

Proportional Rankings

Piotr Skowron
 TU Berlin, Germany
 p.k.skowron@gmail.com

Martin Lackner
 University of Oxford, UK
 martin.lackner@cs.ox.ac.uk

Markus Brill
 TU Berlin, Germany
 markus.brill@campus.tu-berlin.de

Dominik Peters
 University of Oxford, UK
 dominik.peters@cs.ox.ac.uk

Edith Elkind
 University of Oxford, UK
 edith.elkind@cs.ox.ac.uk

Abstract

We extend the principle of proportional representation to rankings: given approval preferences, we aim to generate aggregate rankings so that cohesive groups of voters are represented proportionally in each initial segment of the ranking. Such rankings are desirable in situations where initial segments of different lengths may be relevant, e.g., in recommender systems, for hiring decisions, or for the presentation of competing proposals on a liquid democracy platform. We define what it means for rankings to be proportional, provide bounds for well-known aggregation rules, and experimentally evaluate the performance of these rules.

1 Introduction

Consider a population with dichotomous (approval) preferences over a set of 300 alternatives. Assume that 50% of the population approves of the first 100 alternatives, 30% of the next 100 alternatives, and 20% of the last 100. Suppose that we want to obtain a ranking of the alternatives that reflects the preferences of the population. Perhaps the most straightforward approach is to use *Approval Voting (AV)*, which ranks the alternatives from the most frequently approved to the least frequently approved. A striking property of the resulting ranking is that half of the population does not approve any of the first 100 alternatives in the ranking.

In many scenarios, rankings produced by Approval Voting are highly unsatisfactory; rather, it is desirable to interleave alternatives supported by different (sufficiently large) groups. For instance, consider an anonymous user running a search engine to find results for the query “Armstrong”. If 50% of the users performing this search would like to see results for Neil Armstrong, 30% for Lance Armstrong, and 20% for Louis Armstrong, it is not desirable to only place results referring to Neil Armstrong in the top part of the ranking; rather, results related to each of the Armstrongs should be displayed in appropriately high positions.

There are numerous applications where such diversity within collective rankings is desirable. For instance, consider recommendation systems which aim at accommodating different types of users (the estimated preferences of these user types should be represented proportionally to their likelihood),

a human-resource department providing hiring recommendations when it is unclear how many positions are to be filled, or committee elections with some additional structure (e.g., when we want to elect a committee with a chairman and vice-chairman, as well as substitute members).

Another application in which diversity in rankings is relevant is *liquid democracy* [Behrens *et al.*, 2014]. A defining feature of liquid democracy is that all participants are allowed (and encouraged) to contribute to the decision-making process. This may lead to situations where a very large number of proposals need to be considered. Since it cannot be expected that every participant studies all available alternatives before making a decision, the order in which competing alternatives are presented plays a crucial role [Behrens *et al.*, 2014, Chapter 4.10]. In particular, the LIQUIDFEEDBACK system [Behrens *et al.*, 2014], used by several political parties and other organisations, implements one of the rules (reverse SeqPAV) that we analyze in this paper. Thus, obtaining a better understanding of proportionality in rankings has direct relevance to real-world applications.

The idea of rankings that are diversified with respect to the preferences of a population appears in several streams of research. For example, in the context of search engines, this aim is often referred to as *diversifying search results* [Welch *et al.*, 2011; Santos *et al.*, 2015; Kingrani *et al.*, 2015; Wang *et al.*, 2016]. There are also models which incorporate this idea into online advertising (see the work of Hu *et al.* [2011], and the references therein). In the context of liquid democracy, Behrens *et al.* observe that using AV gives rise to what they call the “noisy minorities” problem: relatively small groups of very active participants can “flood” the system with their contributions, creating the impression that their opinion is much more popular than it actually is. This is problematic insofar as other alternatives (that are potentially much more popular) run the risk of being “buried” and not getting sufficient exposure. Behrens *et al.* suggest that, in order to prevent this problem, the ranking mechanism needs to ensure that the order adequately reflects the opinions of the participants.

In this paper, we propose an abstract model applicable to all these scenarios, and initiate a formal study of the problem of finding a proportional collective ranking. In our study, we use tools from political and social sciences, and in particular we adapt the concept of *proportional representation* [Balinski and Young, 1982; Monroe, 1995] to the case of rankings.

Informally, proportional representation requires that the extent to which a particular preference or opinion is represented in the outcome should be proportional to the frequency with which this preference or opinion occurs within the population. For instance, proportional representation is often a requirement in the context of *parliamentary elections*, where candidates are grouped into political parties and voters express preferences over parties. If a party receives, say, 20% of the votes, then proportional representation requires that this party should be allocated (roughly) 20% of the parliamentary seats (see the work of Laslier, 2012, for a discussion of arguments for and against proportional representation in a political context).

The concept of proportional representation can be extended to the case of rankings in a very natural way. Informally, we say that a ranking τ is proportional if each prefix of τ , viewed as a subset of alternatives, satisfies (some form of) proportional representation. In more detail, inspired by the concepts of *Extended* or *Proportional Justified Representation* that have been recently introduced in the context of multiwinner elections [Aziz et al., 2017; Sánchez-Fernández et al., 2017], we require that, for each “sufficiently large” group of voters with consistent preferences, a proportional number of alternatives approved by this group is placed appropriately highly in the ranking. The position of such alternatives in the ranking depends on the level of their support, indicated by the size and the cohesiveness of the respective group of voters.

Our Contribution. (i) We formalize the concept of proportionality of a ranking and introduce a quantitative way of measuring it, (ii) we observe that several known multiwinner rules satisfying *committee monotonicity* can be viewed as ranking rules, (iii) we provide theoretical bounds on the quality of proportional representation of several ranking rules, and (iv) we experimentally evaluate ranking rules.

Related Work. Proportional representation is traditionally studied in scenarios where a *subset* of alternatives (such as a parliament or a committee) needs to be chosen (see the survey by Faliszewski et al. [2017]). The setting that is most commonly studied is that of *closed party lists*, where alternatives are grouped into pairwise disjoint parties and voters are restricted to selecting a single party [Gallagher, 1991]. In this setting, providing proportional representation reduces to solving an *apportionment* problem [Balinski and Young, 1982; Pukelsheim, 2014; Brill et al., 2017b]. In an influential paper, Monroe [1995] generalized the concept of proportional representation to settings where voters’ preferences are given as rank-orderings of the set of alternatives. For approval preferences, concepts capturing proportional representation have recently been studied by Aziz et al. [2017] and Sánchez-Fernández et al. [2017].

The setting considered in this paper differs from all of the above settings in that we are interested in *ranking* the alternatives, rather than choosing a subset of them. To the best of our knowledge, proportional rankings based on approval preferences have not been considered in the literature. In the context of linear (i.e., rank-order) preferences, proportional rankings are discussed by Schulze [2011]. However, this paper neither

proposes a measure of proportionality, nor does it compare the proportionality provided by different rules.

An extended version of this paper, including all proofs, can be found on arXiv [Skowron et al., 2016].

2 Preliminaries

For $s \in \mathbb{N}$, we write $[s] = \{1, \dots, s\}$. For each set X , we let $\mathcal{S}(X)$ and $\mathcal{S}_k(X)$ denote the set of all subsets of X , and the set of all k -element subsets of X , respectively.

Let $N = [n]$ be a finite set of *voters* and A a finite set of m *alternatives*. Each voter $i \in N$ approves a nonempty subset of alternatives $A_i \subseteq A$. For each $a \in A$, we write N_a for the set of voters who approve a , i.e., $N_a = \{i \in N : a \in A_i\}$. We refer to $|N_a|$ as the *approval score* of a . A list $P = (A_1, \dots, A_n)$ of approval sets, one for each voter $i \in N$, is called a *profile* on (A, N) .

A *ranking* is a linear order over A . For a ranking r and for $k \in [m]$, we denote the k -th element in r by r_k . Thus, r can be represented by the list (r_1, r_2, \dots, r_m) . Given a ranking $r = (r_1, \dots, r_m)$ and $k \in [m]$, we let $r_{\leq k} = \{r_1, \dots, r_k\}$ denote the subset of A consisting of the top k elements according to r . An (*approval-based*) *ranking rule* f maps a profile P on (A, N) to a ranking $f(P)$ over A .

In what follows, we consider several ranking rules that can be obtained by adapting existing multiwinner rules. An (*approval-based*) *multiwinner rule* takes as input a profile P on (A, N) and an integer $k \in [m]$ and outputs a k -element subset of A , referred to as the *winning committee*. A generic adaptation of multiwinner rules to ranking rules is possible whenever the respective multiwinner rule \mathcal{R} has the property that $\mathcal{R}(P, k-1) \subseteq \mathcal{R}(P, k)$ for all $k \leq m$ (this property is known as *committee monotonicity* or *house monotonicity*): if this is the case, we can produce a ranking by placing the unique alternative in $\mathcal{R}(P, k) \setminus \mathcal{R}(P, k-1)$ in position k .

Fix a profile P . We consider the following ranking rules:¹

Approval Voting (AV). AV ranks the alternatives in order of their approval score, so that $|N_{r_1}| \geq \dots \geq |N_{r_m}|$.

Sequential Thiele Rules. This family of rules has been proposed by Danish polymath Thorvald N. Thiele [1895] (and is also known as Reweighted Approval Voting). It is parameterized by *weight vectors*, i.e., sequences (w_1, w_2, \dots) of non-negative reals. For a given weight vector $\mathbf{w} = (w_1, w_2, \dots)$ and a subset $S \subseteq A$ of alternatives, define the *w-Thiele score* of S as $w(S) = \sum_{i \in N} \sum_{j=1}^{|A_i \cap S|} w_j$. Usually, it is assumed that $w_1 \geq w_2 \geq \dots$; the intuition behind this constraint is that voters prefer sets S that contain more of their approved alternatives, but there are decreasing marginal returns to adding further such alternatives to S . The sequential w-Thiele rule constructs a ranking iteratively, starting with the empty partial ranking $r = ()$. In step $k \in [m]$, it appends to r an alternative a with maximum *marginal contribution* $w(r_{\leq k-1} \cup \{a\}) - w(r_{\leq k-1})$ among all yet unranked alternatives. Many interesting rules belong to this family for suitable

¹All rules that we describe may have to break ties at some point in the execution; we adopt an adversarial approach to tie-breaking, i.e., we say that a ranking rule satisfies a property only if it satisfies it for all possible ways of breaking ties.

weight vectors \mathbf{w} . For example, AV can be viewed as the sequential \mathbf{w} -Thiele with $\mathbf{w} = (1, 1, 1, \dots)$, *Sequential Proportional Approval Voting (SeqPAV)* is defined by $\mathbf{w}_{\text{PAV}} = (1, 1/2, 1/3, \dots)$, *Greedy Chamberlin–Courant (Greedy CC)* is defined by $\mathbf{w} = (1, 0, 0, \dots)$, and for every $p > 1$, the p -geometric rule is defined by $\mathbf{w} = (1/p, 1/p^2, 1/p^3, \dots)$.

Reverse SeqPAV. This rule is a bottom-up variant of SeqPAV; it has been introduced by Thiele [1895] (in the multiwinner setting) and was independently proposed by Behrens *et al.* [2014] under the name *harmonic weighting*. Initially it sets $S = A$ and $r = ()$. Then, at each step it picks an alternative a minimizing $w_{\text{PAV}}(S) - w_{\text{PAV}}(S \setminus \{a\})$, removes it from S and prepends it to r , i.e., the rule builds a ranking starting with the lowest-ranked alternative.

Phragmén’s Rule. Swedish mathematician Edvard Phragmén [1895] proposed a committee selection rule that can be phrased as a load balancing procedure: every alternative incurs a *load* of one unit, and the load of alternative a has to be distributed among all voters in N_a . Phragmén’s rule constructs a ranking iteratively, starting with the empty partial ranking $r = ()$. Initially, the load of each voter is 0. At each step, the rule picks a yet unranked alternative and distributes its associated load of 1 over the voters who approve it; the alternative and the load distribution scheme are chosen so as to minimize the maximum load across all voters. This alternative is then appended to the ranking r . (For details, see the work of Mora and Oliver [2015], Janson [2016], and Brill *et al.* [2017a]).

3 Measures of Proportionality

In this section, we define a measure of proportionality for rankings and then extend it to ranking rules. In what follows, let P be a profile on (A, N) with $|A| = m$ and $|N| = n$.

Given a group of voters $N' \subseteq N$ and a set of alternatives $S \subseteq A$, a natural measure of the group’s “satisfaction” provided by S is the average number of alternatives in S that are approved by a voter in N' . Thus, we define

$$\text{avg}(N', S) = \frac{1}{|N'|} \sum_{i \in N'} |A_i \cap S|.$$

We refer to $\text{avg}(N', S)$ as the *average representation of N' with respect to S* . To extend this idea to rankings, we consider the case where the subset S is an initial segment of a given ranking r , i.e., $S = r_{\leq k}$ for some $k \in [m]$. Intuitively, every group N' wants to have an average representation $\text{avg}(N', r_{\leq k})$ that is as large as possible, for all $k \in [m]$. Now, in our model, whether a group of voters deserves to be represented in the top positions of a ranking depends on two parameters: its relative size and its cohesiveness, i.e., the number of alternatives that are unanimously approved by the group members. This motivates the following definition.

Definition 1. (Significant group) Fix a profile P on (A, N) . The proportion of a group $N' \subseteq N$ is $\alpha(N') = |N'|/|N|$, and the cohesiveness of N' is given by $\lambda(N') = |\bigcap_{i \in N'} A_i|$. Given $\alpha \in (0, 1]$ and $\lambda \in [m]$, we say that a group N' is (α, λ) -significant in P if $|N'| = \lceil \alpha n \rceil$ and $\lambda(N') \geq \lambda$.

The following definition captures the intuitively compelling idea that a group can demand to be represented in the top positions of the ranking in proportion to its significance, as long as the number of demanded alternatives does not exceed the cohesiveness of the group.

Definition 2. (Justifiable demand) The justifiable demand of a group $N' \subseteq N$ with respect to the top k positions of a ranking is defined as

$$\text{jd}(N', k) = \min(\lfloor \alpha(N') \cdot k \rfloor, \lambda(N')).$$

For example, if a group contains 25% of the voters and has a cohesiveness of 3, it has a justifiable demand of 1 with respect to the top four positions, a justifiable demand of 2 with respect to the top eight positions, and a justifiable demand of 3 with respect to the top twelve positions, which is also its maximum justifiable demand.

It would be desirable to find a rule that provides every group with an average representation that meets the group’s justifiable demand. However, the following example shows that this is not always possible.

Example 1. Let $A = \{a, b, c\}$ and $n = 6$ and consider the profile P given by $A_1 = \{a\}$, $A_2 = \{a, b\}$, $A_3 = \{b\}$, $A_4 = \{b, c\}$, $A_5 = \{c\}$, and $A_6 = \{a, c\}$. Consider the ranking $r = (a, b, c)$. The group $N' = \{4, 5, 6\}$ has $\alpha(N') = 1/2$, $\lambda(N') = 1$. Therefore, its justifiable demand with respect to the top two positions is $\text{jd}(N', 2) = \min(\lfloor 1/2 \cdot 2 \rfloor, 1) = 1$. However, its average representation with respect to the top two positions of r is only $\text{avg}(N', \{a, b\}) = 2/3$. Since P is completely symmetric with respect to the alternatives, we can find such a group for every other ranking as well.

Example 1 shows that it may not be feasible to provide each group with the level of representation that meets its justifiable demand, but it might be possible to guarantee a large fraction of it. For instance, in Example 1 one can ensure that $\text{avg}(N', k) \geq 2/3 \cdot \text{jd}(N', k)$ for all groups N' and for all $k \leq 3$. This observation leads to the following definition.

Definition 3. (Optimal ranking) We define the quality of a ranking r for a profile P as

$$q_P(r) = \min_{\substack{k \in [m], N' \subseteq N: \\ \text{jd}(N', k) > 0}} \frac{\text{avg}(N', r_{\leq k})}{\text{jd}(N', k)}.$$

An optimal ranking for P is a ranking in $\arg \max_r q_P(r)$.

We observe that $q_P(r)$ is well-defined, i.e., there is always a pair (N', k) with $\text{jd}(N', k) > 0$. Indeed, take an alternative a with maximum approval score: as $A_i \neq \emptyset$ for all $i \in N$, we have $|N_a| \geq n/m$ by the pigeonhole principle. Further, $\lambda(N_a) \geq 1$, so $\text{jd}(N_a, m) \geq \min(\lfloor \frac{n}{m} \cdot m \rfloor, 1) = 1$.

Example 2. Consider the profile from Example 1 and the ranking $r = (a, b, c)$. For the first position in this ranking (i.e., for $k = 1$), there exists no group of voters with a positive justifiable demand, since such a group would need to consist of all the voters who all approve at least one common alternative. For $k = 2$, the justifiable demand of three groups of voters is equal to 1. These groups are $N_1 = \{1, 2, 6\}$, $N_2 = \{2, 3, 4\}$, and $N_3 = \{4, 5, 6\}$, with commonly approved alternatives a , b , and c , respectively. The average representation of groups N_1 ,

N_2 , and N_3 is equal to $4/3$, $4/3$, and $2/3$, respectively. There is no group with justifiable demand of 2 with respect to the top two positions. Finally, for $k = 3$ there are 9 groups with justifiable demand of 1 (these are $\{1, 2\}$, $\{1, 6\}$, $\{2, 3\}$, $\{2, 4\}$, $\{2, 6\}$, $\{3, 4\}$, $\{4, 5\}$, $\{4, 6\}$, and $\{5, 6\}$); each such group has an average representation of at least one. There is no group with a justifiable demand of two—such a group would need to contain four voters, who all approve at least two common alternatives—and, clearly, no group with a justifiable demand of three. Thus, the quality of our ranking is $q_P(r) = 2/3$, and this is an optimal ranking.

Unfortunately, good rankings are hard to compute.

Theorem 1. *Given a profile P , it is NP-hard to decide whether there exists a ranking r with $q_P(r) \geq 1$.*

Proof. We give a reduction from VERTEX COVER. An instance of this problem is given by a graph $G = (V, E)$ and an integer ℓ . It is a ‘yes’-instance if there is a subset of ℓ vertices $S \subseteq V$ such that each edge in E contains some vertex in S . We can assume that G is 3-regular, as VERTEX COVER remains NP-hard in this case [Garey and Johnson, 1979].

Given a vertex cover instance (G, ℓ) with $G = (V, E)$, we construct an instance of the problem of finding an optimal ranking as follows. Since the degree of each vertex is exactly 3, we can assume that $3\ell \geq |E|$: instances with $3\ell < |E|$ are trivially ‘no’-instances. We set $A = V \cup \{d\}$, where d is a dummy alternative. For every edge $\{a, b\} \in E$ we create a voter that approves $\{a, b\}$. We also add $3\ell + 3 - |E| \geq 3$ dummy voters who approve d only. Consequently, $n = |N| = 3\ell + 3$. Since the graph is 3-regular, we have $m = |A| = 2|E|/3 + 1 \leq 3\ell$. This defines the profile P . It can be shown that G has a vertex cover of size ℓ if and only if there exists a ranking r with $q_P(r) \geq 1$, i.e., $\text{avg}(N', r_{\leq k}) \geq \text{jd}(N', k)$ for all $N' \subseteq N$ and $k \in [m]$. For details, see the full version of the paper [Skowron *et al.*, 2016]. \square

Further, a fairly straightforward reduction from the Maximum k -Subset Intersection problem [Xavier, 2012] shows that even the problem of deciding whether there exists an (α, λ) -significant group of voters is NP-hard.

Proposition 1. *The problem of deciding whether there exists an (α, λ) -significant group of voters is NP-complete.*

The measure $q_P(r)$ can be lifted from individual rankings to ranking rules: we can measure the quality of a ranking rule f as the minimum value of $q_P(f(P))$, over all possible profiles P . However, in our theoretical analysis of ranking rules we take the following approach, which assumes more flexibility on the part of groups of voters that seek to be represented: given a group of proportion α and with cohesiveness at least λ , we ask at which point in the ranking the average satisfaction of this group reaches λ . This guarantee is given by the function $\kappa(\alpha, \lambda)$, which is formally defined as follows.

Definition 4. (κ -group representation) *Let $\kappa(\alpha, \lambda)$ be a function from $((0, 1] \cap \mathbb{Q}) \times \mathbb{N}$ to \mathbb{N} . A ranking r provides κ -group representation (κ -GR) for profile P if for all rational $\alpha \in (0, 1]$, all $\lambda \in \mathbb{N}$, and all voter groups $N' \subseteq N$ that are (α, λ) -significant in P it holds that*

$$\text{avg}(N', r_{\leq \kappa(\alpha, \lambda)}) \geq \lambda.$$

A ranking rule f satisfies κ -group representation (κ -GR) if $f(P)$ provides κ -group representation for every profile P .

Example 3. *Consider again the profile from Example 1 and the ranking $r = (a, b, c)$. For the κ -group representation of this ranking we have $\kappa(1/2, 1) \geq 3$. To establish this, we consider groups consisting of half of the voters who jointly approve at least one alternative, and look for the smallest value of k such that each such group on average approves at least one from the the top k alternatives in r .*

Let us now explore the differences between the two measures, the worst-case quality of the ranking $q_P(f(P))$, and κ -group representation. While $q_P(f(P))$ is just a single number, κ -group representation carries more information. In particular, it makes it possible to express the fact that some groups are better represented in the top parts of the resulting ranking than further below (this can be captured by making κ a convex function of λ). As a more concrete example, assume that $q_P(f(P)) < 1/2$ and consider a group N' that is $(1/10, m/2)$ -significant. Since the justified demand of this group is upper-bounded by $m/10$, $q_P(f(P))$ does not say anything about how far we need to go down the ranking to obtain an average representation greater than $\lambda > m/20$ (the justifiable demand of N' is at most $m/10$ and f guarantees only half of it), while κ -group representation provides such information for each $\lambda \in [1, m/2]$.

In the next section, we investigate guarantees in terms of κ -group representation provided by the ranking rules introduced in Section 2.

4 Theoretical Guarantees for Ranking Rules

We start our analysis with the simplest of our rules, namely, Approval Voting (AV). Interestingly, AV provides very good guarantees to majorities, that is, groups with $\alpha(N') > 1/2$. For example, for $\alpha(N') = 0.8$ it guarantees an average satisfaction of λ within the top $\lceil 1.5\lambda \rceil$ positions. However, for groups with $\alpha(N') \leq 1/2$ it provides no guarantee at all.

Theorem 2. *For rational $\alpha > 1/2$, AV satisfies $\kappa(\alpha, \lambda)$ -group representation for $\kappa(\alpha, \lambda) = \lceil \frac{\lambda\alpha}{2\alpha-1} \rceil$, but fails it for $\kappa(\alpha, \lambda) = \lceil \frac{\lambda\alpha}{2\alpha-1} \rceil - 1$. However, for $\alpha < 1/2$, AV does not satisfy $\kappa(\alpha, \lambda)$ -group representation for any function $\kappa(\alpha, \lambda)$.*

Proof. We first prove that for $\alpha > 1/2$, AV satisfies $\kappa(\alpha, \lambda)$ -group representation for $\kappa(\alpha, \lambda) = \lceil \frac{\lambda\alpha}{2\alpha-1} \rceil$. Consider a group of voters $N' \subseteq N$ with $n' = |N'| = \lceil \alpha n \rceil$, $\lambda(N') \geq \lambda$. Let $k = \lceil \frac{\lambda\alpha}{2\alpha-1} \rceil$, and let r be the ranking returned by AV. Assume for the sake of contradiction that $\text{avg}(N', r_{\leq k}) < \lambda$. Then there exists some alternative h that is approved by all voters in N' , but does not appear in the top k positions of r . Thus, each alternative in $r_{\leq k}$ is approved by at least $\lceil \alpha n \rceil$ voters and hence by at least $2\lceil \alpha n \rceil - n$ voters in N' . Hence,

$$\text{avg}(N', r_{\leq k}) \geq k \cdot \frac{2\lceil \alpha n \rceil - n}{\lceil \alpha n \rceil} \geq k \cdot \frac{2\alpha n - n}{\alpha n} \geq \lambda.$$

This contradiction proves our first claim.

Now, we show that AV does not satisfy $\kappa(\alpha, \lambda)$ -group representation for $\kappa(\alpha, \lambda) = \lceil \frac{\lambda\alpha}{2\alpha-1} \rceil - 1$. Fix $\alpha \in (0, 1] \cap \mathbb{Q}$ and $\lambda \in \mathbb{N}$ and let $k = \lceil \frac{\lambda\alpha}{2\alpha-1} \rceil - 1 < \frac{\lambda\alpha}{2\alpha-1}$. As $\alpha \in \mathbb{Q}$, there

exist integers x and n such that $\alpha = x/n$. Let $A = A' \cup A''$, where $|A'| = |A''| = k$ and $A' \cap A'' = \emptyset$. Consider a profile on (A, N) that contains two groups of voters G and F with sizes $|G| = |F| = x = \alpha n$, such that G and F have the smallest possible intersection. That is, for $\alpha \leq 1/2$ the sets G and F are disjoint, and for $\alpha > 1/2$ we have $|G \cap F| = (2\alpha - 1)n$. Suppose that each voter in G approves all alternatives in A' and each voter in F approves all alternatives in A'' . AV may rank A' in the first k positions. Thus,

$$\text{avg}(F, r_{\leq k}) = k \cdot \frac{(2\alpha - 1)n}{\alpha n} < \lambda.$$

Finally, let us prove that for $\alpha \leq 1/2$, AV does not satisfy $\kappa(\alpha, \lambda)$ -group representation for any function $\kappa(\alpha, \lambda)$. For the sake of contradiction let us assume that this is not the case. Let us fix $\alpha \in (0, 1]$, $\lambda \in \mathbb{N}$, and let $k = \kappa(\alpha, \lambda)$. Using the same idea as in the previous paragraph, we obtain an instance where there is a set of voters N' with $|N'| = \alpha n$, $\lambda(N') = |N'| \geq \lambda$ such that no voter in N' approves any of the alternatives appearing in the top k positions of the ranking returned by AV. This gives a contradiction. \square

In contrast, Phragmén's rule, SeqPAV and p -geometric rules give reasonable guarantees for *all* values of α and λ .

Theorem 3. *SeqPAV satisfies $\kappa(\alpha, \lambda)$ -group representation for $\kappa(\alpha, \lambda) = \lceil \frac{2(\lambda+1)^2}{\alpha^2} \rceil$.*

Proof. Fix $\alpha \in (0, 1]$, $\lambda \in \mathbb{N}$, a profile P , and a group of voters $N' \subseteq N$ such that $n' = |N'| = \lceil \alpha n \rceil$ and $\lambda(N') \geq \lambda$. Set $k = \lceil \frac{2(\lambda+1)^2}{\alpha^2} \rceil$. Let r be the ranking returned by SeqPAV.

Assume for the sake of contradiction that $\text{avg}(N', r_{\leq k}) < \lambda$ and set $z = n' \cdot \text{avg}(N', r_{\leq k})$; note that $z < \lambda n'$. By our assumption, at the end of each step $t \in [k]$ there exists some alternative that is approved by all voters in N' , but has not been ranked yet; let h be some such alternative. For $t \in [k+1]$, consider the moment just before the t -th step of SeqPAV. Let $a_i(t)$ denote the number of alternatives selected so far that appear in A_i , and let $T(t) = \sum_{i \in N} \frac{1}{a_i(t)+1}$. Note that $n = T(1) \geq T(2) \geq \dots \geq T(k+1) \geq 0$. For each $t \in [k+1]$ we have $\sum_{i \in N'} a_i(t) \leq z$. Using the arithmetic mean–harmonic mean inequality, we infer

$$\frac{z + n'}{n'} \geq \frac{\sum_{i \in N'} (a_i(t) + 1)}{n'} \geq \frac{n'}{\sum_{i \in N'} \frac{1}{a_i(t)+1}},$$

that is,

$$\sum_{i \in N'} \frac{1}{a_i(t) + 1} \geq \frac{(n')^2}{z + n'} > \frac{(n')^2}{n'(\lambda + 1)} = \frac{n'}{\lambda + 1}.$$

Consider the alternative a selected by SeqPAV at step t , $t \in [k]$. Without loss of generality, assume that $N_a = \{1, \dots, s\}$. Since alternative h is available at this step, it has to be the case that SeqPAV favors a over h , i.e., for the harmonic weight vector $\mathbf{w} = (1, 1/2, 1/3, \dots)$ we have $w(r_{\leq t-1} \cup \{a\}) - w(r_{\leq t-1}) \geq w(r_{\leq t-1} \cup \{h\}) - w(r_{\leq t-1})$. This implies

$$\sum_{i=1}^s \frac{1}{a_i(t) + 1} \geq \sum_{i \in N'} \frac{1}{a_i(t) + 1} > \frac{n'}{\lambda + 1}.$$

Thus, we have

$$\begin{aligned} T(t) - T(t+1) &= \sum_{i=1}^s \left(\frac{1}{a_i(t) + 1} - \frac{1}{a_i(t) + 2} \right) \\ &= \sum_{i=1}^s \frac{1}{(a_i(t) + 1)(a_i(t) + 2)} \geq \sum_{i=1}^s \frac{1}{2(a_i(t) + 1)^2}. \end{aligned}$$

Applying the Cauchy–Schwarz inequality to $(1, \dots, 1)$ and $(\frac{1}{a_1(t)+1}, \dots, \frac{1}{a_s(t)+1})$, we obtain

$$\begin{aligned} \sum_{i=1}^s \frac{1}{2(a_i(t) + 1)^2} &\geq \frac{1}{2s} \left(\sum_{i=1}^s \frac{1}{a_i(t) + 1} \right)^2 \\ &> \frac{1}{2n} \left(\frac{n'}{\lambda + 1} \right)^2 \geq \frac{n\alpha^2}{2(\lambda + 1)^2}. \end{aligned}$$

Since the above inequality holds for each $t \in [k]$, we have

$$T(1) - T(k+1) = \sum_{t=1}^k (T(t) - T(t+1)) > \frac{kn\alpha^2}{2(\lambda + 1)^2}.$$

Since $\frac{kn\alpha^2}{2(\lambda+1)^2} \geq n$, we have $T(k+1) < T(1) - n = n - n = 0$, a contradiction. This completes the proof. \square

The technique developed in the proof of Theorem 3 can be used to provide similar bounds for other sequential Thiele rules. For the p -geometric rule we obtain the following bound.

Theorem 4. *For each $p > 1$, the p -geometric rule satisfies $\kappa(\alpha, \lambda)$ -group representation for $\kappa(\alpha, \lambda) = \lceil \frac{p^{\lambda+1}}{\alpha(p-1)} \rceil$.*

A much more involved analysis using a potential function argument gives us a bound for Phragmén's rule.

Theorem 5. *Phragmén's rule satisfies $\kappa(\alpha, \lambda)$ -group representation for $\kappa(\alpha, \lambda) = \lceil \frac{5\lambda}{\alpha^2} + \frac{1}{\alpha} \rceil$.*

Theorem 4 establishes a linear relationship between the proportion of the group α and the guarantee $\kappa(\alpha, \lambda)$. Thus, our bound for the p -geometric rule is better than our bounds for SeqPAV and Phragmén's rule from Theorems 3 and 5, respectively. Further, as suggested by Example 1, a linear relationship is the best we can hope for. Unfortunately, as a tradeoff we obtain an exponential relationship between the required amount of representation λ and the guarantee $\kappa(\alpha, \lambda)$. Nevertheless, Theorem 4 shows that if we are only interested in optimizing $\kappa(\alpha, \lambda)$ for a *constant* value of $\lambda > 1$ (e.g., if we know that no agent desires more than λ objects in total), then the p -geometric rule provides very good guarantees for groups of different sizes. Furthermore, it can be shown that by taking $p = \lambda/(\lambda + 1)$ we get the best possible guarantee in this setting (i.e., for a fixed λ).

For Reverse SeqPAV, we have not been able to establish an analogous lower bound. However, we can obtain a bound on $\kappa(\alpha, \lambda)$ for each $\alpha \in (0, 1]$ and for *some* sufficiently large λ .

Theorem 6. *Let $\alpha \in (0, 1]$, $\lambda \in \mathbb{N}$, and let $N' \subseteq N$ be an (α, λ) -significant group. Let r be the ranking returned by Reverse SeqPAV. Then, there exists $y \geq \lambda$ such that $\text{avg}(N', r_{\leq y/\alpha}) \geq y$.*

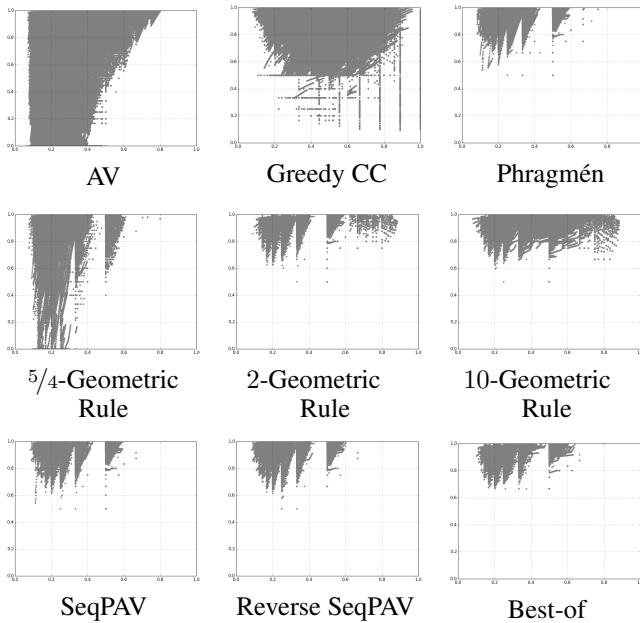


Figure 1: Violations of justified demand encountered in our data sets. The x -axis shows the proportion of the respective group of voters, $\alpha(N')$, and the y -axis shows the quotient of average representation and justified demand, $y_r(N', k)$. Roughly, ranking rules with more gray points perform worse according to our metrics; lower points correspond to more severe violations of proportionality. Points to the left correspond to small groups, points to the right correspond to large groups with unmet justified demand.

5 Experimental Evaluation of Ranking Rules

The results of the previous section provide worst-case guarantees for several interesting ranking rules. We now complement these results with upper bounds, i.e., observed “violations” of the justified demand of voter groups. To this end, we consider a large number of synthetic preferences (various variants of the Impartial Culture model, Mallows’ model, and Urn model) and real-world preference data sets taken from PrefLib [Mattei and Walsh, 2013], and analyze the representation offered by ranking rules. In total, our experiments are based on 315,500 instances. Due to space constraints we only give a very brief description of the experiments and a short discussion of what we learned. A more complete description of the experimental setting and an analysis of the results are provided in the full version of the paper [Skowron *et al.*, 2016].

5.1 Measures of Quality of Group Representation

In our experiments we record every *violation* of justified demand: If, for a ranking r and a $k \in [m]$, there exists a group N' with $\text{avg}(N', r_{\leq k}) < \text{jd}(N', k)$, we set $y_r(N', k) = \frac{\text{avg}(N', r_{\leq k})}{\text{jd}(N', k)}$ and plot a point at $(\alpha(N'), y_r(N', k))$ indicating this violation. Figure 1 shows these plots for different ranking rules, including the “best-of” rule, which selects a ranking r that has the highest quality $q_P(r)$ among the rankings generated by our rules (note that this need not be the optimal ranking according to Definition 3). Violations displayed in the lower part of the plots have a small ratio y_r and thus are more

severe. Note that several points may originate from the same ranking, and different rankings may produce the same point. Hence, these plots do not display *how often* violations occur but rather *in which regions* (small/large groups, minor/major violations) violations have been recorded.

5.2 Results of the Simulations

Our experiments indicate that (i) among the rules we consider, the 2-geometric rule, SeqPAV, Reverse SeqPAV, and Phragmén’s rule are best suited to generate proportional rankings, and (ii) there is no single best choice among these four rules (the best-of rule outperforms all of them). In particular it is noteworthy that these four rules achieve a quality of at least $1/2$ on our test instances, which is only surpassed by the best-of rule (which achieves at least $2/3$). Unfortunately, the best-of rule is certainly not practical, as it is very expensive to compute $q_P(r)$. Further experiments and theoretical results are required to determine which (polynomial-time computable) rule is the best choice (for a given data set). Additional insight can be gained from the plots of geometric rules: The $5/4$ -geometric rule resembles the plot for AV, and the 10-geometric rule shows mainly violations in the top part similarly to Greedy CC. This is a consequence of the fact that the 1-geometric rule is exactly AV and that the p -geometric rule approaches Greedy CC as $p \rightarrow \infty$. Whether $p = 2$ is the sweet spot between these two extremes remains to be determined.

6 Conclusions

In this paper, we have formalized a fundamental concept that is relevant to many real-life applications: *proportional rankings* can provide diversified search results, can accommodate different types of users in recommendation systems, can support decision-making processes under liquid democracy, and can even produce committees with an internal hierarchical structure. Our formalization of this problem allows us to leverage voting rules introduced as far back as the 19th century and apply them to these modern application scenarios.

After evaluating the proportionality of several appealing ranking rules both theoretically and experimentally, we identified four such rules that appear to perform very well in this regard: the 2-geometric rule, Sequential Proportional Approval Voting and its reverse variant, and Phragmén’s rule. However, none of these rules clearly outperforms the other three, and there remains a need for an in-depth analysis to determine which rules are most appropriate for various applications.

While all four of these rules are polynomial-time computable, we have shown that the optimal rule (i.e., the rule that outputs rankings maximizing the quality measure q_P) is NP-hard to compute. It would be desirable to develop ways to compute the outputs of this rule in reasonable time for practical instances, and to identify other ranking rules that may provide an even better approximation to the optimal rule than the rules we have considered in this work.

Acknowledgements

This work was supported by the European Research Council (ERC) under grant number 639945 (ACCORD) and by the Alexander von Humboldt Foundation.

References

- [Aziz *et al.*, 2017] H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 48(2):461–485, 2017.
- [Balinski and Young, 1982] M. Balinski and H. P. Young. *Fair Representation: Meeting the Ideal of One Man, One Vote*. Yale University Press, 1982. (2nd Edition [with identical pagination], Brookings Institution Press, 2001).
- [Behrens *et al.*, 2014] J. Behrens, A. Kistner, A. Nitsche, and B. Swierczek. *The Principles of LiquidFeedback*. 2014.
- [Brill *et al.*, 2017a] M. Brill, R. Freeman, S. Janson, and M. Lackner. Phragmén’s voting methods and justified representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, pages 406–413, 2017.
- [Brill *et al.*, 2017b] M. Brill, J.-F. Laslier, and P. Skowron. Multiwinner approval rules as apportionment methods. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, pages 414–420, 2017.
- [Faliszewski *et al.*, 2017] P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Multiwinner voting: A new challenge for social choice theory. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 2. 2017. To appear.
- [Gallagher, 1991] M. Gallagher. Proportionality, disproportionality and electoral systems. *Electoral Studies*, 10(1):33–51, 1991.
- [Garey and Johnson, 1979] M. Garey and D. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, 1979.
- [Hu *et al.*, 2011] B. Hu, Y. Zhang, W. Chen, G. Wang, and Q. Yang. Characterizing search intent diversity into click models. In *Proceedings of the 20th International Conference on World Wide Web*, pages 17–26, 2011.
- [Janson, 2016] S. Janson. Phragmén’s and Thiele’s election methods. Technical Report arXiv:1611.08826 [math.HO], arXiv.org, 2016.
- [Kingrani *et al.*, 2015] S. K. Kingrani, M. Levene, and D. Zhang. Diversity analysis of web search results. In *Proceedings of the ACM Web Science Conference*, pages 43:1–43:2, 2015.
- [Laslier, 2012] J.-F. Laslier. Why not proportional? *Mathematical Social Sciences*, 63(2):90–93, 2012.
- [Mattei and Walsh, 2013] N. Mattei and T. Walsh. Preflib: A library for preferences. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory*, pages 259–270, 2013.
- [Monroe, 1995] B. L. Monroe. Fully proportional representation. *The American Political Science Review*, 89(4):925–940, 1995.
- [Mora and Oliver, 2015] X. Mora and M. Oliver. Eleccions mitjançant el vot d’aprovació. El mètode de Phragmén i algunes variants. *Butlletí de la Societat Catalana de Matemàtiques*, 30(1):57–101, 2015.
- [Phragmén, 1895] E. Phragmén. *Proportionella val. En val-teknisk studie*. Svenska spösmål 25. Lars Hökersbergs förlag, Stockholm, 1895.
- [Pukelsheim, 2014] F. Pukelsheim. *Proportional Representation: Apportionment Methods and Their Applications*. Springer, 2014.
- [Sánchez-Fernández *et al.*, 2017] L. Sánchez-Fernández, E. Elkind, M. Lackner, N. Fernández, J. A. Fisteus, P. Basanta Val, and P. Skowron. Proportional justified representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, pages 670–676, 2017.
- [Santos *et al.*, 2015] R. L. T. Santos, C. MacDonald, and I. Ounis. Search result diversification. *Foundations and Trends in Information Retrieval*, 9(1):1–90, 2015.
- [Schulze, 2011] M. Schulze. Free riding and vote management under proportional representation by the single transferable vote. Available at <http://m-schulze.9mail.de/schulze2.pdf>, 2011.
- [Skowron *et al.*, 2016] P. Skowron, M. Lackner, M. Brill, D. Peters, and E. Elkind. Proportional rankings. Technical Report arXiv:1612.01434 [cs.GT], arXiv.org, 2016.
- [Thiele, 1895] T. N. Thiele. Om flerfoldsvalg. *Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandling*, pages 415–441, 1895.
- [Wang *et al.*, 2016] Y. Wang, Z. Luo, and Y. Yu. Learning for search results diversification in Twitter. In *Proceedings of the 17th International Conference on Web-Age Information Management*, pages 251–264, 2016.
- [Welch *et al.*, 2011] M. J. Welch, J. Cho, and C. Olston. Search result diversity for informational queries. In *Proceedings of the 20th International Conference on World Wide Web*, pages 237–246, 2011.
- [Xavier, 2012] E. C. Xavier. A note on a maximum k-subset intersection problem. *Information Processing Letters*, 112(12):471–472, 2012.