Characterizing Fairness in Societal Resource Allocation

Doctoral Consortium

Tasfia Mashiat
George Mason University
Fairfax, VA, USA
tmashiat@gmu.edu

ABSTRACT

Societal biases can lead to disparate impacts on, and treatment of, different demographic groups. This can have substantial effects on the outcomes of public resource allocation scenarios like child welfare, housing allocations for homeless persons, etc. In recent years, there has been an increasing interest in devising algorithms for allocating public resources. However, ensuring the fairness and equitability of these algorithms is challenging since the definition of fairness is highly intersectional, multi-modal, and domain-specific. Moreover, the allocation of these resources is dynamic and time-dependent in nature. While there exist several notions of fairness in the Machine Learning (ML) literature, their applicability to Fair Division (FD) of resources is limited. In our research, we aim to bridge the gap between the Fair ML and economic theories of FD for public resource allocation. More specifically, we will study different application areas such as policing, homelessness, and eviction, devise fair algorithms and metrics for these resources, and evaluate their effectiveness in public policymaking.

KEYWORDS

Fairness in Machine Learning, Societal Resource Allocation, AI for Social Good

ACM Reference Format:
Tasfia Mashiat. 2023. Characterizing Fairness in Societal 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 BACKGROUND AND MOTIVATION

Public systems that manage the social problem of resource allocation often have to make complex decisions under great uncertainty. They have to make decisions about resource assignment to individuals while maintaining the underlying capacity constraints, which often become difficult due to poor assessment of needs and inefficient understanding of the compatibility between services and individuals [1, 2]. For example, severe discrimination has been observed in child welfare services, which in turn can create permanent turmoil in a child's life [3].

The notion of fairness in fair division literature can be broadly defined in three separate ways: *Envy-freeness, Proportionality, and Equitability* [4]. These metrics model the problem of assigning a set of resources ($K = \{1, 2, ..., k\}$) to a set of agents ($N = \{1, 2, ..., n\}$) where each agent has its preferences for each of the resources based on a valuation function v. *Envy-freeness* ensures that no participating agent wishes to swap their resources with another agent, thus allocation A_i for agent i will be envy-free if $\forall i, j \in N : v_i(A_i) \geq v_i(A_j)$. *Proportionality* refers that in a system with n agents, each agent will receive $1/n$ th of the agent's total utility for all of the resources. Thus an allocation A_i is proportional for agent i if $\forall i \in N : v_i(A_i) \geq 1/n$ [5]. Finally, *Equitable* allocation implies that each agent will receive equal utilities [6], an allocation A_i is equitable for agent i if $\forall i, j \in N : v_i(A_i) = v_j(A_j)$.

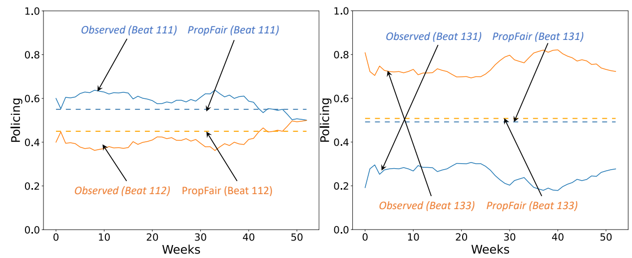

(a) (111,112) **(b) (131,133)**

Figure 1: Observed and predicted assignment of police officers as estimated by our method. Police Beats (111,112) and (131,133) are neighbouring beats of District1. Y axis shows the proportion of policing in the pair of beats.

While these metrics ensure optimal allocation based on utility function, they do not measure biases against different demographic groups. However, it is imperative to ensure equitable distributions of scarce resources among different demographic groups. Prior research has shown that algorithmic methods can promote disparate impacts [7, 8]. Several recent works have focused on these concerns and proposed algorithmic fairness notions such as demographic parity [9], equalized odds [10], equal opportunity [10], and counterfactual fairness [11]. However, most of these notions are defined for static environments [12], and do not account for allocation systems that are dynamic in nature where the supply and demand of resources continuously shift over time.

Motivated by the lack of fairness notions in this domain, we are working on three different societal resource allocation problems, Homelessness, Policing, and Eviction. In these three domains, the decision algorithm has to allocate resources among different demographic groups. We combine the concepts of FD and ML literature to apply in the allocation of societal resources, more specifically show how these metrics result in different allocation decisions.

2 CURRENT WORK

Fairness in Predictive Policing: Predictive Policing algorithms assign police officers across a city based on historical crime data. Past

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.

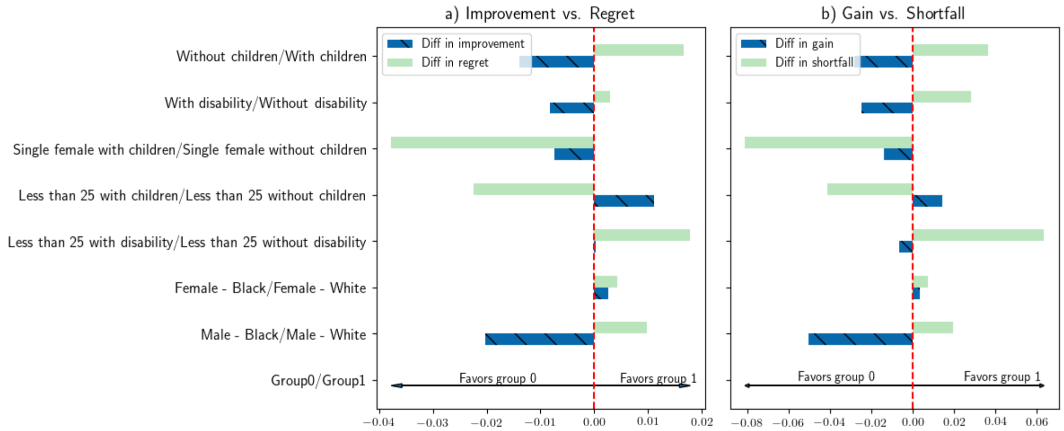

Figure 2: Fairness trade-off in the observed assignment of homeless services. This compares which demographic group is favored by the assignment depending on the fairness metric. Trade-offs occur when improvement favors one group and regret the other one (left panel) or when shortfall favors one group and gain the other (right panel).

works show that these algorithms can suffer from multiple drawbacks, such as feedback loop (past decisions affecting subsequent decisions), and unfairness (in terms of protected attributes) [8, 13].

This project presents a solution for both of these issues by incorporating causal inference in the algorithm. We define a causal framework to allocate police officers in a fixed geographic area of a city called "Police Beat". We demonstrate that this causality-driven framework was able to allocate police officers optimally across beats within a policing district based on the true underlying criminality. Moreover, we apply this algorithm (*PropFair*) in a novel dataset obtained from the City of Chicago (through Freedom of Information Act requests) and show that there is significant variability in policing allocation across different neighboring beats for underlying criminality (Figure 1). The work is accepted as an extended abstract at AAMAS 2023 [14].

Fairness in Homeless Service Delivery: According to federal guidelines, homelessness can be defined as the absence of stable and permanent residences which include shelters, inhabitable places such as cars, parks, and various public access places. Since 2007, more than 550,000 people experience homelessness on regular basis, while 1.5 million people seek homeless delivery services each year [15]. Though each year congress provides funds to the homeless delivery services, the funds, and efforts are insufficient to handle the undergoing homelessness crisis in the United States [16, 17]. Moreover, the allocation of these services has a significantly different impact on utilities for people of different demographic groups. In this work, we show that one group performing comparatively better than the other for one metric can do worse for a different metric (Figure 2). Even the performance can also shift if the metrics are converted from additive to multiplicative for gained utilities. We defined two metrics, Improvement, and Regret, by adapting the idea of equitable allocation of fair division literature. We defined Improvement as the metric of how well an agent is doing for the best-suited option for an allocation and Regret as the metric of how well an agent is doing in an allocation based on the worst possible allocation, similar definitions are presented for multiplicative measures (Gain and Shortfall). We also applied these metrics to a novel dataset collected from the homeless management information system (HMIS) of a metropolitan area from 2007 through 2014. We demonstrate that one group can appear privileged based on Improvement over the other group, whereas it may appear less privileged by the allocation than the other group on the Regret metric for an assignment policy. The work is published at FAccT 2022 [18].

3 FUTURE PLAN AND CONCLUSION

In the future, we would like to pursue the following research directions:

Robust fairness in dynamic police allocation?

Our current causal model can handle both feedback loops and unfairness in terms of demographics. However, we are continuing our project to understand whether there are any other hidden confounders that can impact the allocation policy. We are also performing extensive experiments for bench-marking our method with existing policing algorithms such as predictive policing and polya urn model [13].

How to specify appropriate fairness notion for Eviction?

Thousands of US residents get evicted from their houses each year. These decisions are typically bi-level meaning, they come from the landlords and the courts where they can disagree on various grounds. This makes it a complex decision-making process, where social biases can creep in. To this end, we obtained a real world eviction dataset and currently working on modeling the process and characterizing fairness.

Finally, we are extending our work on homelessness. Though our work [18] shows the trade-offs in fairness metrics, we are working on how these metrics can be used to formulate a fairness guideline.

In summary, we hope to continue working on novel societal resource allocation problems. We believe our research will contribute to characterizing fairness in these domains, generate new computational models, and inform future research for how to ensure fairness in critical societal resource allocation.

REFERENCES

[1] D. Gubits, M. Shinn, M. Wood, S. R. Brown, S. R. Dastrup, and S. H. Bell, "What interventions work best for families who experience homelessness? impact estimates from the family options study," *Journal of Policy Analysis and Management*, vol. 37, no. 4, pp. 835–866, 2018.

[2] M. Shinn, S. R. Brown, B. E. Spellman, M. Wood, D. Gubits, and J. Khadduri, "Mismatch between homeless families and the homelessness service system," *Cityscape (Washington, DC)*, vol. 19, no. 3, p. 293, 2017.

[3] A. Chouldechova, D. Benavides-Prado, O. Fialko, and R. Vaithianathan, "A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions," in *Conference on Fairness, Accountability and Transparency*. PMLR, 2018, pp. 134–148.

[4] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, and M. Kyropoulou, "The efficiency of fair division," *Theory of Computing Systems*, vol. 50, no. 4, pp. 589–610, 2012.

[5] J. Robertson and W. Webb, *Cake-cutting algorithms: Be fair if you can*. CRC Press, 1998.

[6] K. Cechlárová and E. Pillárová, "On the computability of equitable divisions," *Discrete Optimization*, vol. 9, no. 4, pp. 249–257, 2012.

[7] T. Bolukbasi, K.-W. Chang, J. Y. Zou, V. Saligrama, and A. T. Kalai, "Man is to computer programmer as woman is to homemaker? debiasing word embeddings," *Advances in neural information processing systems*, vol. 29, pp. 4349–4357, 2016.

[8] K. Lum and W. Isaac, "To predict and serve?" *Significance*, vol. 13, no. 5, pp. 14–19, 2016.

[9] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel, "Fairness through awareness," in *Proceedings of the 3rd innovations in Innovations in Theoretical Computer Science Conference*, 2012, pp. 214–226.

[10] M. Hardt, E. Price, and N. Srebro, "Equality of opportunity in supervised learning," *Advances in neural information processing systems*, vol. 29, pp. 3315–3323, 2016.

[11] M. J. Kusner, J. R. Loftus, C. Russell, and R. Silva, "Counterfactual fairness," *arXiv preprint arXiv:1703.06856*, 2017.

[12] L. T. Liu, S. Dean, E. Rolf, M. Simchowitz, and M. Hardt, "Delayed impact of fair machine learning," in *International Conference on Machine Learning*. PMLR, 2018, pp. 3150–3158.

[13] D. Ensign, S. A. Friedler, S. Neville, C. Scheidegger, and S. Venkatasubramanian, "Runaway feedback loops in predictive policing," in *Conference on Fairness, Accountability and Transparency*, 2018, pp. 160–171.

[14] T. Mashiat, X. Gitiaux, H. Rangwala, P. Fowler, and S. Das, "Counterfactually fair dynamic assignment: A case study on policing," in *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems*, 2023.

[15] M. Henry, R. Watt, L. Rosenthal, A. Shivji *et al.*, "The 2018 annual homeless assessment report (ahar) to congress. part 1: Point-in-time estimates of homelessness," *The US Department of Housing and Urban Development*, 2018.

[16] R. Vaithianathan, E. Putnam-Hornstein, N. Jiang, P. Nand, and T. Maloney, "Developing predictive models to support child maltreatment hotline screening decisions: Allegheny county methodology and implementation," *Center for Social Data Analytics*, 2017.

[17] A. R. Kube, S. Das, and P. J. Fowler, "Fair and efficient allocation of scarce resources based on predicted outcomes: Implications for homeless service delivery," 2019.

[18] T. Mashiat, X. Gitiaux, H. Rangwala, P. Fowler, and S. Das, "Trade-offs between group fairness metrics in societal resource allocation," in *2022 ACM Conference on Fairness, Accountability, and Transparency*, 2022, pp. 1095–1105.