I Will Have Order! Optimizing Orders for Fair Reviewer Assignment

Extended Abstract

Justin Payan
University of Massachusetts Amherst
Amherst, United States
jpayan@umass.edu

Yair Zick
University of Massachusetts Amherst
Amherst, United States
yzick@umass.edu

ABSTRACT

We study mechanisms that allocate reviewers to papers in a fair and efficient manner. We model reviewer assignment as an instance of a fair allocation problem, presenting an extension of the classic round-robin mechanism, called Reviewer Round Robin (RRR). Round-robin mechanisms are a standard tool to ensure envy-free up to one item (EF1) allocations. However, fairness often comes at the cost of decreased efficiency. To overcome this challenge, we carefully select an approximately optimal round-robin order. Applying a relaxation of submodularity, γ-weak submodularity, we show that greedily inserting papers into an order yields a $(1+\gamma^2)$-approximation to the maximum welfare attainable by our round-robin mechanism under any order. Our approach outputs highly efficient EF1 allocations for three real conference datasets, offering comparable performance to state-of-the-art paper assignment methods in fairness, efficiency, and runtime, while providing the only EF1 guarantee.

KEYWORDS

Fair Allocation; Envy-Free up to One Item; Reviewer Assignment

ACM Reference Format:
Justin Payan and Yair Zick. 2022. I Will Have Order! Optimizing Orders for Fair Reviewer Assignment: Extended Abstract. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), Online, May 9–13, 2022*, IFAAMAS, 3 pages.

1 OVERVIEW

Peer review plays a prominent role in nearly all aspects of academia. Given the broad application of peer review and its significant gatekeeping role, it is imperative that this process remains as objective as possible. For this process to be successful, reviewers must possess the proper expertise for their assigned papers. Overall assignment accuracy maintains quality standards for academic publications. However, it is imperative that we do not sacrifice review quality on some papers to obtain higher overall matching scores. Papers with poorly matched reviewers may be unfairly rejected or receive unhelpful feedback, causing the authors real harm. We thus desire algorithms which are globally accurate and fair.

Our principal fairness criterion is *envy-freeness*: one paper envies another paper if it prefers the other's assigned reviewers over its own. Although papers cannot directly compare their assigned reviewers, envy-freeness and its relaxations preclude large, unjustified disparities in reviewer-paper alignment scores. Reducing envy also ensures that we cannot improve low-scoring papers without significantly harming those papers at the top.

It is generally not possible to obtain envy-free allocations for indivisible items [4], so we focus on the relaxed criterion of envy-freeness up to one item (EF1) [6, 12]. EF1 allocations require that when a paper i has higher affinity for the reviewers of a paper j, it is due to a single, high-affinity reviewer rather than a complete imbalance in outcomes. In standard fair allocation settings, the well-known *round-robin* (RR) mechanism produces EF1 allocations by assigning each agent its highest-scoring remaining item, one at a time in a fixed agent order [7]. The additional constraints of reviewer assignment break round-robin's EF1 guarantees, so we present *Reviewer Round Robin* (RRR) for EF1 reviewer assignments.

While RR mechanisms are known to satisfy fairness constraints, their efficiency guarantees depend on the order in which players pick items. Consider a stylized setting with papers i and j and reviewers r_1 and r_2: paper i has affinity 5 for both reviewers, while paper j has affinity 10 for r_1 and 0 for r_2. A round-robin mechanism that assigns to i first might assign r_1 to i, leaving j with r_2. Assigning to j first results in a much better outcome, without compromising on fairness. It is NP-hard to maximize welfare under round-robin [1, 2], and maximizing welfare subject to EF1 in general is not approximable in polynomial time [3]. We thus answer the question: Can we identify *approximately optimal* player orders?

Using techniques from submodular optimization, we run a combinatorial search for orders of papers that yield high efficiency allocations for picking-sequence mechanisms like RRR. We optimize a function on partial paper sequences, which varies according to the welfare of the allocation resulting from the picking sequence. This function is not submodular in general, but we can capture its distance from submodularity via a variable γ. Our main theoretical result (Theorem 3.1), which is of independent interest to the fair division community, shows that a simple greedy approach maximizes this function up to a factor of $1+\gamma^2$. Please refer to the more extensive version of the paper for algorithm details and proofs[1].

2 PRELIMINARIES

We are given a set of papers $N = \{1, \dots, n\}$, and a set of reviewers $R = \{r_1, r_2, \dots, r_m\}$. Each paper i has an affinity function over reviewers $v_i : R \to \mathbb{R}_{\geq 0}$. The affinity typically models alignment between reviewer expertise and paper topics, but can incorporate other elements like reviewer bids, conflicts of interest, and author suggestions; several works study how these values are generated

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.

[1] Available on arxiv at https://arxiv.org/abs/2108.02126. A version of this paper was also included in the Workshop on Cooperative AI at NeurIPS 2021 [13].

(see the survey by Wang et al. [15]), and are orthogonal to our work. We assume *additive* valuations, where for a paper $i \in N$ and a subset $S \subseteq R$, $v_i(S) = \sum_{r \in S} v_i(r)$. An *assignment* or *allocation* of reviewers to papers is an ordered tuple $\mathcal{A} = (A_1, A_2, \ldots A_n)$ where each *bundle* $A_i \subseteq R$ is a set of distinct reviewers assigned to paper i. Each reviewer $r \in R$ has an upper bound u_r on the number of papers they can review, and each paper requires k reviewers.

We seek allocations that maximize utilitarian social welfare subject to the EF1 constraint. An allocation $\mathcal{A}$ is EF1 if for all pairs of papers i and j, $\exists r \in A_j$ such that $v_i(A_j \setminus \{r\}) \leq v_i(A_i)$. The *utilitarian social welfare* (USW) of an allocation $\mathcal{A}$ is equal to $\sum_{i \in N} v_i(A_i)$.

For round-robin, we define an *order* on papers $i \in N$ as a tuple $O = (S, o)$, where $S \subseteq N$ is the set of papers in the order and $o : S \to [|S|]$ is a permutation on S mapping papers to positions. We can think of an order $O = (S, o)$ as an ordered list $[o_1, o_2, \ldots o_{|S|}]$ such that $o_l = o^{-1}(l)$ for all positions l. We use the notation $O + i$ to indicate the order (S', o') that appends i to the end of O.

A set function $f : 2^E \to \mathbb{R}$ is *monotone* if for all $X \subseteq Y \subseteq E$, $f(X) \leq f(Y)$. Given set function f, set $X \subseteq E$, and element $e \in (E \setminus X)$, denote the marginal gain of adding e to X under f as $\rho_e^f(X) = f(X + e) - f(X)$ or simply $\rho_e(X)$ if f is understood. Given a monotone, non-negative set function f, we say that f is γ-weakly submodular if for all $X \subseteq Y \subseteq E$ and $e \in E \setminus Y$, $\gamma \rho_e^f(X) \geq \rho_e^f(Y)$.

3 FAIR AND EFFICIENT REVIEWER ASSIGNMENT

To ensure EF1, our allocations draw upon the simple and well-known round-robin mechanism for assigning goods to agents. We have the additional constraint that papers must select at most k distinct reviewers. We demonstrate via an example that standard round-robin, modified in the obvious way to meet this constraint, can fail to be EF1. Our Reviewer Round Robin algorithm produces reviewer assignments which satisfy all constraints and are EF1. Any time a paper would select a reviewer such that EF1 would be violated, we require the paper to select a different reviewer.

We also introduce variants of Reviewer Round Robin when papers require different numbers of reviewers and when each reviewer must receive a minimum number of papers to review. These modified problem settings are often used by conference organizers to account for late reviews or borderline papers, or to ensure more balanced workloads for the reviewers. The case when papers require different numbers of reviewers is of particular interest, as we guarantee the fairness criterion of *weighted* envy-freeness up to one item (WEF1) from Chakraborty et al. [8] instead of EF1. The algorithm to ensure WEF1 differs from round-robin, since papers are assigned reviewers according to a shifting priority order rather than in a fixed sequence each round. There are numerous ties in priority, and we optimize the order in which these ties are resolved.

We present a greedy approach to maximize the USW of our picking sequence-based mechanisms by optimizing over the *ordering* of the papers. Given one of our mechanisms $\mathcal{M}$ (i.e. RRR), we define a function $\text{USW}_{\mathcal{M}}(O)$, which represents the USW from running $\mathcal{M}$ on agents in the order O. The greedy algorithm maintains an order O, always adding the paper i which maximizes $\text{USW}_{\mathcal{M}}(O + i)$.

Our proof technique draws from optimization of submodular set functions. We identify orders O_P with sets of tuples P, where P consists of tuples (i, j) mapping papers to positions in the order. We create a function $f(P) = \text{USW}_{\mathcal{M}}(O_P)|P|^{\alpha}$ over tuples, where α is selected as the smallest value such that $f(P)$ is monotonically increasing. We show that our greedy algorithm is equivalent to greedily maximizing $f(P)$, which is γ-weakly submodular. Our main result, Theorem 3.1, proves the $(1 + \gamma^2)$-approximation factor.

Theorem 3.1. *Suppose that f is the monotonically increasing, γ-weakly submodular function $f(P) = USW_{\mathcal{M}}(O_P)|P|^{\alpha}$. The set P^{alg} returned by the greedy algorithm satisfies $f(P^{alg}) \geq \frac{1}{1+\gamma^2} f(P^*)$, where O_{P^*} is the optimal paper order for $\mathcal{M}$.*

When $\gamma = 1$ (and thus f is submodular), Theorem 3.1 yields a $\frac{1}{2}$-approximation guarantee, which beats the $\frac{1}{3}$-approximation guarantee provided by Fisher et al. [10]. This improvement is possible because our algorithm can always select a tuple (i, j) appending i to the end of the order O_P, enabling a more tailored analysis. The greedy algorithm is a tight $\frac{1}{2}$-approximation for submodular maximization in the *unconstrained* regime [5], which our result matches even though we operate in a constrained (albeit less general) space.

4 EXPERIMENTS

We run experiments[2] on three publicly available[3] conference datasets: MIDL, CVPR 2018, and one older CVPR iteration. We compare against the widely-used Toronto Paper Matching System or TPMS (which maximizes USW with no fairness criterion) [9], FairFlow (currently implemented in OpenReview[4]) [11], and PeerReview4All or PR4A (a fair assignment algorithm used by ICML 2020) [14].

In terms of welfare, our approach consistently outperforms FairFlow, though PR4A has slightly better welfare. All of the baselines violate EF1 on both CVPR conferences, with FairFlow and TPMS showing a large number of violations. We also calculated additional global fairness metrics, finding our approach outperforms FairFlow and TPMS but not PR4A. It appears that the local fairness of EF1 may require trade-offs in these other global notions of fairness.

5 CONCLUSION

Our algorithm ensures EF1, while remaining highly competitive on welfare and other fairness guarantees. Moreover, the greedy algorithm is easy to implement and understand, lending it further appeal. Our approach of optimizing over *orders* for picking sequences is of independent interest, and may inspire further study of optimal picking sequences. Finally, there are many applications of different fairness, efficiency, robustness, or non-manipulability constraints from the fair allocation literature to problems in peer review, which could bring some much-needed rigor to this fundamental process.

ACKNOWLEDGMENTS

We thank Purujit Goyal, Melisa Bok, and Andrew McCallum from OpenReview for help developing our problem statement, Vignesh Viswanathan for teaching us about submodular optimization, and the reviewers of the Cooperative AI Workshop at NeurIPS 2021, WINE 2021, and AAMAS 2022.

[2] Implementations of our approach and all baselines are available at https://github.com/justinpayan/ReviewerAssignmentCode.
[3] https://github.com/iesl/fair-matching/tree/master/data
[4] 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)*. 468–474.

[3] Siddharth Barman, Ganesh Ghalme, Shweta Jain, Pooja Kulkarni, and Shivika Narang. 2019. Fair Division of Indivisible Goods Among Strategic Agents. In *Proceedings of the 18th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. 1811–1813.

[4] Sylvain Bouveret, Yann Chevaleyre, and Nicolas Maudet. 2016. Fair Allocation of Indivisible Goods. In *Handbook of computational social choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia (Eds.). Cambridge University Press, Chapter 12, 284–309.

[5] Niv Buchbinder, Moran Feldman, Joseph Naor, and Roy Schwartz. 2012. A Tight Linear Time (1/2)-Approximation for Unconstrained Submodular Maximization. In *Proceedings of the 53rd Symposium on Foundations of Computer Science (FOCS)*. 649–658.

[6] Eric Budish. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy* 119 (12 2011), 1061–1061.

[7] 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.

[8] Mithun Chakraborty, Ayumi Igarashi, Warut Suksompong, and Yair Zick. 2020. Weighted Envy-Freeness in Indivisible Item Allocation. In *Proceedings of the 19th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. 231–239.

[9] 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*.

[10] M. L. Fisher, G. L. Nemhauser, and L. A. Wolsey. 1978. An analysis of approximations for maximizing submodular set functions—II. In *Polyhedral Combinatorics*, M. L. Balinski and A. J. Hoffman (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 73–87.

[11] 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.

[12] Richard J Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*. 125–131.

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

[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.

[15] Fan Wang, Ben Chen, and Zhaowei Miao. 2008. A survey on reviewer assignment problem. In *Proceedings of the 21st International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems*. 718–727.