

Algorithmic Fairness in Temporal Resource Allocation

Doctoral Consortium

Ashwin Kumar

Washington University in St. Louis

Saint Louis, MO, USA

ashwinkumar@wustl.edu

ABSTRACT

There has been a significant body of research on improving social welfare in resource allocation, but much of it has focused on single-shot allocation scenarios, where a given pool of resources must be divided equitably. In contrast, my research aims to address the unique challenges posed by temporal resource allocation problems that involve many repeated allocations, with both resources and beneficiaries able to re-enter the market at different points in time. Automated algorithms are often employed to guide resource allocation in these scenarios by estimating and comparing utilities of different allocations, making algorithmic fairness a concern as well. In this work, I aim to improve long-term social welfare in addition to maximizing the utility of such systems through the lens of pre-, in-, and post-processing fairness. I propose a simple incentive-based approach for post-processing fairness with black-box value functions, outperforming existing baselines in a ridesharing application. I discuss two other research thrusts using fairness-aware dataset balancing for pre-processing fairness and learning non-myopic fairness policies for in-processing fairness. Combining all of these approaches, my goal is to present a holistic view of improving social welfare in temporal resource allocation through the lens of algorithmic fairness.

KEYWORDS

Fairness; Resource Allocation; Fair Allocation; Reinforcement Learning

ACM Reference Format:

Ashwin Kumar. 2023. Algorithmic Fairness in Temporal Resource Allocation: Doctoral Consortium. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 3 pages.

1 INTRODUCTION

The design of fair resource allocation schemes is central to the acceptance of algorithmic decision making for solving societal problems. In typical resource allocation settings, the goal is to divide a pool of resources amongst agents to maximize utility as well as social welfare. There exist many notions of this welfare objective, like max-min fairness and proportional fairness [7]. Many real-world scenarios (e.g., ridesharing, refugee resource allocation, homelessness prevention), take the form of a repeated resource allocation problem, where resources arrive over time windows, and decisions have to be made with limited or no information about the future. Resources, as well as agents, may re-enter or stay in

the market for multiple time steps. In this work, we term these as Temporal Resource Allocation Problems (**TRAPs**). In order to make non-myopic decisions, TRAPs often employ Value Function Approximators (**VFAs**) to predict the expected future utility of certain matchings. When machine learning techniques are used to estimate these values, they introduce the possibility of algorithmic bias.

Mitigation techniques for algorithmic fairness usually fall into one of three categories [5], (1)**pre-processing**: modifying the input of the learning algorithm to make it fairer; (2)**in-processing**: Adding fairness constraints to the objective or modifying the learned VFA to be fair e.g., by moving the decision boundary; (3)**post-processing**: treating the VFA as a black box and modifying the outputs to enforce fairness. In my research, I aim to explore each of these categories in search of solutions for improving social welfare and overall utility in TRAPs by improving the fairness of VFAs.

2 RESEARCH DIRECTIONS

2.1 Post-Processing Fairness: Simple Incentives

Post-processing fairness is of use when a complicated but effective VFA leads to disparate division of resources. In my initial approach, I looked at the use of simple group-based incentives to skew the utilities predicted by the VFA in favor of disadvantaged groups, thus moving the allocation for a time window towards a fairer solution. As an application area for this approach, I looked at ridesharing matching in urban environments. In ridesharing matching, passengers need to be matched to available taxis to minimize their waiting time, for the objective of maximizing the fraction of people served (i.e., the service rate) by a fixed fleet of taxis. The optimal matching for each time window is found by optimizing the total expected values of passenger assignments (from a VFA). However, when looking at a real world dataset, we found state-of-the-art VFAs to be partial in its service to regions with low demand [3, 8]. Specifically, when grouping passengers by origin-destination region pairs, we found a huge disparity in the service rates. Drawing from the notion of Statistical Parity [1], we designed a simple additive term to serve as a gradient step in the space of all matches, moving towards a fairer allocation based on historical service rates.

Concretely, given a set of vehicles \mathcal{V} and a set of passenger requests \mathcal{R} , we generate a set of feasible actions A_i for vehicle i , where $a_i^k \in A_i$ is a subset of requests. Let $U(i, a)$ denote the predicted utility of assigning vehicle i action a , the objective is to maximize this utility across all vehicles subject to capacity constraints.

Let z_j denote the historical service rate of group j , and \mathbf{Z} denote the set of group service rates. To move towards statistical parity, we minimize the variance in these service rates (a variance of zero will satisfy the fairness requirements). Specifically, we compute the gradient of the variance $\text{var}(\mathbf{Z})$ with respect to the assignment

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

\mathcal{A} , $\frac{\partial}{\partial \mathcal{A}} \text{var}(\mathbf{Z})$. If we assume that the average of the metric over all groups is stable (i.e., $\frac{\partial}{\partial \mathcal{A}} \bar{z} \approx 0$, a reasonable assumption if a long enough history is included), then we can find an assignment for a modified utility function that accounts for the gradient of the variance with respect to the assignment \mathcal{A} :

$$\begin{aligned}
 U'(i, a) &= U(i, a) - \lambda \frac{\partial}{\partial \mathcal{A}} \text{var}(\mathbf{Z}) \\
 &= U(i, a) - \frac{1}{|\mathbf{Z}|} \lambda \frac{\partial}{\partial \mathcal{A}} \sum_{z_j \in \mathbf{Z}} (z_j - \bar{z})^2 \\
 &= U(i, a) + \beta \sum_{z_j \in \mathbf{Z}} (\bar{z} - z_j) \frac{\partial z_j}{\partial \mathcal{A}} \tag{1}
 \end{aligned}$$

The constant β serves as a trade-off parameter for the value of fairness. If we assume the effects of each action to be independent in the assignment, Eq.1 can be simplified by taking the derivative with respect to the action a instead of the assignment \mathcal{A} . Intuitively, this method adds an incentive for groups which have a below-average metric value.

Doing this for a state-of-the-art VFA, we find we can significantly improve passenger-side fairness, with minimal loss in the overall efficiency. We are also able to show that for a high enough β , this method guarantees that the worst-off passenger group sees an improvement at each time step [4].

We also applied the same idea towards minimizing the income disparity in drivers, where the incentive now depends on their relative historical income level. Experiments show this method also gives better fairness-efficiency trade-offs than existing methods. Combining both passenger and driver side fairness, we were able to get pareto-dominating solutions with highly competitive efficiency.

We also found application for this approach in allocation of homelessness resources for homelessness prevention. In this problem, we need to allocate homeless people to different interventions, each with different probabilities of preventing re-entry into homelessness. This probability is calculated using a counterfactual BART model [2], and people enter the system at different points of time, thus making it a TRAP. With the goal of equalizing post-decision re-entry probabilities of different demographics (groups), we applied the simple incentives approach to modify the predicted probabilities. Initial results suggest favorable effects of this approach.

2.2 In-processing Fairness: Non-Myopic Fairness

The incentives approach in the previous section forms a myopic method of enforcing fairness, where we modify utilities to elicit a fairer allocation. This leads directly to the next thrust of research: Is there a way to select actions that improve long-term fairness?

Existing work [6] tries to add variance as a cost to the optimization objective in the ridesharing context. This is not stable, as the variance term changes scale over time. This also prevents us from creating useful hyperparameters to specify the tradeoff between fairness and efficiency. Further, in our experiments, we observed that our myopic post-processing solution outperformed this method.

In initial experiments, training with the myopic incentive score added to the reward function also yielded minimal improvements

to the learned VFA’s fairness. Since the incentive function is not additive across multiple time steps, optimizing it as a discounted reward does not elicit a true representation of the overall fairness objective.

My goal in this project is to identify and develop a scalable metric for fairness that leads to the minimization of fairness over time, and that can be decomposed over subsequent time steps. We will learn this metric in parallel with the VFA, since the long-term fairness is dependent on the environment dynamics as dictated by the VFA. Concretely, this will allow us to combine the long-term utility and fairness of an action, and allow us to tune the trade-off weight as a bonus. It is not clear whether it is possible to create an additive fairness utility that is optimal, so part of my research will involve figuring out existence conditions. If such a metric does not exist, I plan to develop heuristic solutions with competitive approximation ratios.

2.3 Pre-processing Fairness: Bias Bootstrapping

While modifying the assignment to make it fair serves as a post-hoc solution, it may be possible to modify the training data for the VFA so that it leads to the prediction of fair utilities. In the third frontier for my research, I look at methods for value estimation that are trained in an online fashion, akin to the temporal nature of TRAPs. Since training data for such methods is often generated by playing out trajectories that (mostly) follow the policy of the current VFA, the data distribution will reflect its biases, leading to a feedback loop, which we term Bias Bootstrapping.

As an example, active learning is often used to identify data points to query for the ground truth, when such queries are expensive, and then trained using the updated dataset. If the current predictor is strongly biased against group A (all false negatives, say), using active sampling for selecting individuals to “query” and label manually can lead to a dataset without new members of group A, which leads to the next classifier being biased as well. My preliminary experiments show that these feedback loops exist in paradigms like Active Learning and Reinforcement Learning.

In this project, I plan to design modifications to experience replay and active sampling that mitigate this issue of Bias Bootstrapping. Early experiments suggest that increasing population diversity improves the convergence, but further investigation is needed to pin-point the loss in terms of the learning value associated with such methods.

3 CONCLUSION

The intersection of temporal resource allocation and algorithmic fairness holds the potential for many diverse approaches to improve equity. I propose evaluating and improving the social welfare in temporal resource allocation using ideas from algorithmic fairness, giving a unique insight into the interplay between the two. Existing results show that this is a problem of interest and establish a bedrock for solution approaches. Through the three-pronged search for fairness interventions, this research will provide a holistic overview for fairness in TRAPs.

ACKNOWLEDGMENTS

This research is partially supported by the National Science Foundation under awards 1812619, 1905558, 1939677, 2020289, 2118201, and 2214141; by ARO grant W911NF1810208; and by Amazon. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the sponsoring organizations, agencies, or the United States government.

REFERENCES

- [1] Michael Feldman, Sorelle A. Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. 2015. Certifying and Removing Disparate Impact (*KDD '15*). Association for Computing Machinery, 259–268. <https://doi.org/10.1145/2783258.2783311>
- [2] Amanda Kube, Sanmay Das, and Patrick J Fowler. 2019. Allocating interventions based on predicted outcomes: A case study on homelessness services. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 622–629.
- [3] Ashwin Kumar, Sanket Shah, Meghna Lowalekar, Pradeep Varakantham, Alvitta Ottley, and William Yeoh. 2021. FairVizARD: A Visualization System for Assessing Fairness of Ride-Sharing Matching Algorithms. ICAPS System Demo.
- [4] Ashwin Kumar, Yevgeniy Vorobeychik, and William Yeoh. 2023. Using Simple Incentives to Improve Two-Sided Fairness in Ridesharing Systems. *Under Review*.
- [5] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)* 54, 6 (2021), 1–35.
- [6] Naveen Raman, Sanket Shah, and John Dickerson. 2021. Data-Driven Methods for Balancing Fairness and Efficiency in Ride-Pooling. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*.
- [7] Brianna Richardson and Juan E Gilbert. 2021. A Framework for Fairness: A Systematic Review of Existing Fair AI Solutions. *arXiv preprint arXiv:2112.05700* (2021).
- [8] Sanket Shah, Meghna Lowalekar, and Pradeep Varakantham. 2020. Neural Approximate Dynamic Programming for On-Demand Ride-Pooling. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 507–515.