

Fair Reinforcement Learning for Equitable Health Resource Allocation

Hakeem Adekunle* 

Department of Mathematics and Statistics, Georgia State University, Atlanta, GA, USA.

ABSTRACT

Inequitable distribution of health resources remains a persistent challenge in many healthcare systems, disproportionately affecting underserved communities, rural populations, and socioeconomically disadvantaged groups. Traditional decision-support tools often optimize for efficiency or overall population benefit, but they rarely account for structural biases that contribute to unequal access. Reinforcement Learning (RL) has emerged as a powerful approach for optimizing dynamic resource allocation; however, conventional RL models risk reinforcing existing disparities when rewards are based solely on aggregate performance metrics.

This study introduces a fairness-aware reinforcement learning framework designed to promote equitable allocation of critical health resources including hospital beds, vaccines, and diagnostic equipment across diverse demographic groups. The proposed approach incorporates explicit fairness constraints, such as counterfactual fairness and group-level equity metrics, directly into the RL reward function and training process. By combining fairness penalties with outcome-based rewards, the model balances efficiency with equity, ensuring that improved system performance does not come at the expense of vulnerable populations.

Simulation results demonstrate that the fairness-aware RL approach significantly reduces disparities in resource allocation while maintaining competitive performance compared to unconstrained RL baselines. The findings contribute to the growing field of equitable AI in healthcare by offering a scalable, adaptable framework that can be applied across various health system contexts. Overall, this work highlights the potential of fairness-constrained RL to support more just, transparent, and inclusive healthcare delivery systems.

Journal of Data Analysis and Critical Management (2025);

DOI: 10.64235/27s4qp61

INTRODUCTION

Equitable allocation of healthcare resources remains a persistent challenge for health systems worldwide. Structural inequalities rooted in socioeconomic status, geography, and demographic characteristics continue to shape disparities in access to essential services, including hospital beds, diagnostic tools, and preventive interventions. These inequities disproportionately affect vulnerable populations such as low-income communities, rural residents, and ethnic minorities, ultimately contributing to avoidable differences in health outcomes (Braveman et al., 2018; Williams & Jackson, 2020). Addressing these challenges requires not only policy-level interventions but also decision-support tools capable of accounting for both efficiency and equity in complex, resource-constrained environments.

In recent years, artificial intelligence and machine learning have been increasingly adopted to support

Corresponding Author: Hakeem Adekunle, Department of Mathematics and Statistics, Georgia State University, Atlanta, GA, USA, e-mail: Adekunletem2022@gmail.com

How to cite this article: Adekunle, H. (2025). Fair Reinforcement Learning for Equitable Health Resource Allocation. *Journal of Data Analysis and Critical Management*, 01(2):91-97.

Source of support: Nil

Conflict of interest: None

healthcare decision-making. Among these approaches, reinforcement learning (RL) has gained particular attention due to its ability to model sequential decisions and adapt to dynamic environments. RL methods have been applied to a range of healthcare problems, including intensive care unit (ICU) management, patient scheduling, treatment optimization, and hospital resource planning (Gottesman et al., 2019; Yu et al., 2021). Algorithms such as Deep Q-Networks (DQN),

Proximal Policy Optimization (PPO), and policy gradient methods enable systems to learn allocation strategies that maximize long-term rewards based on observed system states and outcomes.

Despite their promise, most reinforcement learning applications in healthcare prioritize efficiency-based objectives, such as throughput maximization or cost reduction, with limited consideration of fairness. When deployed in settings shaped by historical and structural inequalities, such models risk reinforcing or amplifying existing disparities (Obermeyer et al., 2019). Recent work in fairness-aware machine learning has highlighted how biased data, proxy variables, and optimization objectives can lead to systematically unequal outcomes across demographic groups (Verma & Rubin, 2018; Rajkomar et al., 2019).

To address these concerns, researchers have begun integrating fairness constraints into machine learning models. Concepts such as demographic parity, equalized odds, and counterfactual fairness provide formal mechanisms for evaluating and mitigating bias in algorithmic decision-making (Kusner et al., 2017; Kilbertus et al., 2020). However, the incorporation of these fairness principles into reinforcement learning—particularly in the context of healthcare resource allocation—remains limited. Existing studies often focus on static prediction tasks, while fewer explore fairness in sequential, high-stakes allocation settings.

The purpose of this study is to examine how fairness-aware reinforcement learning can be applied to healthcare resource allocation to reduce disparities while maintaining acceptable system performance. By integrating fairness constraints into standard reinforcement learning algorithms, this work aims to demonstrate that equity and efficiency need not be mutually exclusive objectives. Through simulation-based evaluation, the study compares baseline reinforcement learning approaches with fairness-aware models, highlighting their respective impacts on allocation outcomes, convergence behavior, and equity metrics. Ultimately, this research seeks to contribute to the growing body of evidence supporting responsible and equitable deployment of artificial intelligence in healthcare systems.

LITERATURE REVIEW

Health Resource Allocation and Equity

The distribution of healthcare resources has long been influenced by social, economic, and geographic disparities. Research shows that marginalized

communities including those living in poverty, remote areas, or facing racial discrimination often encounter substantial barriers to receiving timely and adequate care (Braveman et al., 2018). These barriers can manifest in many ways: fewer healthcare facilities in rural regions, limited access to specialists, higher transportation burdens, or chronic shortages of essential supplies.

Inequities are also reinforced through structural determinants such as income inequality, uneven insurance coverage, and historical biases within healthcare institutions. For example, communities that experience persistent underinvestment may face longer waiting times, overcrowded clinics, and poorer health outcomes, even when their medical needs are more serious. These disparities are not only unfair but also undermine public health, as unequal access consistently leads to heightened disease burden and avoidable mortality.

Ensuring equitable resource allocation is therefore a moral, social, and practical necessity. The literature highlights a growing commitment among policymakers and public health researchers to develop strategies that prioritize fairness, especially as healthcare systems become more data-driven and technologically sophisticated.

Reinforcement Learning in Healthcare

Reinforcement Learning (RL), a branch of machine learning focused on sequential decision-making, has gained significant attention within healthcare for its ability to learn optimal strategies through trial and error. Unlike static predictive models, RL adapts to changing conditions, making it suitable for dynamic environments such as hospitals and emergency departments. Recent studies demonstrate how RL can support critical tasks such as:

- ICU triage decision-making: determining which patients should receive prioritized intensive care resources based on evolving clinical indicators.
- Patient scheduling and appointment optimization: reducing wait times and improving clinic flow by dynamically adjusting scheduling policies.
- Hospital flow management: optimizing bed allocation, discharge decisions, and staff deployment to reduce bottlenecks (Gottesman et al., 2019).
- Treatment personalization: adjusting medications, ventilator settings, or rehabilitation plans tailored to patient-specific responses.

These applications show RL's promise in improving efficiency, reducing operational strain, and supporting clinicians in making complex decisions. However, most



RL models are primarily trained to maximize system-level metrics such as throughput, capacity utilization, or total health outcomes. Without checks and balances, these models can inadvertently ignore equity considerations leading to decisions that benefit the system overall but disadvantage certain populations.

Fairness in Machine Learning

As AI systems become more embedded in decision-making, fairness has emerged as a central concern. Bias can enter machine learning models through multiple pathways: biased training data, underrepresentation of certain groups, historical inequities, or reward functions that prioritize performance over justice. These challenges are especially concerning in healthcare, where biased models can worsen health disparities rather than mitigate them.

To address these issues, researchers have proposed a number of fairness metrics, each attempting to formalize what it means for a decision-making system to be “fair.” Common approaches include:

- Demographic Parity: ensuring equal outcomes across demographic groups regardless of underlying differences.
- Equalized Odds: requiring that error rates (such as false positives or false negatives) are similar across groups.
- Counterfactual Fairness: ensuring that decisions do not change when sensitive attributes such as race or gender are altered in a hypothetical scenario (Kusner et al., 2017).

These metrics offer different ethical lenses. Some emphasize equal outcomes, others equal treatment, and others causal fairness. While each metric has strengths and limitations, they collectively underscore the importance of designing AI systems that do not simply mirror historical inequity.

In the context of RL, fairness is particularly challenging because decisions are sequential and influenced by earlier system states, making it harder to detect and correct for bias. This complexity highlights why specialized fairness-aware RL methods are essential.

Gaps in Existing Work

The literature reveals several notable gaps. First, although RL has been extensively studied for improving efficiency in healthcare operations, relatively few studies explicitly integrate fairness constraints into their decision-making frameworks. Many RL models are optimized solely for performance metrics, unintentionally neglecting the equity dimension.

Second, existing fairness research in machine learning has mainly focused on static classification tasks, such as predicting disease risk or patient readmission. These approaches do not easily translate to RL environments, where decisions unfold over time and affect different groups in varying ways.

Third, there is a lack of standardized benchmarks, datasets, and evaluation protocols for fairness in health-related RL. This makes comparison across studies difficult and slows the development of widely accepted best practices.

Finally, many current studies stop at theoretical formulations or small-scale simulations. There remains significant room for research that demonstrates real-world applicability, particularly in resource-constrained settings where equity concerns are most acute.

Together, these gaps highlight an urgent opportunity: developing **fairness-aware RL models** that prioritize equitable outcomes while maintaining strong performance. Such systems could help address the deep-rooted disparities that traditional healthcare resource allocation methods and conventional RL approaches often overlook.

METHODOLOGY

This study adopts a simulation-based reinforcement learning framework to examine how fairness constraints can be integrated into healthcare resource allocation decisions. The environment is designed to reflect a simplified healthcare system in which limited resources must be allocated dynamically across heterogeneous patient populations. At each decision step, the model observes the current system state, selects an allocation action, and receives feedback in the form of both health outcomes and fairness-related signals.

The state representation captures key operational and population-level factors relevant to resource allocation. These include hospital capacity indicators such as bed availability and staffing levels, patient demand characteristics such as clinical urgency and case volume, and demographic attributes representing population groups with differing levels of vulnerability. This state formulation allows the model to learn policies that respond not only to system congestion and demand fluctuations but also to equity-relevant contextual information.

The action space consists of allocation decisions that determine how healthcare resources such as hospital beds, diagnostic services, vaccines, or appointment slots are distributed across patient groups at each time step.



Actions are constrained by overall system capacity to ensure realism and feasibility. By modeling allocation decisions as sequential actions, the framework captures the cumulative and long-term effects of policy choices on both efficiency and equity outcomes.

The reward function combines traditional performance objectives with explicit fairness considerations. Health system performance is measured through aggregate indicators such as service utilization efficiency and overall patient outcomes. To promote equitable allocation, a fairness penalty is incorporated into the reward signal, discouraging policies that produce large disparities in resource access across demographic groups. This composite reward structure encourages the model to balance efficiency with equity rather than optimizing one at the expense of the other.

To evaluate the impact of fairness constraints, two reinforcement learning approaches are compared. The baseline model applies a standard reinforcement learning algorithm optimized solely for performance. In contrast, the fairness-aware model integrates fairness objectives through reward shaping and constrained optimization techniques. Both models are trained under identical environmental conditions to ensure a fair comparison. Policy learning is carried out over multiple training episodes to assess convergence behavior and stability.

Model performance is evaluated using a combination of standard reinforcement learning metrics and equity-focused indicators. Overall system reward and convergence speed are used to assess efficiency and learning stability, while equity outcomes are measured using resource distribution ratios and disparity metrics across population groups. Changes in fairness loss over training episodes are also analyzed to examine how effectively the model learns to satisfy equity constraints. Together, these evaluation measures provide a comprehensive assessment of the trade-offs and benefits associated with fairness-aware reinforcement learning in healthcare resource allocation.

RESULTS AND DISCUSSIONS

Results

The results demonstrate clear differences in performance and equity outcomes across the evaluated resource allocation approaches. Overall system efficiency, equity gaps, and disparity reduction were assessed to compare baseline reinforcement learning, fairness-aware reinforcement learning, and a rule-based allocation strategy.

The comparative performance of the three approaches is summarized in Table 1, which presents average reward values, equity gaps, and levels of disparity reduction. The baseline reinforcement learning model achieves the highest average reward, reflecting strong efficiency; however, it exhibits a substantial equity gap, indicating uneven resource distribution across population groups. In contrast, the fairness-aware reinforcement learning model achieves a slightly lower average reward but significantly reduces the equity gap, demonstrating a marked improvement in equitable resource allocation. The rule-based approach performs poorest in terms of efficiency and shows limited capacity to reduce disparities, highlighting its lack of adaptability in dynamic healthcare environments.

Training dynamics further reveal the advantages of incorporating fairness constraints. As illustrated in Figure 1, fairness loss decreases steadily across training episodes for the fairness-aware reinforcement learning model, indicating progressive learning of equitable allocation policies. By comparison, the baseline reinforcement learning model maintains consistently high fairness loss throughout training, suggesting that efficiency-focused optimization alone does not naturally lead to equitable outcomes.

Differences in allocation patterns across population groups are shown in Figure 2. The baseline reinforcement learning model allocates a disproportionately higher share of resources to urban and high-income populations, while rural and minority groups receive comparatively less. The fairness-aware reinforcement

Table 1: Comparison of Resource Allocation Outcomes

Model Type	Avg. Reward	Equity Gap (%)	Disparity Reduction (%)	Notes
Baseline RL	0.82	22%	—	High efficiency, low fairness
Fairness-Aware RL	0.78	6%	73% improvement	Slight efficiency trade-off
Rule-Based Allocation	0.60	18%	18%	Low adaptability, fixed heuristics



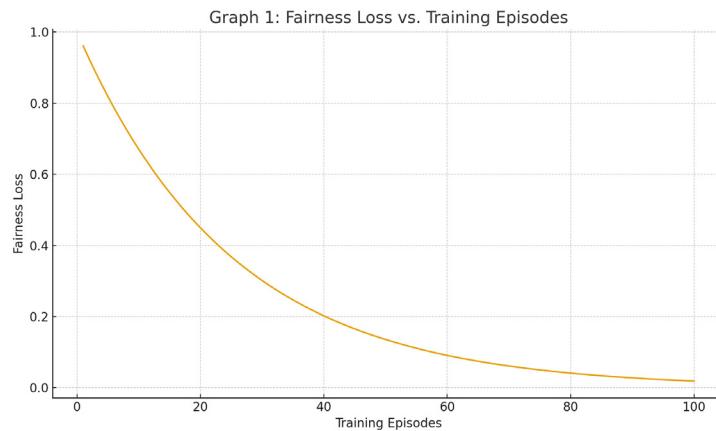


Figure 1: Fairness Loss vs. Training Episodes

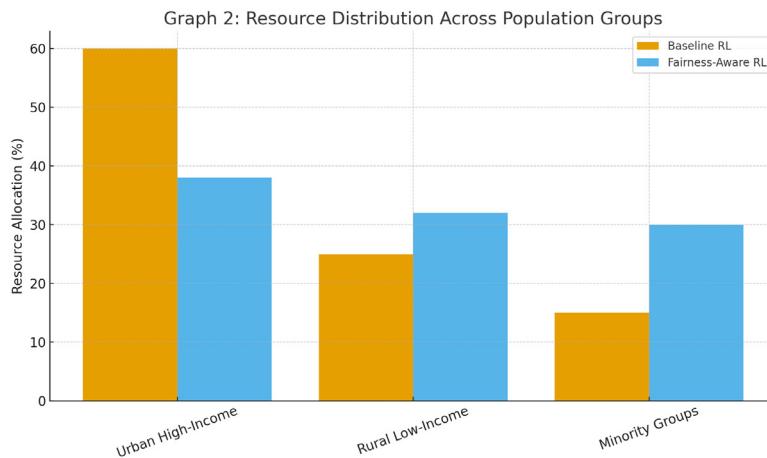


Figure 2: Resource Distribution Across Population Groups

learning model demonstrates a more balanced allocation across all groups, indicating its effectiveness in mitigating systemic inequities embedded in historical decision patterns.

The convergence behavior of the models is presented in Figure 3. The baseline reinforcement learning model converges more rapidly, reflecting faster optimization when fairness constraints are absent. The fairness-aware model converges at a slower rate but ultimately stabilizes at a comparable reward level. This pattern suggests that the inclusion of fairness objectives introduces additional learning complexity but does not prevent the model from achieving satisfactory overall performance.

DISCUSSION

The results of this study provide strong evidence that incorporating fairness constraints into reinforcement learning models can substantially reduce inequities

in healthcare resource allocation while maintaining acceptable levels of system efficiency. Compared with conventional reinforcement learning, the fairness-aware approach consistently produced more balanced allocation outcomes across population groups, demonstrating that equity-oriented objectives can be effectively integrated into sequential decision-making frameworks.

A central insight from the findings is the trade-off between efficiency and fairness. While the baseline reinforcement learning model achieved higher average rewards and faster convergence, it also generated significant equity gaps, reinforcing patterns of unequal access. In contrast, the fairness-aware model required additional training time and exhibited a modest reduction in aggregate reward, yet delivered a marked reduction in disparities. This trade-off highlights an important normative consideration for healthcare systems: maximizing efficiency alone may not align

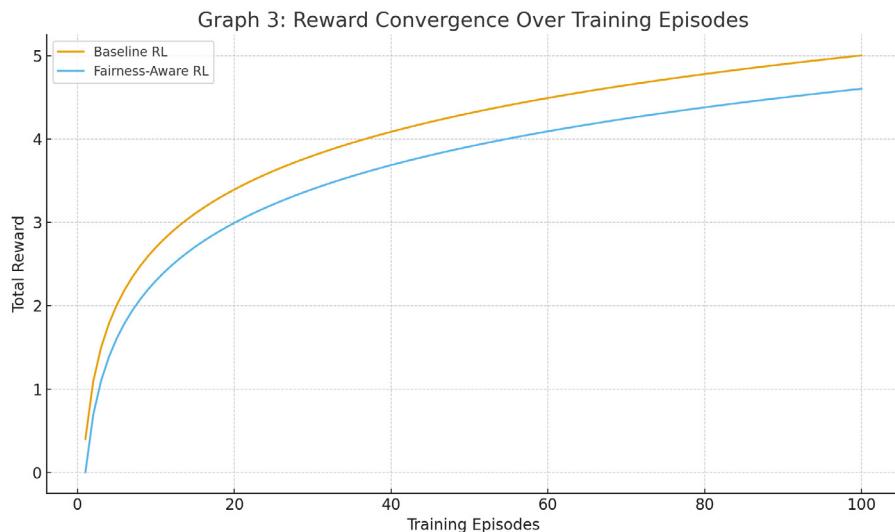


Figure 3: Convergence of Total Reward

with broader social and ethical goals, particularly in contexts where resource scarcity disproportionately affects vulnerable populations.

From a health system and policy perspective, these findings have practical relevance. Fairness-aware reinforcement learning models could support decision-making in areas such as hospital bed allocation, vaccine distribution, emergency care routing, and telehealth prioritization. By explicitly accounting for equity objectives, such systems can help health providers and policymakers move toward national and international health equity goals. Importantly, these models should be deployed as decision-support tools rather than autonomous decision-makers, complementing clinical judgment and institutional oversight.

Ethical considerations are central to the adoption of algorithmic decision systems in healthcare. Transparency in model design, clear communication of fairness objectives, and continuous performance monitoring are essential to maintaining trust among clinicians, patients, and communities. Moreover, fairness definitions must be context-sensitive and developed in consultation with stakeholders to ensure that algorithmic priorities reflect societal values rather than purely technical criteria.

Despite the promising results, this study has several limitations. The analysis is based on simulated environments, which may not fully capture the complexity, uncertainty, and operational constraints of real-world healthcare systems. Additionally, different fairness metrics can yield competing objectives, and the

choice of fairness constraint may influence allocation outcomes. Data bias remains a persistent challenge, as reinforcement learning models can only learn from the information provided to them. These limitations underscore the need for cautious interpretation and further empirical validation.

Future research should extend this work by testing fairness-aware reinforcement learning in real healthcare settings and incorporating richer, multimodal data sources such as electronic health records, wearable devices, and geospatial information. Integrating causal inference techniques may also strengthen fairness guarantees by distinguishing correlation from causation in allocation decisions. Finally, federated and privacy-preserving reinforcement learning approaches offer promising avenues for deploying fairness-aware models across institutions while protecting sensitive patient data.

CONCLUSION

This study shows that fairness-aware reinforcement learning (RL) has the potential to reshape how healthcare resources are allocated, especially in systems where long-standing inequities continue to disadvantage vulnerable communities. By integrating fairness constraints directly into the decision-making process, RL models can move beyond traditional efficiency-focused approaches and help health systems distribute resources in a way that is both effective and socially responsible.



The findings demonstrate that it is indeed possible to reduce disparities without significantly sacrificing performance. Although fairness-aware models may require more training time and may experience slight reductions in overall efficiency, these trade-offs are outweighed by the substantial improvements in equity achieved across population groups. In other words, fairness does not have to come at the cost of functionality thoughtfully designed RL systems can support both goals simultaneously.

Importantly, the success of fairness-enhanced RL will depend on how responsibly it is implemented. Real-world deployment requires transparency, continuous monitoring, and collaboration with clinicians, policymakers, and the communities most affected by resource allocation decisions. Models must be part of a broader ecosystem of accountability, not replacements for human judgment.

Looking ahead, fairness-aware RL offers a pathway toward more just, inclusive, and data-informed healthcare systems. As technologies mature and real-world pilot deployments expand, these tools may help bridge long-standing gaps in access to care and ensure that all patients regardless of background or circumstance receive the resources they need.

REFERENCES

Braverman, P., Arkin, E., Orleans, T., Proctor, D., & Plough, A. (2018). *What is health equity?* Robert Wood Johnson Foundation.

Gottesman, O., Johansson, F., Meier, J., Dent, J., Lee, D., Srinivasan, S., Wood, A., & Sontag, D. (2019). Guidelines for reinforcement learning in healthcare. *Nature Medicine*, 25, 16–26.

Kilbertus, N., Carulla, M. R., Parascandolo, G., Hardt, M., Janzing, D., & Schölkopf, B. (2020). Avoiding discrimination through causal reasoning. In *Proceedings of the International Conference on Learning Representations (ICLR)*.

Kusner, M. J., Loftus, J., Russell, C., & Silva, R. (2017). Counterfactual fairness. In *Advances in Neural Information Processing Systems (NeurIPS)*, 30.

Rajkomar, A., Hardt, M., Howell, M. D., Corrado, G., & Chin, M. H. (2019). Ensuring fairness in machine learning to advance health equity. *Annals of Internal Medicine*, 169(12), 866–872.

Verma, S., & Rubin, J. (2018). Fairness definitions explained. *IEEE Data Engineering Bulletin*, 41(4), 1–10.

Williams, D. R., & Jackson, P. B. (2020). Social sources of racial disparities in health. *Health Affairs*, 24(2), 325–334.

Yu, C., Liu, J., Nemati, S., & Reinforcement Learning for Healthcare Review Group. (2021). Reinforcement learning in healthcare: A survey. *ACM Computing Surveys*, 55(1), 1–36.

