abs/1912.03802 (2019). arXiv:1912.03802 http://arxiv.org/abs/1912.03802

[112] Tianmin Shu, Caiming Xiong, and Richard Socher. 2017. Hierarchical and Interpretable Skill Acquisition in Multi-task Reinforcement Learning. *CoRR* abs/1712.07294 (2017). arXiv:1712.07294 http://arxiv.org/abs/1712.07294

[113] Umer Siddique, Paul Weng, and Matthieu Zimmer. 2020. Learning Fair Policies in Multi-Objective (Deep) Reinforcement Learning with Average and Discounted Rewards. In *International Conference on Machine Learning*. PMLR, 8905–8915.

[114] Ashudeep Singh, Yoni Halpern, Nithum Thain, Konstantina Christakopoulou, EH Chi, Jilin Chen, and Alex Beutel. 2020. Building healthy recommendation sequences for everyone: A safe reinforcement learning approach. In *FAccTRec Workshop*.

[115] Ashudeep Singh and Thorsten Joachims. 2018. Fairness of exposure in rankings. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2219–2228.

[116] Daniel P. Skarlicki and Robert Folger. 1997. Retaliation in the workplace: The roles of distributive, procedural, and interactional justice. *Journal of Applied Psychology* 82, 3 (1997), 434–443.

[117] Aleksandrs Slivkins. 2019. Introduction to Multi-Armed Bandits. *Foundations and Trends® in Machine Learning* 12, 1-2 (2019), 1–286. https://doi.org/10.1561/2200000068

[118] Till Speicher, Muhammad Ali, Giridhari Venkatadri, Filipe Nunes Ribeiro, George Arvanitakis, Fabrício Benevenuto, Krishna P. Gummadi, Patrick Loiseau, and Alan Mislove. 2018. Potential for Discrimination in Online Targeted Advertising. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency (Proceedings of Machine Learning Research, Vol. 81)*, Sorelle A. Friedler and Christo Wilson (Eds.). PMLR, 5–19. https://proceedings.mlr.press/v81/speicher18a.html

[119] Harini Suresh and John V. Guttag. 2019. A Framework for Understanding Unintended Consequences of Machine Learning. *CoRR* abs/1901.10002 (2019). arXiv:1901.10002 http://arxiv.org/abs/1901.10002

[120] Richard S. Sutton and Andrew G. Barto. 2018. *Reinforcement Learning: An Introduction* (second ed.). MIT Press.

[121] Csaba Szepesvári. 2010. *Algorithms for Reinforcement Learning*. Morgan and Claypool. https://doi.org/10.2200/S00268ED1V01Y201005AIM009

[122] Mohammad Sadegh Talebi and Alexandre Proutiere. 2018. Learning Proportionally Fair Allocations with Low Regret. *Proc. ACM Meas. Anal. Comput. Syst.* 2, 2, Article 36 (June 2018), 31 pages. https://doi.org/10.1145/3224431

[123] Maryam Tavakol. 2020. Fair Classification with Counterfactual Learning. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2073–2076.

[124] the Department of Economic and Social Affairs, the United Nations. 2015. The World's Women 2015 Trends and Statistics.

[125] Sukhadeo Thorat and Katherine S. Neuman (Eds.). 2012. *Blocked by Caste: Economic Discrimination in Modern India*. Oxford University Press.

[126] Le Tong, Yangyi Chen, Xin Zhou, and Yifu Sun. 2021. QoE-Fairness Tradeoff Scheme for Dynamic Spectrum Allocation Based on Deep Reinforcement Learning. In *The 5th International Conference on Computer Science and Application Engineering*. 1–7.

[127] Ruibo Tu, Xueru Zhang, Yang Liu, Hedvig Kjellström, Mingyan Liu, Kun Zhang, and Cheng Zhang. 2020. How Do Fair Decisions Fare in Long-term Qualification?. In *Thirty-fourth Conference on Neural Information Processing Systems*.

[128] Anastasios Valkanis, Georgios Papadimitriou, Georgia Beletsioti, Emmanouel Varvarigos, and Petros Nicopolitidis. 2022. Efficiency and fairness improvement for elastic optical networks using reinforcement learning-based traffic prediction. *Journal of Optical Communications and Networking* 14, 3 (2022), 25–42.

[129] Abhinav Verma, Vijayaraghavan Murali, Rishabh Singh, Pushmeet Kohli, and Swarat Chaudhuri. 2018. Programmatically interpretable reinforcement learning. In *International Conference on Machine Learning*. PMLR, 5045–5054.

[130] Eugene Vinitsky, Aboudy Kreidieh, Luc Le Flem, Nishant Kheterpal, Kathy Jang, Cathy Wu, Fangyu Wu, Richard Liaw, Eric Liang, and Alexandre M Bayen. 2018. Benchmarks for reinforcement learning in mixed-autonomy traffic. In *Conference on robot learning*. PMLR, 399–409.

[131] Sandra Wachter and Brent Mittelstadt. 2019. A Right to Reasonable Inferences Re-Thinking Data Protection Law in the Age of Big Data and AI. *Columbia Business Law Review* 2019, 2 (2019), 494–620.

[132] Guang Wang, Shuxin Zhong, Shuai Wang, Fei Miao, Zheng Dong, and Desheng Zhang. 2021. Data-Driven Fairness-Aware Vehicle Displacement for Large-Scale Electric Taxi Fleets. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*. IEEE, 1200–1211.

[133] Jane X Wang, Edward Hughes, Chrisantha Fernando, Wojciech M Czarnecki, Edgar A Duéñez-Guzmán, and Joel Z Leibo. 2019. Evolving Intrinsic Motivations for Altruistic Behavior. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*. 683–692.

[134] Lequn Wang, Yiwei Bai, Wen Sun, and Thorsten Joachims. 2021. Fairness of Exposure in Stochastic Bandits. In *Proceedings of the 38th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 10686–10696. https://proceedings.mlr.press/v139/wang21b.html

[135] Mei Wang and Weihong Deng. 2020. Mitigating bias in face recognition using skewness-aware reinforcement learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9322–9331.

[136] Yuyan Wang, Xuezhi Wang, Alex Beutel, Flavien Prost, Jilin Chen, and Ed H. Chi. 2021. *Understanding and Improving Fairness-Accuracy Trade-Offs in Multi-Task Learning*. Association for Computing Machinery, New York, NY, USA, 1748–1757. https://doi.org/10.1145/3447548.3467326

[137] Min Wen, Osbert Bastani, and Ufuk Topcu. 2021. Algorithms for Fairness in Sequential Decision Making. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 130)*, Arindam Banerjee and Kenji Fukumizu (Eds.). PMLR, 1144–1152. https://proceedings.mlr.press/v130/wen21a.html

[138] Min Wen, Osbert Bastani, and Ufuk Topcu. 2021. Algorithms for Fairness in Sequential Decision Making. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 1144–1152.

[139] Ning Xie, Gabrielle Ras, Marcel van Gerven, and Derek Doran. 2020. Explainable Deep Learning: A Field Guide for the Uninitiated. *CoRR* abs/2004.14545 (2020). arXiv:2004.14545 https://arxiv.org/abs/2004.14545

[140] Meguru Yamazaki and Miki Yamamoto. 2021. Fairness Improvement of Congestion Control with Reinforcement Learning. *Journal of Information Processing* 29 (2021), 592–595.

[141] Mingqi Yuan, Qi Cao, Man-on Pun, and Yi Chen. 2020. Fairness-Oriented User Scheduling for Bursty Downlink Transmission Using Multi-Agent Reinforcement Learning. *arXiv preprint arXiv:2012.15081* (2020).

[142] Rich Zemel, Yu Wu, Kevin Swersky, Toni Pitassi, and Cynthia Dwork. 2013. Learning Fair Representations. In *Proceedings of the 30th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 28)*, Sanjoy Dasgupta and David McAllester (Eds.). PMLR, Atlanta, Georgia, USA, 325–333. https://proceedings.mlr.press/v28/zemel13.html

[143] Chongjie Zhang and Julie A Shah. 2014. Fairness in multi-agent sequential decision-making. In *Advances in Neural Information Processing Systems*. 2636–2644.

[144] Dell Zhang and Jun Wang. 2021. Recommendation Fairness: From Static to Dynamic. *arXiv preprint arXiv:2109.03150* (2021).

[145] Haodi Zhang, Zihang Gao, Yi Zhou, Hao Zhang, Kaishun Wu, and Fangzhen Lin. 2019. Faster and Safer Training by Embedding High-Level Knowledge into Deep Reinforcement Learning. *CoRR* abs/1910.09986 (2019). arXiv:1910.09986 http://arxiv.org/abs/1910.09986

[146] Xueru Zhang and Mingyan Liu. 2020. Fairness in Learning-Based Sequential Decision Algorithms: A Survey. *CoRR* abs/2001.04861 (2020). arXiv:2001.04861 https://arxiv.org/abs/2001.04861

[147] Han Zhao and Geoffrey J. Gordon. 2019. *Inherent Tradeoffs in Learning Fair Representations*. Curran Associates Inc., Red Hook, NY, USA.

[148] Qinyun Zhu and Jae Oh. 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)*. IEEE, 460–466.

[149] Matthieu Zimmer, Claire Glanois, Umer Siddique, and Paul Weng. 2021. Learning fair policies in decentralized cooperative multi-agent reinforcement learning. In *International Conference on Machine Learning*. PMLR, 12967–12978.

A   LIST OF SYMBOLS

$\mathcal{A}$	Set of arms/actions
C	Context
δ	Confidence in a Probably Approximate Correct(PAC) bound
Δ_{gap}	Difference between the maximum and the second maximum expected rewards in Celis et al. [18]
Δ_{gap}	Lower bound on the difference in expected reward of the optimal action and any other extremal action in a convex set seen by the learner[63]
D	Diameter of the Markov Chain in Zhu and Oh [148]
d	Dimension of a feature vector
γ	Discount factor
ϵ	Accuracy in a Probably Approximate Correct(PAC) bound

ϵ Tolerance w.r.t. demographic parity in Wen et al. [137]

η Design parameter which controls the priority of meeting the fairness requirement in Li et al. [76]

κ $1 - rad\sqrt{\frac{2\log{(2dT/\delta)}}{T\lambda}}$ [63]

λ Measure of how quickly an agent can learn from uniformly random actions [63]

m Maximum number of arms that can be played simultaneously in Li et al. [76]

K Number of agents in Zhu and Oh [148]

N Number of episodes

rad Radius of a ball

$\mathcal{S}$ Set of states

f Problem-dependent value from Jabbari et al. [59]

L Upper bound on the $L2$ norm of all the feature vectors [56]

T Time horizon

t Index of a time step

B THEORY OF FAIR-RL

Joseph et al. [63] consider contextual bandits in which each arm (i.e., individual) is represented by a vector of d features. For this setting, they propose an algorithm called RIDGEFAIR$_\text{m}$ with a regret bound of $\tilde{O}(d|\mathcal{A}|^2\sqrt{T})$. Secondly, they provide an algorithm for infinite bandits. Recall that, in this setting the learner sees a convex set of arms contained in a ball of radius rad. The provided algorithm has a regret bound of $\tilde{O}\left(\frac{rad^6\Psi^2}{\kappa^2\lambda^2\Delta_{gap}^2}\right)$ where λ measures how quickly an agent can learn from uniformly random actions, $\kappa = 1 - rad\sqrt{\frac{2\log{(2dT/\delta)}}{T\lambda}}$, and Δ_{gap} is a lower bound on the difference in expected reward of the optimal action and any other extremal action in a convex set seen by the learner.

Liu et al. [85] propose using individual fairness in stochastic bandits. To that end, they impose: i) *smooth fairness* constraint - two arms with similar distribution should be selected with similar probability, and ii) *calibrated fairness* constraint - sample each arm with probability equal to its reward being the greatest. The authors elucidate how these two constraints address the first two critiques of meritocratic fairness presented earlier. As it turns out, this notion of exact calibrated fairness is not possible to achieve during the learning phase, hence the authors aim to minimize expected cumulative amount by which the algorithm violates calibrated fairness constraint across time-steps. They propose a variation of Thompson sampling adhering to smooth fairness for total variation distance and prove a $\tilde{O}\left((|\mathcal{A}|T)^{2/3}\right)$ upper bound on fairness regret. Furthermore, they also modify their algorithm for dueling bandits and prove a $\tilde{O}\left(|\mathcal{A}|^{4/3}T^{2/3}\right)$ bound on fairness regret.

Gillen et al. [46] consider individual fairness constraints governed by a Mahalanobis similarity metric in linear contextual bandits. The authors assume that the algorithm has access to an oracle that can ascertain what it means to be fair, but cannot explicitly enunciate the fairness metric. The fairness constraint necessitates that the difference between the probabilities with which any two actions are taken is bounded by the distance between their contexts. After the learner makes its decisions, it receives feedback from the oracle specifying for which pairs of contexts the fairness constraint was violated. The goal is to guarantee near-optimal regret (with respect to the best fair policy), while violating the fairness constraints as infrequently as possible. They propose an algorithm, which – i) achieves regret of $\tilde{O}\left(|\mathcal{A}|^2d^2\log{(T)} + d\sqrt{T}\right)$ with a high probability, where d is the dimension of the context vectors, and ii) violates

the fairness constraint by more than ϵ on at most $O\left(|\mathcal{A}|^2 d^2 \log(d/\epsilon)\right)$ steps. The above guarantees are obtained by setting $\epsilon = O(1/T)$. Other trade-offs between regret and fairness violations are also possible.

In the setting explored by Schumann et al. [111], the reward for each arm belonging to protected groups is additionally corrupted with a bias term. The performance measure is cumulative regret ignoring the bias terms. The authors propose an algorithm called GroupFairTopInterval which can accommodate both their definitions of group fairness – demographic parity and proportional parity.

Wen et al. [137] develop fair decision-making policies in discounted MDPs. They consider the episodic case where the system is reset after a fixed number of steps. They distinguish between the quality of outcomes for the decision-maker and the quality of outcomes for an individual with the latter being termed as *individual rewards*. Individual rewards are assumed to be state independent. The goal is to learn an optimal policy that does not favor the majority sub-population over the minority sub-population in terms of individual rewards. In particular, their goal works with demographic parity and equal opportunity [50]. The performance bound is the expected regret in terms of cumulative reward for the decision-maker over total number of episodes N. They use explore-then-commit strategy [72] where the number of exploration episodes is computed using the number of steps in an episode, number of states, an upper bound on the reward, and other problem-dependant factors. They guarantee that the above strategy satisfies demographic parity with arbitrary tolerance ϵ and prove $O\left(\left(N^{2/3} + 1/\epsilon^2\right)\log(1/\delta)\right)$ upper bound on the expected regret with probability $1 - \delta$. A noteworthy caveat for this bound is that the reference optimal policy is required to satisfy a stricter level of fairness.

Li et al. [76] explore a variation of classical stochastic bandits called combinatorial sleeping bandits in which multiple arms (upto m) can be simultaneously played and an arm could sometimes be "sleeping" (i.e., unavailable). The fairness criterion is to have a minimum selection fraction for each individual arm. They propose a variation of UCB algorithm [9] called Learning with Fairness Guarantee (LFG) which satisfies asymptotic fairness guarantee. As a performance measure, they use time-average regret which is the difference between the expected per-step reward of the optimal policy satisfying the fairness criterion and that of the considered algorithm. They provide an *instance-independent*[3] upper bound on the time-average regret of $\frac{|\mathcal{A}|}{2\eta} + O\left(\frac{\sqrt{m|\mathcal{A}|T\log T} + |\mathcal{A}|}{T}\right)$ for LFG, where η is a design parameter which controls the priority of meeting the fairness requirement. The authors enlist real-time scheduling in wireless networks, ad placement in online advertising systems, and task assignment in crowd-sourcing platforms as practical applications and present experimental results on simulations showing competent performance of LFG.

Patil et al. [97] consider stochastic bandits where the minimum selection fraction of each arm is limited to $\left[0, \frac{1}{|\mathcal{A}|-1}\right)$. The performance measure is called fairness-aware regret in which the expected cumulative reward of the considered algorithm is compared to that of an optimal policy satisfying the fairness constraint. They propose a meta-algorithm called FAIR-Learn which either selects the arm straying the most from the fairness constraint if such an arm exists or selects an arm according to the MAB algorithm used as a black-box. The authors prove that the fairness guarantees hold uniformly over time and provide a $O(\sqrt{T\log T})$ bound on the fairness-awareness regret.

Wang et al. [134] aim to achieve merit-based fairness of exposure to the items while optimizing utility to the users in classical stochastic bandits and stochastic linear bandits. The fairness constraint requires that each arm receives an amount of exposure proportional to its merit, where merit is quantified through an application-dependent merit function given to the algorithm. They characterize the merit function – i) to have some minimum positive constant merit, and ii) to be Lipschitz continuous. The necessity of these two conditions is established by proving a $O(T)$

---

[3]See definition 7 in Slivkins [117]

lower bound on expected fairness regret when either of the conditions is relaxed. They distinguish between *reward regret* – the gap between the expected cumulative reward of a policy and that of the optimal fair policy, and *fairness regret* – the cumulative L1 distance between a policy and the optimal fair policy. For stochastic bandits, they propose a frequentist algorithm called FairX-UCB and a Bayesian algorithm called FairX-TS. For FairX-UCB, they prove a high-probability bound of $\tilde{O}\left(LipschitzConstant \cdot \sqrt{|\mathcal{A}|T}/(MinimumMerit)\right)$ on fairness regret and a high-probability bound of $\tilde{O}\left(\sqrt{|\mathcal{A}|T}\right)$ on reward regret. For FairX-TS, they prove the same respective bounds for Bayesian fairness regret and Bayesian reward regret. They further extend the algorithms to FairX-LinUCB and FairX-LinTS for stochastic linear bandits. For both algorithms, they prove a bound of $\tilde{O}\left(d\sqrt{T}\right)$ on the reward regret (resp. Bayesian reward regret) where d is the dimensionality of the context vectors. In addition, they prove a high-probability bound of $\tilde{O}\left(LipschitzConstant d\sqrt{T}/(MinimumMerit)\right)$ on fairness regret for FairX-LinUCB, and for FairX-LinTS, they prove a bound of $\tilde{O}\left(LipschitzConstant\sqrt{dT}/(MinimumMerit)\right)$ on Bayesian fairness regret. Furthermore, the experimental results on synthetic data and on real-world data from "Yahoo! Today" Module [77] show that the proposed algorithms provide better fairness in terms of exposure to the items.

Ghodsi and Mirfakhar [45] provide a RL solution for fair and efficient allocation of a divisible resource among a population with distinct preferences. They model this problem using adversarial bandits where the learner chooses the next fraction of the resource to be distributed and then an adversary specifies each individual's (i.e. arm's) valuation of that fraction. The agent allocates the fraction to one of the arms and receives a reward dependent on the selected arm's valuation. Fairness measure of a policy is the difference between the highest accumulated valuation among all the arms and the lowest. Fairness-regret of a policy is the difference between the fairness measure of an optimum offline policy and the said policy. They propose an algorithm based on EXP3 [10] which attains sub-linear bounds for both cumulative regret and fairness-regret if the adversary is restrained to only assign certain kind of evaluations.