Fair Allocation Problems in Reviewer Assignment

Doctoral Consortium

Justin Payan
University of Massachusetts Amherst
Amherst, MA, USA
jpayan@umass.edu

ABSTRACT

Peer review is an incredibly important and revered process for establishing scientific rigor and facilitating effective discourse. However, many conventions of this process are ad-hoc and could benefit heavily from a more intentional design. I present our existing work applying fair allocation techniques to a major problem in peer review, the assignment of reviewers. We propose a simple but flexible algorithm for fair reviewer assignment, based on the classic Round Robin picking sequence. In addition, I discuss two other areas in peer review that are ripe for fair allocation and mechanism design, reviewer bidding and two-sided fair reviewer assignment.

KEYWORDS

Reviewer Assignment; Fair Allocation; Fair Two-Sided Matching; Preference Elicitation

ACM Reference Format:

Justin Payan. 2022. Fair Allocation Problems in Reviewer Assignment: Doctoral Consortium. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), Online, May 9–13, 2022,* IFAAMAS, 3 pages.

1 INTRODUCTION

Peer review is a fundamental component of the research process, and researchers have been calling for improvements to the peer review system for years [7, 11]. Our work focuses on a central component of the peer review process, the assignment of reviewers to papers. We often see unfairness to authors of papers in less popular or interdisciplinary areas, unbalanced workloads for reviewers, and clustering of reviewer bids on papers with popular topics. Our work seeks to solve these problems, among others, using definitions and techniques from the field of fair allocation.

We first aim for fair and efficient assignments of reviewers to papers. Existing assignment algorithms (present in systems like EasyChair, CMT, and OpenReview) first compute affinities between reviewers and papers, using a mixture of reviewer bids, keyphrase matching, collaborative filtering, and textual analysis. These systems then maximize the overall affinity of the assignment, subject to some constraints [4]. However, maximizing overall affinity can result in unfair results for individual papers. Using the fairness notion of *envy-freeness* from the fair allocation literature, we seek reviewer assignments that are envy-free up to one reviewer (EF1). We apply a modification of the classic Round Robin allocation procedure to achieve EF1 allocations under the constraints of reviewer assignment. However, the welfare of Round Robin depends crucially on the selection order. We prove that greedily optimizing the selection order yields a $(1+\gamma^2)$-approximation to the best selection order for Round Robin, where γ represents the distance from submodularity of the welfare function over selection orders.

We also describe proposed work in two related problem areas, two-sided fairness for reviewer assignment and reviewer bidding.

1.1 Related Work

Two simultaneous papers by Kobren et al. [6] and Stelmakh et al. [14] address the unfairness of the affinity-maximizing solution for reviewer assignment (exemplified by the well-known TPMS [4]). Stelmakh et al. [14] maximize the egalitarian welfare (the welfare of the worst-off paper) for reviewer assignment, while Kobren et al. [6] establish a threshold for the minimum score for any paper. We might also consider the assignment from the reviewers' perspectives. Meir et al. [8] propose a system where reviewers are incentivized to bid a certain number of points, and papers with fewer bids are worth more points. These bids can be directly interpreted as valuations, unlike affinities which are less directly interpretable. We would ultimately like to guarantee fairness to reviewers and papers simultaneously. We might apply something like the algorithm of Patro et al. [9], which uses a Round Robin procedure to ensure maximin share (MMS) on one side and EF1 on the other side.

Finally, Stelmakh [13] and Shah [12] present a large number of applications of fair allocation and mechanism design in reviewer assignment and peer review generally; there is still room for innovation in the problems they address and in others yet to be studied.

2 FAIR REVIEWER ASSIGNMENT

This section describes the results in [10], which presents a method for fair and efficient assignment of reviewers to papers (also accepted to AAMAS 2022 as an extended abstract). Consider a set of n papers N and a set of m reviewers R. Papers have a required number of reviewers k, and each reviewer r can review at most u_r papers. We seek a partition of the reviewers $\mathcal{A} = (A_1, A_2, \ldots, A_n)$ such that the papers' requirements are met exactly ($|A_i| = k$ for all i) and reviewer loads are within the required range (for all r, $\sum_{i \in N} |A_i \cap \{r\}| \leq u_r$). We assume access to a set of affinity scores $v_i(r)$ for all $i \in N$ and $r \in R$, which are often computed in systems like OpenReview, CMT, or EasyChair as previously mentioned (see [4] for an example). A paper i values a set of reviewers S additively, that is, $v_i(S) = \sum_{r \in S} v_i(r)$.

In general, we will try to allocate reviewers to maximize the utilitarian social welfare (USW):

$$\mathrm{USW}(\mathcal{A}) = \sum_{i \in N} v_i(A_i). \tag{1}$$

Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9–13, 2022, Online.

However, simply maximizing welfare can lead to poor results for individual papers. Hence, we seek to maximize welfare subject to the fairness constraint of envy-freeness up to one item (EF1). An allocation is envy-free up to 1 item if for all i and j, $v_i(A_i) \geq v_i(A_j \setminus \{r\})$ for some $r \in A_j$.

The classic Round Robin picking sequence is well-known to be EF1 in the general allocation setting [3]. In our setting, we would fix an order of papers and allow the papers to pick their favorite remaining reviewer at each round. This picking sequence can fail to be EF1 in our context, since papers may not select the same reviewer more than once. However, we can simply check for EF1 violations across a subset of papers before allowing a selection to occur - we call this updated algorithm Reviewer Round Robin.

Although Round Robin is guaranteed to be EF1 for general fair allocation problems (and Reviewer Round Robin is EF1 for reviewer assignment), the welfare of the resulting allocation depends on the order in which papers choose. As it is NP-hard to find the optimal order for Round Robin [1, 2], we use a greedy approach to find an approximately optimal order. We maintain a partial order over a subset of papers, and at each step we add the remaining paper whose addition yields the maximum welfare when running Round Robin. We show that partial orders can be represented as sets of tuples of the form $(paper, position)$. Valid orders are thus equivalent to sets in the intersection of two partition matroids, where one enforces uniqueness of papers in the order and the other enforces uniqueness of positions. We define a notion of γ-weak submodularity for monotonically non-decreasing functions f, which requires that for all sets $X \subseteq Y$ and $e \notin Y$, $\gamma(f(X \cup \{e\}) - f(X)) \geq f(Y \cup \{e\}) - f(Y)$. The main result of [10] states that if a monotonically non-decreasing transformation of the welfare function over the set of $(paper, position)$ tuples is γ-weakly submodular, then the welfare of the greedy Round Robin order is a $(1 + \gamma^2)$-approximation to optimal.

There are several drawbacks to the above approach. First, identifying the greedy choice of paper in the picking sequence requires fully iterating through the Round Robin picking sequence for each paper at every step. Thus the algorithm takes $O(n^3k)$ time. In addition, Round Robin represents a restricted set of all EF1 allocations, so guaranteeing high welfare relative to the optimal Round Robin picking sequence may not be meaningful in all contexts. We thus extend the notion of optimizing picking sequences to the broader class of recursively balanced picking sequences. A picking sequence is recursively balanced if all agents pick k times before any agent picks $k + 1$ times, for all k. Recursively balanced picking sequences are still EF1, since for any agents i and j, i does not envy j after removal of j's first item chosen. Just as with Reviewer Round Robin, we must modify the algorithm to remain EF1 for reviewer assignment. Our greedy algorithm picks the best paper to go next in the recursively balanced picking sequence at each step. Finding the greedy choice is as simple as checking the affinity between each available paper and the reviewer it would choose, and the overall runtime becomes $O(n^2k)$. We also find empirically that the total welfare tends to be much higher under the greedy recursively balanced picking sequence compared to the greedy Round Robin picking sequence. Although this second approach does not appear in [10], we are in the process of integrating a version of the greedy recursively balanced picking sequence into the OpenReview matcher[1] and hope to include it in future versions of the paper.

3 OTHER APPLICATIONS

3.1 Reviewer Bidding

Fair allocation algorithms typically assume access to complete and truthful valuations. Although prior work [4, 6, 10, 14] interprets reviewer-paper affinities as paper valuations, this interpretation can break down in some cases (for example, a reviewer with high affinity may decide to reject the paper). We might address this discrepancy by eliciting bids from reviewers then allocating *papers to reviewers* according to the bids. Meir et al. [8] encourage reviewers to bid on papers with low demand by giving reviewers a suggested quota and paying more for papers with fewer bids. This approach does not necessarily elicit the minimum number of bids needed to ensure a fair and efficient allocation. We might instead directly elicit bids that are likely to increase the welfare of the final allocation. With any bidding mechanism, we will still obtain a limited preference profile for each reviewer, necessitating fair allocation algorithms that operate under uncertainty. We might also want to encourage reviewers to bid on papers with high affinity, so that reviewers bid on papers in their expertise. We are also interested in testing bidding mechanisms for chore assignment using small tasks on Amazon Mechanical Turk or in a controlled lab setting.

3.2 Fair Two-Sided Matching

Existing work on reviewer assignment algorithms is either fair to papers or reviewers, but not both. Fair two-sided matching papers have only considered the scenario when both sides consider the other side to be goods, not chores. In fair two-sided reviewer assignment, the papers consider the reviewers to be goods, while the reviewers consider the papers to be chores. This is an interesting novel setting for the fair two-sided matching problem. We might hope to achieve EF1 for both sides, since Freeman et al. [5] study this solution concept and do not rule it out when both sides are considered goods. We also may be able to design an algorithm with simpler two-sided fairness guarantees, such as guaranteeing that reviewers are all assigned nearly the same number of papers while ensuring maximum egalitarian welfare or EF1 for the papers.

4 CONCLUSION

The fields of mechanism design and fair allocation provide powerful frameworks for many of the common problems in automated peer review support systems. These approaches offer provable guarantees, along with easy-to-understand algorithms that scale to our increasingly larger conferences. We hope to provide new algorithms (and integrate them into OpenReview) as options for conference organizers to compare and contrast with existing approaches. Despite these algorithms' provable guarantees on metrics of interest, the ultimate test of their effectiveness will come from using and evaluating them in practice. In the end, we hope that our work will enable conference organizers to sculpt the best environment for conference participants and broader scientific advancement.

[1] https://github.com/openreview/openreview-matcher

REFERENCES

[1] Haris Aziz, Thomas Kalinowski, Toby Walsh, and Lirong Xia. 2016. Welfare of sequential allocation mechanisms for indivisible goods. In *Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI)*. 787–794.

[2] Haris Aziz, Toby Walsh, and Lirong Xia. 2015. Possible and necessary allocations via sequential mechanisms. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*.

[3] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D Procaccia, Nisarg Shah, and Junxing Wang. 2019. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation (TEAC)* 7, 3 (2019), 1–32.

[4] Laurent Charlin and Richard Zemel. 2013. The Toronto paper matching system: an automated paper-reviewer assignment system. In *Proceedings of the 2013 ICML Workshop on Peer Reviewing and Publishing Models*.

[5] Rupert Freeman, Evi Micha, and Nisarg Shah. 2021. Two-Sided Matching Meets Fair Division. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*. 203–209.

[6] Ari Kobren, Barna Saha, and Andrew McCallum. 2019. Paper matching with local fairness constraints. In *Proceedings of the 25th International Conference on Knowledge Discovery and Data Mining (KDD)*. 1247–1257.

[7] Alison McCook. 2006. Is peer review broken? Submissions are up, reviewers are overtaxed, and authors are lodging complaint after complaint about the process at top-tier journals. What's wrong with peer review? *The scientist* 20, 2 (2006), 26–35.

[8] Reshef Meir, Jérôme Lang, Julien Lesca, Nicholas Mattei, and Natan Kaminsky. 2021. A market-inspired bidding scheme for peer review paper assignment. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*.

[9] Gourab K Patro, Arpita Biswas, Niloy Ganguly, Krishna P Gummadi, and Abhijnan Chakraborty. 2020. Fairrec: Two-sided fairness for personalized recommendations in two-sided platforms. In *Proceedings of the 29th The Web Conference (TheWebConf)*. 1194–1204.

[10] Justin Payan and Yair Zick. 2021. I Will Have Order! Optimizing Orders for Fair Reviewer Assignment. In *Workshop on Cooperative AI at NeurIPS 2021*.

[11] Drummond Rennie. 2016. Let's make peer review scientific. *Nature News* 535, 7610 (2016), 31.

[12] Nihar B Shah. 2019. *Principled methods to improve peer review*. Technical Report. Carnegie Mellon University.

[13] Ivan Stelmakh. 2021. Towards Fair, Equitable, and Efficient Peer Review. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, Vol. 35. 15736–15737.

[14] Ivan Stelmakh, Nihar B Shah, and Aarti Singh. 2019. PeerReview4All: Fair and accurate reviewer assignment in peer review. In *Proceedings of the 30th International Conference on Algorithmic Learning Theory (ALT)*. 828–856.