I Will Have Order! Optimizing Orders for Fair Reviewer Assignment

Justin Payan and **Yair Zick**
University of Massachusetts Amherst
{jpayan, 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 Greedy Reviewer Round Robin (GRRR) 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.

1 Introduction

Peer review plays a prominent role in nearly all aspects of academia. It serves a number of functions – selecting the best manuscripts, assessing originality, providing feedback, and more [Mulligan *et al.*, 2013]. Given the broad application of peer review and its significant gatekeeping role, it is imperative that this process remains as objective as possible.

One important parameter is whether reviewers possess the proper expertise for their assigned papers. Selecting reviewers for submitted papers is therefore a crucial first step of any reviewing process. In large conferences such as NeurIPS/ICML/AAAI/IJCAI, reviewer assignment is largely automated through systems such as the Toronto Paper Matching System (TPMS) [Charlin and Zemel, 2013], Microsoft CMT[1], or OpenReview[2]. Inappropriately assigned reviewers may lead to failures: misinformed decisions, reviewer disinterest, and a general mistrust of the peer-review process.

[1]https://cmt3.research.microsoft.com/

[2]https://openreview.net

Accuracy and fairness are two important criteria in reviewer assignment [Shah, 2019], and fair division in general [Bouveret *et al.*, 2016]. 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 for the papers. To accomplish these goals, we consider the fair reviewer assignment problem from the lens of *fair allocation*.

Our principal fairness criterion is *envy*: 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. State-of-the-art fair reviewer assignment algorithms either maximize the minimum paper score or ensure all scores exceed some threshold [Stelmakh *et al.*, 2019; Kobren *et al.*, 2019]. Although this approach ensures a higher floor, there may still be many low-scoring papers that could benefit from trading reviewers with high-scoring papers. Low envy 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 [Bouveret *et al.*, 2016], so we focus on the relaxed criterion of envy-freeness up to one item (EF1) [Budish, 2011; Lipton *et al.*, 2004]. EF1 allocations have the property that whenever 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. Maximizing welfare subject to EF1 is NP-hard and is not approximable in polynomial time [Barman *et al.*, 2019]. In standard fair allocation settings, the well-known *round-robin* (RR) mechanism produces EF1 allocations by setting an order of agents, and letting them select one item at a time. Due to constraints on reviewer selection, round-robin is not EF1 for reviewer assignment. We thus present a variation on classic RR, which we term *Reviewer Round Robin* (RRR).

While RR mechanisms are known to satisfy fairness constraints, their efficiency guarantees are highly dependent on the order in which players pick items. For example, consider a stylized setting where there are two papers (i and j) and two reviewers (r_1 and r_2): paper i has an affinity score of 5 for both reviewers, while paper j has a score of 10 for r_1 and 0 for r_2. A round-robin mechanism that lets i pick first runs the risk of having i pick r_1, leaving j with r_2. Letting j pick first results in a much better outcome, without compromising on fairness. It is generally difficult to identify optimal picking sequences [Bouveret and Lang, 2011; Kalinowski *et al.*, 2013; Aziz *et al.*, 2015; Aziz *et al.*, 2016]. We thus answer the question: Can we identify *approximately optimal* paper orders?

1.1 Our Contributions

We run a combinatorial search for orders of papers that yield high efficiency allocations for picking-sequence mechanisms like RRR. To this end, we examine the problem of finding an optimal paper order via the lens of *submodular optimization*. 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 3.2), which is of independent interest to the fair division community, shows that a simple greedy approach maximizes this function up to a factor of $1 + \gamma^2$. We call this approach Greedy Reviewer Round Robin (GRRR). To our knowledge, ours is the first work to greedily optimize picking sequences by allowing the locally optimal agent to pick at each step.

We compare our GRRR algorithm with four other state-of-the-art paper assignment frameworks on three real-world conference datasets. Not only is RRR/GRRR the only provably EF1 approach, it is considerably faster than the only other method with similar fairness metrics. Finally, GRRR maintains high utility guarantees, offering comparable performance to the state-of-the-art benchmarks.

Many proofs, detailed descriptions of algorithm variants, and experimental results are presented in the full version of the paper [Payan and Zick, 2021]. At the time of publication we are also in the process of integrating one such variant, called FairSequence, in OpenReview.

1.2 Related Work

The reviewer assignment problem has been widely studied. Most works model the problem as a mixed-integer linear program maximizing affinity between reviewers and papers; the Toronto Paper Matching System (TMPS) is the most notable work with this formulation [Charlin and Zemel, 2013]. The affinity typically models alignment between reviewer expertise and paper topics, but can incorporate other relevant notions like reviewer bids, conflicts of interest, and author suggestions; several works study how these values are generated and are orthogonal to our work. Affinities are generally considered a good proxy for value at both an individual and collective level, since higher-affinity reviewers will typically be more qualified for and interested in the paper, resulting in more detailed and accurate reviews. Affinity scores are universally available in systems like TPMS, Microsoft CMT, or OpenReview, and it is standard practice to use these affinity scores to compute welfare and fairness measures [Stelmakh *et al.*, 2019; Kobren *et al.*, 2019].

A number of prior works consider fairness objectives in peer review, though none of them consider envy-freeness up to one item. The most recent approaches maximize the minimum paper score, or maximize the sum of scores subject to a minimum individual score threshold [Kobren *et al.*, 2019; Stelmakh *et al.*, 2019].

A few works in fair allocation are relevant for fair reviewer assignment, but they have important limitations in our setting. Aziz *et al.* [2019] present an algorithm which attempts to output a W-satisfying EF1 allocation for constraint(s) W. When W includes a minimum threshold for welfare, their approach is somewhat similar to ours. However, rather than greedily maximizing welfare by letting the locally optimal agent pick, they let any agent pick as long as W can still be achieved. Biswas and Barman [2018] present a modification of the round-robin mechanism that assigns a complete EF1 allocation when items are partitioned into categories and agents can receive a limited number of items from each category, but overall number of items per agent is unlimited. Dror *et al.* [2021] study fair allocation under matroid constraints, but only for identical or binary valuations, less than four agents, or a single uniform matroid constraint.

Our application of submodular optimization to optimizing orders for round-robin is inspired by previous work on fair allocation with submodular *valuations* [Babaioff *et al.*, 2021; Benabbou *et al.*, 2020]. Prior work has also studied maximization of approximately submodular functions, though none has combined matroid constraints with a definition of approximate submodularity similar to ours [Das and Kempe, 2011; Gölz and Procaccia, 2019].

Existing work shows the hardness of maximizing welfare for EF1 and picking sequence allocations. Aziz *et al.* [2019] show maximizing welfare subject to EF1 is NP-hard, and Barman *et al.* [2019] show the same problem is not even polynomial-time approximable. Aziz *et al.* [2016] show that the problem of determining if a given welfare is possible under a picking sequence of a certain class is NP-complete for some classes of picking sequences (but not round-robin).

2 Preliminaries

We represent reviewer assignment as a problem of allocating indivisible goods, with papers as agents and reviewers as goods. To simplify notation, given a set X and an element y, we often write $X + y$ instead of $X \cup \{y\}$. We are given a set of papers $N = \{1, \ldots, n\}$, and a set of reviewers $R = \{r_1, r_2, \ldots, r_m\}$. Each paper i has an affinity function over reviewers $v_i : R \to \mathbb{R}_{\geq 0}$, which defines the alignment between the paper and the reviewer.

Papers generally receive more than one reviewer, so we define affinity functions over sets of reviewers. We assume *additive* functions, 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 $A_i \subseteq R$ is a set of distinct reviewers assigned to paper i. We can refer to A_i as paper i's *bundle*, and $v_i(A_i)$ as paper i's *valuation* under $\mathcal{A}$.

Each reviewer $r \in R$ has an upper bound u_r on the number of papers they can review, and may have lower bounds

l_r, which help ensure a more even distribution of work. No reviewer can be assigned to the same paper twice. Given an allocation $\mathcal{A}$, we denote the number of papers to which a reviewer r is assigned as $c_r^{\mathcal{A}} = \sum_{i \in N} |\{r\} \cap A_i|$, or just c_r when $\mathcal{A}$ is clear from context. Each paper i requires k_i reviewers. Often we have a fixed k such that $k_i = k$ for all i. An allocation is *complete* if every paper i is assigned k_i distinct reviewers (and *incomplete* otherwise).

We now discuss our notion of fairness. An allocation $\mathcal{A}$ is considered *envy-free* if for all pairs of papers i and j, $v_i(A_j) \leq v_i(A_i)$. This criterion is not achievable in general (consider the example of two papers and one reviewer r whose upper bound is $u_r = 1$), so we relax the criterion. An allocation $\mathcal{A}$ is envy-free up to one item (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* ("utilitarian welfare" or "USW") of an allocation $\mathcal{A}$ is the sum of the papers' valuations under that allocation. USW is a natural objective in the context of reviewer assignment, and has been used in many prior works on this topic [Charlin and Zemel, 2013; Kobren *et al.*, 2019; Stelmakh *et al.*, 2019].

We also use the *Nash social welfare* or "NSW" as a second measure of efficiency in our experimental evaluation. The NSW is equal to the product of papers' valuations. When some agents have valuation of 0, we instead report the number of papers with 0 valuation and the Nash welfare of the remaining papers. Nash welfare is another common efficiency measure, and allocations with high NSW provide a balance of efficiency and fairness [Caragiannis *et al.*, 2019].

For round-robin, we define an *order* on papers $i \in N$ as a tuple $\mathcal{O} = (S, o)$, where $S \subseteq N$ is the set of papers in the order and $o : S \rightarrow [|S|]$ is a permutation on S mapping papers to positions. We slightly abuse notation and say that a paper $i \in \mathcal{O}$ if $i \in S$. For any $i, j \in \mathcal{O}$, we say that $i \prec_{\mathcal{O}} j$ if and only if $o(i) < o(j)$. We sometimes write $i \prec j$ when $\mathcal{O}$ is clear from context. We can think of an order $\mathcal{O} = (S, o)$ 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 $\mathcal{O} + i$ to indicate the order (S', o') that appends i to the end of $\mathcal{O}$. Formally, $S' = S \cup \{i\}$, $o'(j) = o(j)$ for $j \in S$, and $o'(i) = |S'|$.

3 Fair and Efficient Reviewer Assignment

We show how to obtain EF1 reviewer assignments when all papers have equal demands $k_i = k$, and reviewers do not have lower bounds l_r; we handle the more general case in the full version. Our algorithm draws upon the simple and well-known round-robin mechanism for assigning goods to agents. Given an ordered list of agents, round-robin proceeds in rounds. In each round, we iterate over the agents in the assigned order, assigning each agent its highest valued remaining good. The process terminates when all goods have been assigned. This allocation is EF1 for additive valuations by a simple argument [Caragiannis *et al.*, 2019]. For any agent i, we divide the item assignments into rounds specific to that agent i. Agent i prefers its own bundle to the bundle of any agent $j \succ i$, and it prefers its own bundle to that of any agent $j' \prec i$ if we ignore j''s good from the first round.

We have an additional constraint not present in the setting

Algorithm 1 Reviewer Round Robin (RRR)

```
Require: Reviewers R, reviewer upper limits u_r, paper order O, affinity functions v_i, bundle size limit k
  Initialize allocation A as A_i ← ∅ for all papers i ∈ O,
  Initialize the first-assigned reviewer F_i ← ∅ for all i
  Initialize the attempted set S_i ← ∅ for all i
  for Round t ∈ {1, ... k} do
     for i ∈ O in increasing order do
        for Reviewer r in decreasing order of v_i(r) (break ties lexicographically) do
           if c_r^A < u_r, and r ∉ A_i then
              Attempt to assign r to i (S_i ← S_i ∪ {r})
              if No j ≺ i with r ∈ S_j envies A_i ∪ {r} and no j ≻ i with r ∈ S_j envies (A_i ∪ {r}) \ F_i then
                 A_i ← (A_i ∪ {r})
                 If t = 1, F_i ← {r}
                 Move to the next paper
        If no new reviewer is assigned to i, return A
  return A
```

examined by Caragiannis *et al.* [2019]: papers must be assigned at most k distinct reviewers. A trivial modification of round-robin allows us to satisfy the cardinality constraint — proceed for exactly k rounds, then stop. We might naively update round-robin to satisfy the distinctness constraint as well, by assigning each paper the best reviewer they do not already have. However, the argument from Caragiannis *et al.* [2019] fails. To see why, suppose a paper i is assigned a reviewer r in one round. In the next round, i may still prefer r over any other reviewer, but we cannot assign it. We will be forced to assign i a much worse reviewer, giving another paper that desired "second copy" of r.

We present a modification of round-robin that produces reviewer assignments which satisfy all constraints and are EF1. Reviewer Round Robin or RRR (Algorithm 1) forbids any assignment that violates a crucial invariant for proving EF1. This invariant derives from the proof of EF1 in the additive case. Any time we would assign a reviewer such that EF1 would be violated, we forbid the assignment and instead assign a different reviewer. EF1 violations can only arise when another paper preferred that reviewer but the assignment was forbidden, either because it had been assigned already, or because it would have caused an EF1 violation for that paper as well. We always attempt to assign reviewers in preference order. Thus when we attempt to assign a reviewer r to paper i, we only need to check for EF1 violations against other papers to which we have attempted to assign r in the past. Theorem 3.1 asserts the correctness of RRR.

Theorem 3.1. *RRR terminates with an EF1 allocation where papers receive at most k distinct reviewers, and no reviewer r is assigned to more than u_r papers.*

It is possible for RRR to return an incomplete allocation when a complete one exists, but it is straightforward to show that RRR always returns a complete, EF1 allocation when the number of reviewers is large.

Proposition 3.1. *Given a reviewer assignment problem with*

m reviewers, n papers, and k paper bundle size limits, where $m \geq kn$, RRR returns a complete and EF1 allocation.

We hypothesize in Conjecture 3.1 that it is not always possible to assign all papers k reviewers and still achieve EF1; the existence of an instance for which no complete allocation is EF1 is left as an open problem.

Conjecture 3.1. *There exists a reviewer assignment problem with papers N, reviewers R, affinity functions $\{v_i\}_{i \in N}$, reviewer load bounds $\{u_r\}_{r \in R}$, and paper bundle size limit k such that there exists a complete allocation, but no complete allocation is EF1.*

Papers may sometimes require different numbers of reviewers, especially in real conference settings where conference organizers may run reviewer assignment multiple times to account for late reviews, borderline papers, and other mitigating circumstances. In addition, conference organizers might wish to require that each reviewer receives a minimum number of papers to review. These reviewer lower bounds ensure more balanced workloads for the reviewers. To satisfy these additional real-world constraints, we introduce variants of RRR that allow for variable paper demands k_i and reviewer lower bounds l_r. The full version of the paper contains detailed descriptions and analyses of these variants.

3.1 Optimizing Orders for RRR

We have shown how to provably obtain EF1 allocations of reviewers to papers, but have not offered any welfare guarantees so far. We first state that maximizing welfare under round-robin orders is NP-hard, which we prove in the full version by closely following existing techniques [Aziz *et al.*, 2015; Aziz *et al.*, 2016].

Proposition 3.2. *Maximizing welfare subject to round-robin (and RRR) is NP-hard.*

In this section, we present a simple greedy approach to approximately maximize the USW of our picking sequence by optimizing over the *ordering* of the papers. We present results using RRR (Algorithm 1), but all results apply equally well to the variants which handle variable paper demands and reviewer lower bounds. In each case, we can apply the results from this section to approximately maximize the utilitarian welfare for any of these mechanisms, over all orderings $\mathcal{O}$.

We define a function $\text{USW}_{\text{RRR}}(\mathcal{O}, k, R, \{u_r\}_{r \in R}, \{v_i\}_{i \in N})$, which represents the USW from running RRR on agents in the order $\mathcal{O}$ with reviewers R, reviewer upper limits u_r, affinity functions v_i, and paper bundle size limit k. When it is clear from context, we will drop most of the arguments, writing $\text{USW}_{\text{RRR}}(\mathcal{O})$ to indicate that we run RRR with the order $\mathcal{O}$ and all other parameters defined by the current problem instance. Our algorithm, which we call Greedy Reviewer Round Robin (GRRR), maintains an order $\mathcal{O}$, always adding the paper i which maximizes $\text{USW}_{\text{RRR}}(\mathcal{O} + i)$. It returns an order on papers, which can be directly input to RRR to obtain an EF1 allocation of reviewers. This algorithm is very simple and flexible. It admits trivial parallelization, as the function USW_{RRR} can be independently computed for each paper. One can also reduce runtime by subsampling the remaining papers at each step. Subsampling weakens the approximation guarantee in theory; while we do not attempt to analyze the approximation ratio of the subsampling approach in this work, we run our largest experiments with this variant, and still obtain high-welfare allocations. Let us now establish the welfare guarantees of GRRR.

We first review important concepts and terms used in the proof. A *matroid* [Oxley, 2011] is a pair $(E, \mathcal{I})$ with ground set E and independent sets $\mathcal{I}$, which must satisfy $\emptyset \in \mathcal{I}$. Independent sets must satisfy the inclusion property: $\forall A \subseteq B \in \mathcal{I}$, $A \in \mathcal{I}$, and the exchange property: $\forall A, B \in \mathcal{I}$ with $|A| < |B|$, $\exists e \in B \setminus A$ such that $A \cup \{e\} \in \mathcal{I}$. A *partition matroid* is defined using categories $B_1, B_2, \dots B_b$ such that $B_i \cap B_j = \emptyset$ for all i, j and $\bigcup_{1 \leq i \leq b} B_i = E$, and capacities $d_1, d_2, \dots d_b$; the independent sets are $\mathcal{I} = \{I \subseteq E : \forall i, |I \cap B_i| \leq d_i\}$. Given two matroids over the same ground set $(E, \mathcal{I}_1)$ and $(E, \mathcal{I}_2)$, the intersection of the two matroids is the pair $(E, \{I : I \in (\mathcal{I}_1 \cap \mathcal{I}_2)\})$. The intersection of two matroids may not be a matroid [Oxley, 2011].

We also use the notion of a submodular set function; submodular functions formalize the notion of diminishing marginal gains. For a set function $f : 2^E \to \mathbb{R}$, a set $X \subseteq E$, and an element $e \in (E \setminus X)$, we can write 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 from context. Given a set E, a function $f : 2^E \to \mathbb{R}$ is *submodular* if for all $X \subseteq Y \subseteq E$ and $e \in E \setminus Y$, $\rho_e^f(X) \geq \rho_e^f(Y)$. A set function is *monotone* if for all $X \subseteq Y \subseteq E$, $f(X) \leq f(Y)$. We define the notion of γ-weak submodularity for monotone, non-negative functions. Given a monotone, non-negative function $f : 2^E \to \mathbb{R}_{\geq 0}$, we say that f is γ-weakly submodular if for all $X \subseteq Y \subseteq E$ and $e \in E \setminus Y$, $\gamma \rho_e^f(X) \geq \rho_e^f(Y)$. When $\gamma = 1$ we recover submodularity, and we always have $\gamma \geq 1$.

GRRR is equivalent to greedily maximizing a γ-weakly submodular function over the intersection of two partition matroids; its guarantee worsens for larger values of γ.

Consider tuples of the form (i, j) where i is a paper and j represents a position in an order. We define a mapping from sets of tuples to orders. Consider the set $E = \{(i, j) : 1 \leq i, j \leq n\}$. Define two partition matroids $(E, \mathcal{I}_1)$ and $(E, \mathcal{I}_2)$, such that $\mathcal{I}_1$ forbids duplicating papers, and $\mathcal{I}_2$ forbids duplicating positions. Define $\mathcal{I}_1$ using categories $B_i^1 = \{(i, j) : 1 \leq j \leq n\}$, and $\mathcal{I}_1 = \{I \subseteq E : \forall i, |I \cap B_i^1| \leq 1\}$. Likewise, $\mathcal{I}_2$ is defined using categories $B_j^2 = \{(i, j) : 1 \leq i \leq n\}$, and $\mathcal{I}_2 = \{I \subseteq E : \forall j, |I \cap B_j^2| \leq 1\}$. Any set P in the intersection of these two matroids can be converted into a paper order $\mathcal{O}_P$ by sorting P on the position elements and outputting the paper elements in that order. Formally, given any set $P \in (\mathcal{I}_1 \cap \mathcal{I}_2)$, we construct an order $\mathcal{O}_P = (S_P, o_P)$ by taking $S_P = \{i \in N : \exists j, (i, j) \in P\}$. For all $(i, j) \in P$, let $j' = |\{(k, l) \in P : l \leq j\}|$ and set $o_P(i) = j'$. We extend this mapping to all subsets of E by sorting on the position elements as a primary key and paper elements as a secondary key, then deleting all but the first tuple for each paper.

With these constructions defined, we observe that maximizing the USW for RRR over a fixed number of rounds k is equivalent to the problem $\max_{P \in (\mathcal{I}_1 \cap \mathcal{I}_2) : |P| = n} \text{USW}_{\text{RRR}}(\mathcal{O}_P)$ for the matroids defined above. We will show that GRRR greedily maximizes a monotonically increasing version of our function over our two partition matroids. Next, we show that when our function is γ-weakly submodular, we can provide a γ-dependent approximation ratio.

To make $\text{USW}_{\texttt{RRR}}(\mathcal{O}_P)$ monotonically increasing, we will multiply by a factor of $|P|^\alpha$, where α is defined as the smallest positive number such that $f(P) = \text{USW}_{\texttt{RRR}}(\mathcal{O}_P)|P|^\alpha$ is monotonically increasing. We first prove that GRRR greedily maximizes $f(P)$. Formally, Lemma 3.1 states that GRRR selects the element i maximizing $f(P+(i,j))$ at each iteration.

Lemma 3.1. *Let $f(P) = USW_{RRR}(\mathcal{O}_P)|P|^\alpha$ for some α such that f is monotonically non-decreasing. Suppose GRRR selects paper i_t at each round t, resulting in a set of tuples P_t. Then for all t, (i_t, t) maximizes $f(P_{t-1}+(i,j))$ over all (i,j) such that $P_{t-1}+(i,j) \in (\mathcal{I}_1 \cap \mathcal{I}_2)$.*

The greedy algorithm for maximizing $f(P)$ terminates when $|P| = n$, so we must also ensure GRRR terminates with an order on all n papers. Although GRRR only considers $\text{USW}_{\texttt{RRR}}(\mathcal{O})$, which may not be monotonically increasing, by construction it runs until reaching a full order over all papers. Therefore, GRRR is equivalent to greedily maximizing f.

We are now ready to prove the $1+\gamma^2$ approximation ratio for GRRR (Theorem 3.2). Our proof is inspired by the proof in [Fisher *et al.*, 1978] that a similar greedy algorithm gives a $\frac{1}{p+1}$-approximation for maximizing a monotone submodular function over the intersection of p matroids. However, the introduction of γ-weak submodularity changes the proof.

Theorem 3.2. *Suppose that f is the monotonically increasing, γ-weakly submodular function $f(P) = USW_{RRR}(\mathcal{O}_P)|P|^\alpha$. The set P^{alg} returned by GRRR satisfies $f(P^{alg}) \geq \frac{1}{1+\gamma^2} f(P^*)$, where $\mathcal{O}_{P^*}$ is the optimal paper order for RRR.*

Proof. Let P_t represent the subset of P^{alg} after the t-th step of GRRR, where we add the element (i_t, t) to P_{t-1}. Let (i_t^*, t) denote the pair in P^* which places paper i_t^* in position t. Denote $L = |P^* \setminus P^{\text{alg}}|$. Consider the elements of $P^* \setminus P^{\text{alg}} = \{(i^*_{t_1}, t_1), \dots (i^*_{t_L}, t_L)\}$, ordered so that $t_1 < t_2 < \dots t_L$. Let $P^{\text{alg}} \cup \{(i^*_{t_1}, t_1), \dots (i^*_{t_l}, t_l)\}$ be denoted as P^{alg}_l (with $P^{\text{alg}}_0 = P^{\text{alg}}$). By monotonicity of f, $f(P^*)$ is bounded from above by:

$$f(P^{\text{alg}} \cup P^*) = f(P^{\text{alg}}) + \sum_{l=1}^{L} \rho_{(i^*_{t_l}, t_l)}(P^{\text{alg}}_{l-1}). \quad (1)$$

By γ-weak submodularity of f, we have that

$$\rho_{(i^*_{t_l}, t_l)}(P^{\text{alg}}_{l-1}) \leq \gamma \rho_{(i^*_{t_l}, t_l)}(P^{\text{alg}}). \quad (2)$$

Equality 1 and inequality 2 imply that

$$f(P^*) \leq f(P^{\text{alg}}) + \gamma \sum_{l=1}^{L} \rho_{(i^*_{t_l}, t_l)}(P^{\text{alg}}),$$

which (again by monotonicity of f) is bounded by

$$f(P^{\text{alg}}) + \gamma \sum_{t=1}^{n} \rho_{(i^*_t, t)}(P^{\text{alg}}).$$

Alg.	USW	NSW	Min Score	EF1 Viol.
FairFlow	1.67	1.56	0.77	8813
FairIR	2.05	1.84	0.27	35262
TPMS	2.08	1.99	0.00	473545
PR4A	1.96	1.89	0.77	82
GRRR	1.82	1.72	0.00	0

Table 1: Results on the CVPR conference.

By γ-weak submodularity of f, we know that for all t, $\rho_{(i^*_t, t)}(P^{\text{alg}}) \leq \gamma \rho_{(i^*_t, t)}(P_{t-1})$. Thus,

$$f(P^*) \leq f(P^{\text{alg}}) + \gamma^2 \sum_{t=1}^{n} \rho_{(i^*_t, t)}(P_{t-1}). \quad (3)$$

Next, we claim that for all t,

$$\rho_{(i^*_t, t)}(P_{t-1}) \leq \rho_{(i_t, t)}(P_{t-1}). \quad (4)$$

At step t, the greedy algorithm chose to add (i_t, t) to P_{t-1}, with i_t maximizing $f(P_{t-1} + (i_t, t))$. If (i^*_t, j) is not present in P_{t-1} for any j, then the greedy algorithm would have considered adding (i^*_t, t) and determined that i_t was better. Suppose that $(i^*_t, j) \in P_{t-1}$ for some j. The greedy algorithm proceeds by filling positions from left to right, so $j \leq t-1$. By the definition of our mapping from sets to orders, i^*_t will take position j and ignore (i^*_t, t). Thus $\rho_{(i^*_t, t)}(P_{t-1}) = 0 \leq \rho_{(i_t, t)}(P_{t-1})$. In either case, inequality (4) holds. Combining (3) with (4) yields

$$f(P^*) \leq f(P^{\text{alg}}) + \gamma^2 \sum_{t=1}^{n} \rho_{(i_t, t)}(P_{t-1}) = (1+\gamma^2) f(P^{\text{alg}}).$$

□

When $\gamma = 1$ (and thus f is submodular), Theorem 3.2 yields a $\frac{1}{2}$-approximation guarantee, which beats the $\frac{1}{3}$-approximation guarantee provided by Fisher *et al.* [1978]. The greedy algorithm is a tight $\frac{1}{2}$-approximation for submodular maximization in the *unconstrained* regime [Buchbinder *et al.*, 2012], which our result matches even though we operate in a constrained (albeit less general) space.

4 Experimental Results

We run experiments on the three conference datasets used by Kobren *et al.* [2019]: Medical Imaging and Deep Learning (MIDL), Conference on Computer Vision and Pattern Recognition (CVPR), and the 2018 iteration of CVPR. MIDL is an order of magnitude smaller than CVPR and CVPR'18, and CVPR'18 is slightly less challenging than CVPR due to a higher ratio of reviewing slots to paper demands. For space we provide results for CVPR and a summary of the other two conferences here; full results and a description of all datasets can be found in the full version of the paper.

We compare our methods to the FairIR and FairFlow algorithms [Kobren *et al.*, 2019], the Toronto Paper Matching System (TPMS) [Charlin and Zemel, 2013], and PeerReview4All (PR4A) [Stelmakh *et al.*, 2019]. FairFlow is currently implemented in OpenReview, and it is in widespread use. TPMS (also in widespread use) provides an upper bound on welfare without fairness guarantees. PR4A has been used in one conference, ICML 2020 [Stelmakh, 2021]. PR4A has a larger runtime than all other algorithms, and it does not support reviewer lower bounds. FairIR has not been used for any conferences to date, perhaps due to its longer runtime.

Following recent works [Kobren *et al.*, 2019; Stelmakh *et al.*, 2019], we only run one iteration of PR4A on CVPR and CVPR'18. On those two conferences, PR4A maximizes the minimum paper score, but stops before maximizing the next smallest score. We run FairIR and FairFlow using the default configuration of the algorithms in the code from Kobren *et al.* [2019]. We also implemented the Constrained Round Robin algorithm [Aziz *et al.*, 2019], but estimated its runtime to be infeasible after testing on MIDL. Implementations of all algorithms are available on GitHub[3].

We report the USW, NSW, minimum paper score, and number of EF1 violations for each algorithm. We report the USW normalized by the number of papers n (equivalent to the mean paper score), and for NSW we report the n-th root of the product of non-zero paper scores. Thus USW and NSW are directly comparable, representing the arithmetic and geometric mean paper scores respectively. For an allocation $\mathcal{A}$, the number of EF1 violations is the number of pairs of papers $i \neq j$ failing EF1. There are $n^2 - n$ total potential violations.

GRRR is the only algorithm to achieve EF1 on all three conferences. FairFlow, FairIR, and TPMS have very high levels of EF1 violation on CVPR, and all four baselines have some EF1 violations on CVPR and CVPR'18. Although some EF1 violations may be permissible, the large number of violations by FairFlow, FairIR, and TPMS on CVPR implies that many papers received unnecessarily imbalanced assignments relative to other papers.

GRRR consistently outperforms FairFlow on welfare measures across all three conferences, though it is outperformed by FairIR, TPMS, and PR4A. Although GRRR obtains a minimum paper score of 0 on CVPR (as does TPMS), only three papers receive a score of 0. We could solve this by increasing the assignment limits of a few reviewers or allowing a small number of EF1 violations for these three papers.

We estimate α and γ on all three datasets. γ is rather large for CVPR and CVPR'18, but our empirical performance demonstrates the value of our approach. It is also quite possible that other conferences or other application areas would yield welfare functions that are closer to monotonic and submodular, leading to lower values of γ.

Remark 4.1 (Fairness to Reviewers)**.** We take the position throughout this paper that is more appropriate to treat reviewers, rather than papers, as goods. Paper reviewing is generally viewed as a chore, not a benefit (the current work being an obvious exception), whereas papers do benefit from appropriate reviews. A more comprehensive treatment of fairness to both reviewers and papers would require a completely novel approach and is out of the scope of this paper. Still we must verify that GRRR is at least as fair to reviewers as our baselines. For each conference, we compute the distribution of reviewing loads for all algorithms. Our method is relatively consistent with the baselines, and does not introduce a large unfairness in reviewing load. We also test the variant of RRR that incorporates reviewer load lower bounds of $l_r = 2$ for all r (same bounds used by Kobren *et al.* [2019]). On all three conferences, this variant terminates with complete allocations satisfying reviewer lower and upper bounds, while maintaining EF1 guarantees and competitive USW.

To conclude, GRRR ensures EF1, while remaining highly competitive on welfare and other fairness guarantees. Moreover, GRRR is easy to implement and understand, and its simple formulation gives it the flexibility to handle many additional constraints (see full version for multiple variants).

5 Conclusion

The reviewer assignment problem provides interesting constraints, complicating the standard fair allocation setting. We demonstrate that a straightforward round-robin allocation, combined with a novel optimization technique on paper orders, finds EF1 allocations with high USW in the reviewer assignment setting. Our approach of optimizing over *orders* for round-robin allocations is of independent interest, and may inspire further study of optimal round-robin allocations. 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.

Acknowledgements

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 Coop AI at NeurIPS 2021, WINE 2021, AAMAS 2022, GAIW at AAMAS 2022, and IJCAI 2022. This work was performed in part using high performance computing equipment obtained under a grant from the Collaborative R&D Fund managed by the Massachusetts Technology Collaborative.

References

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

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

[Aziz *et al.*, 2019] Haris Aziz, Xin Huang, Nicholas Mattei, and Erel Segal-Halevi. The constrained round robin algorithm for fair and efficient allocation. *arXiv preprint arXiv:1908.00161*, 2019.

[Babaioff *et al.*, 2021] Moshe Babaioff, Tomer Ezra, and Uriel Feige. Fair and truthful mechanisms for dichotomous valuations. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5119–5126, 2021.

[3]https://github.com/justinpayan/ReviewerAssignmentCode

[Barman *et al.*, 2019] Siddharth Barman, Ganesh Ghalme, Shweta Jain, Pooja Kulkarni, and Shivika Narang. Fair division of indivisible goods among strategic agents. In *Proceedings of the 18th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1811–1813, 2019.

[Benabbou *et al.*, 2020] Nawal Benabbou, Mithun Chakraborty, Ayumi Igarashi, and Yair Zick. Finding fair and efficient allocations when valuations don't add up. In *Proceedings of the 13th International Symposium on Algorithmic Game Theory (SAGT)*, pages 32–46, 2020.

[Biswas and Barman, 2018] Arpita Biswas and Siddharth Barman. Fair division under cardinality constraints. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 91–97, 2018.

[Bouveret and Lang, 2011] Sylvain Bouveret and Jérôme Lang. A general elicitation-free protocol for allocating indivisible goods. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 73–78, 2011.

[Bouveret *et al.*, 2016] Sylvain Bouveret, Yann Chevaleyre, and Nicolas Maudet. Fair allocation of indivisible goods. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia, editors, *Handbook of computational social choice*, chapter 12, pages 284–309. Cambridge University Press, 2016.

[Buchbinder *et al.*, 2012] Niv Buchbinder, Moran Feldman, Joseph Naor, and Roy Schwartz. A tight linear time (1/2)-approximation for unconstrained submodular maximization. In *Proceedings of the 53rd Symposium on Foundations of Computer Science (FOCS)*, pages 649–658, 2012.

[Budish, 2011] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119:1061–1061, 12 2011.

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

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

[Das and Kempe, 2011] Abhimanyu Das and David Kempe. Submodular meets spectral: greedy algorithms for subset selection, sparse approximation and dictionary selection. In *Proceedings of the 28th International Conference on Machine Learning (ICML)*, pages 1057–1064, 2011.

[Dror *et al.*, 2021] Amitay Dror, Michal Feldman, and Erel Segal-Halevi. On fair division under heterogeneous matroid constraints. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5312–5320, 2021.

[Fisher *et al.*, 1978] Marshall Fisher, George Nemhauser, and Laurence Wolsey. An analysis of approximations for maximizing submodular set functions—II. In M. L. Balinski and A. J. Hoffman, editors, *Polyhedral Combinatorics*, pages 73–87. Springer Berlin Heidelberg, Berlin, Heidelberg, 1978.

[Gölz and Procaccia, 2019] Paul Gölz and Ariel D Procaccia. Migration as submodular optimization. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 549–556, 2019.

[Kalinowski *et al.*, 2013] Thomas Kalinowski, Nina Narodytska, and Toby Walsh. A social welfare optimal sequential allocation procedure. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 227–233, 2013.

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

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

[Mulligan *et al.*, 2013] Adrian Mulligan, Louise Hall, and Ellen Raphael. Peer review in a changing world: An international study measuring the attitudes of researchers. *Journal of the American Society for Information Science and Technology*, 64(1):132–161, 2013.

[Oxley, 2011] James Oxley. *Matroid Theory*. Oxford Graduate Texts in Mathematics. Oxford Univerity Press, 2nd edition, 2011.

[Payan and Zick, 2021] Justin Payan and Yair Zick. I will have order! Optimizing orders for fair reviewer assignment. *arXiv:2108.02126*, 2021.

[Shah, 2019] Nihar B Shah. Principled methods to improve peer review. Technical report, Carnegie Mellon University, 2019.

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

[Stelmakh, 2021] Ivan Stelmakh. Towards fair, equitable, and efficient peer review. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, pages 15736–15737, 2021.