

Justified Representation in Mutiwinner Voting: Axioms and Algorithms*

Edith Elkind¹

1 Department of Computer Science, Univrsity of Oxford, Oxford, United Kingdom
elkind@cs.ox.ac.uk

Abstract

Suppose that a group of voters wants to select $k \geq 1$ alternatives from a given set, and each voter indicates which of the alternatives are acceptable to her: the alternatives could be conference submissions, applicants for a scholarship or locations for a fast food chain. In this setting it is natural to require that the winning set represents the voters fairly, in the sense that large groups of voters with similar preferences have at least some of their approved alternatives in the winning set. We describe several ways to formalize this idea, and show how to use it to classify voting rules; surprisingly, two voting rules proposed in the XIXth century turn out to play an important role in our analysis.

1998 ACM Subject Classification F.2 - Analysis of algorithms and problem complexity

Keywords and phrases voting, committee selection, axioms, PAV

Digital Object Identifier 10.4230/LIPIcs.CVIT.2016.23

1 Setup

The program committee (PC) of a conference in theoretical computer science is about to decide on the set of invited speakers: the chair has solicited nominations from the program committee and the steering committee, and now the PC members have to cast their votes. To keep things simple, the chair sets up a Doodle poll, and asks each PC member to indicate which of the potential speakers they approve of. There are 36 PC members, and four slots for invited talks.

The conference is very broad in scope: it accepts papers on logic, formal methods, algorithms, and complexity. Each PC member identifies with one of these four areas: there are 14 algorithms researchers, 9 experts on formal methods, 8 complexity theorists, and 5 logicians. Each potential speaker is also associated with a particular research area, though some speakers may appeal to researchers in two or more areas. For instance, candidates A , B , C , D and E are supported by almost all algorithms and complexity researchers, in the sense that each of them gets at least 20 votes from PC members who belong to these two communities, candidate F is supported by all formal methods researchers, one algorithms researcher and one complexity theorist, and there are various other candidates who receive 6 or fewer votes; in particular, each logician suggest a different speaker that has no support from other PC members.

How should the chair decide whom to invite given the PC members' votes? Of course, the simplest option is to count how many times each candidate is approved, and select the

* This work was partially supported by European Research Council (ERC) under grant agreement 639945.



four candidates with the maximum number of approvals (breaking ties, say, by seniority or travel costs). However, in our example this means that the chair would end up selecting four candidates who only appeal to “Theory A” (i.e., algorithms and complexity) researchers, and this does not seem fair: if there are four invited talks, and the formal methods community is represented by $9 = 36/4$ PC members, intuitively the formal methods researchers are entitled to at least one slot (which they would probably like to allocate to candidate F). On the other hand, it seems that at least two slots should be allocated to “Theory A” speakers, as they are supported by more than half of the voters. Finding the right balance among these concerns is not easy. In what follows, we will describe a formal model that helps us reason about such settings. Our presentation is based on a sequence of papers [6, 1, 17, 15, 4, 9, 19] by several overlapping groups of researchers that were published in 2014–2017.

2 Formal Model

We consider a set of n voters $[n] = \{1, \dots, n\}$ who choose among alternatives from the set $A = \{a_1, \dots, a_m\}$. Each voter $i \in [n]$ submits an approval ballot $A_i \subseteq A$, which lists all alternatives that she approves of. The goal is to select a winning set, or *committee*, of size k , where $1 \leq k \leq m$. Throughout this paper, we assume that voters are not strategic.

A *committee selection rule* is a mapping that, given A , k , and the list $(A_i)_{i \in [n]}$, outputs one or more winning committees of size k . Note that we allow ties, but the set of winning committees must be non-empty.

2.1 Examples of Committee Selection Rules

We will now present several committee selection rules, and briefly discuss their properties.

Approval Voting (AV) Under this committee selection rule the score of each alternative $a \in A$ is the number of approvals it gets, i.e., $|\{i \mid a \in A_i\}|$, and the winning committee consists of the k alternatives with the highest scores (if there are ties, we output all committees that can be obtained by breaking these ties in some way). While this is a very simple rule, it may fail to be fair, in the sense that even large minorities of voters may end up not being represented in the winning committee(s). In our example, if AV is used, there would be no formal methods talk.

Proportional Approval Voting (PAV) Under approval voting, each voter is assumed to assign a ‘utility’ of one to each of her approved alternatives in the committee. In contrast, under proportional approval voting, we assume that the ‘marginal utility’ that a voter derives from her j -th approved committee member is $\frac{1}{j}$. Formally, given a committee W and a voter i , we write $u_i(W) = h(|W \cap A_i|)$, where $h(j) = 1 + \frac{1}{2} + \dots + \frac{1}{j}$, and output all committees W that maximize the total ‘utility’ $\sum_{i \in [n]} u_i(W)$. In our example the total utility of $\{A, B, C, D\}$ is $20 \times (1 + \frac{1}{2} + \frac{1}{3} + \frac{1}{4})$. This is less than the total utility of $\{A, B, C, F\}$, which is $20 \times (1 + \frac{1}{2} + \frac{1}{3}) + 9$, so the PAV rule would allocate a slot to the formal methods speaker. This rule has a very long history: it was proposed by a Danish polymath Thorvald N. Thiele in 1895 [20], and was subsequently rediscovered by other researchers (see, e.g., [10]).

Of course, instead of using harmonic weights $h(\cdot)$, one can use another weight function $w(\cdot)$, i.e., set $u_i(W) = w(|W \cap A_i|)$ for some $w : \mathbb{N} \cup \{0\} \rightarrow \mathbb{R}$; we refer to the resulting rule as w -AV. For instance, setting $w(i) = i$ for all $i \in \mathbb{N} \cup \{0\}$ corresponds to Approval Voting. Another interesting option is $w(0) = 0$, $w(1) = 1$, which is sometimes referred to as Chamberlin–Courant Approval Voting (CCAV). The CCAV rule aims to select

committees that represent as many voters as possible, but does not necessarily allocate more representatives to larger groups: in our example, CCAV may select just one of the alternatives in $\{A, B, C, D\}$.

Our (and Thiele's) decision to use harmonic weights may seem mysterious at this point, but we will soon see that there are good reasons for it.

Sequential Proportional Approval Voting (SeqPAV) One may wonder if winning committees under PAV can be computed efficiently. It turns out that this is unlikely: deciding if a given committee should win under PAV is coNP-complete [2, 18]. In fact, this was already intuitively obvious to Thiele (more than 70 years before the formal language to talk about such issues was invented), and he proposed a greedy variant of his rule, which is now referred to as Sequential Proportional Approval Voting (SeqPAV). Under this rule, the committee is formed in k steps; we start with an empty committee $W = \emptyset$, and at each step we add a candidate with the largest marginal utility, where the marginal utility of an alternative $a \in A \setminus W$ is computed as

$$\sum_{i \in [n]} (u_i(W \cup \{a\}) - u_i(W)).$$

In Thiele's original proposal $u_i(W) = h(|W \cap A_i|)$, but the same approach can be used with other weight functions. In particular, under the sequential version of the CCAV rule (SeqCCAV) at each step for each alternative $c \in A \setminus W$ we compute

$$\#c = |\{i \in [n] \mid c \in A_i, A_i \cap W = \emptyset\}|,$$

and add an alternative for which this quantity is maximized.

Monroe Approval Voting (MonAV) and its greedy variant (GrMonAV) Under MonAV, to evaluate a given committee W , we assign scores to all many-to-one matchings between voters and members of W such that each committee member is matched to $\lfloor \frac{n}{k} \rfloor$ or $\lceil \frac{n}{k} \rceil$ voters: the score of a matching is the number of voters who are matched to alternatives they approve. We then select a matching with the maximum score; this score is said to be the score of W . We output committees with the maximum score. Under this rule, $\{A, B, C, F\}$ should be preferred to $\{A, B, C, D\}$: if we choose $\{A, B, C, F\}$, we can match 9 voters to each of A, B , and F , so the score of this committee is at least 27, whereas the score of $\{A, B, C, D\}$ cannot exceed 22.

Now, we do not need to go over all possible matchings to compute the score of a given committee W : a committee score can be computed in polynomial time using network flow techniques. However, finding a committee with the maximum score is computationally hard [14]. Therefore, it is natural to consider a greedy variant of this rule [8]: again, we proceed in k stages, starting with $W = \emptyset$, and at each stage we select an alternative $c \in A \setminus W$ and $\lfloor \frac{n}{k} \rfloor$ or $\lceil \frac{n}{k} \rceil$ unassigned voters to be matched to this alternative, so that as many of these voters as possible approve c . It is instructive to compare SeqCCAV and GrMonAV: while both rules add an alternative and remove a number of voters associated with this alternative at each step, SeqCCAV only removes voters who approve the selected alternative (and may remove a large number of voters), while GrMonAV removes roughly $\frac{n}{k}$ voters, some of which may not approve the selected alternative.

Rules based on fractional vote transfers One difficulty with the GrMonAV rule is that it matches each voter to a single alternative; if a voter does not approve the alternative she is matched to, naturally we cannot expect her to be happy with the selected committee. To mitigate this difficulty, Sánchez Fernández et al. [9] propose a family of rules, which they call EJR-Exact, that operate similarly to GrMonAV, but may assign a fraction of

a voter’s weight to a selected alternative. The name of this family of rules may seem mysterious at this point, but it will be explained in the next section. The specific weight assignment rules associated with EJER-Exact rules are fairly complicated, and we will not discuss them here. Another iterative rule that is based on fractional weight assignment was recently proposed by Aziz and Lee [5]; they call their rule the Expanding Approvals Rule (EAR). We note that the idea of fractional assignment of voters to alternatives is quite well-accepted in the context of voting with ranked ballots (where each voter submits an ordering of the alternatives). In fact, the method used to elect the Senate of the Republic of Ireland uses the Single Transferable Vote rule with fractional vote transfers [21].

Phragmén’s rule and its sequential variants Under all voting rules described so far, for each voter we define her ‘utility’ or ‘satisfaction’ from a given committee, and choose a committee that maximizes this quantity. A Swedish mathematician Lars Edvard Phragmén, who published his work [13] around the same time as Thiele did, proposed a radically different perspective: each committee member is associated with one unit of *load*, which needs to be distributed among voters who support that alternative as evenly as possible; an alternative’s load can only be allocated to voters who approve her, and we aim to minimize the maximum voter load. Thus, for each committee, we find the best load assignment in terms of the maximum voter load and then output committees for which this quantity is minimized. Just as for PAV, finding optimal committees under this rule is NP-hard [7], and Phragmén has developed an iterative version of this rule, which is computationally tractable (again, long before the concept of computational complexity was formalized); the subsequent literature [15] refers to this rule as SeqPhragmen. Another rule that is, in some sense, a dual of Phragmén’s rule (or, more precisely, a sequential version of its dual) has been recently proposed by Sánchez Fernández et al. [16], who view their rule as an extension of D’Hondt’s method for party list representation, and therefore refer to their rule as Open D’Hondt Method.

There are many other approval-based committee selection rules; we refer the reader to a survey by Kilgour [11]. However, the rules listed above seem to be particularly well suited to achieve our representation desiderata.

3 Axioms

We will now formulate our first representation axiom; intuitively, this axiom says that there should be a formal method to talk in the program of our conference. In more detail, as there are n voters and k slots, a group of voters of size $\frac{n}{k}$ is ‘entitled’ to a slot. However, there are many ways of partitioning voters into groups of this size, and it is easy to see that we cannot guarantee representation to every such group. We therefore focus on groups that are *cohesive*, i.e., there exists an alternative approved by all group members; we would like to ensure that no such group is completely disenfranchised. This axiom was proposed by Aziz and Walsh [6], in a workshop paper that started this line of research.

► **Definition 1.** A committee W provides *justified representation (JR)* if for every group of voters V such that $|V| \geq \frac{n}{k}$ and $\bigcap_{i \in V} A_i \neq \emptyset$ there exists a voter $i \in V$ who approves a member of W , i.e., $|A_i \cap W| \geq 1$. A committee selection rule *satisfies justified representation* if every committee in its output provides JR.

It is easy to see that Approval Voting does not satisfy JR; indeed, this is illustrated by our conference example. However, most of the other rules we consider satisfy JR: this is

true for PAV (and for every w -AV rule that satisfies $w(1) = 1$, $w(j) - w(j-1) \leq \frac{1}{j}$) [1], for MonAV and GrMonAV [1], as well as for all variants of Phragmén's rule that we have mentioned [7] (including the ODH rule [16]). A prominent exception is the SeqPAV rule, which may fail JR [1]. Indeed, sequential variants of almost all weighted AV rules fail JR; the only exception is the sequential variant of the CCAV rule (SeqCCAV) [1].

For most of these rules it is not difficult to prove that they satisfy JR; we provide the argument for two rules to illustrate the main ideas.

► **Theorem 2.** *PAV and SeqCCAV satisfy JR.*

Proof. Consider first PAV. Fix a committee W that does not provide JR, i.e., there exists a set of voters V with $|V| \geq \frac{n}{k}$ and $\cap_{i \in V} A_i \neq \emptyset$ such that $A_i \cap W = \emptyset$ for all $i \in V$. Let a be some alternative in $\cap_{i \in V} A_i$. We will argue that there is an alternative $c \in W$ such that $(W \setminus \{c\}) \cup \{a\}$ has higher PAV score than W does, i.e., PAV cannot output W .

Define the *marginal contribution* of an alternative $c \in W$ to be

$$m(c) = \sum_{i \in [n]} (u_i(W) - u_i(W \setminus \{c\})).$$

By changing the order of summation, we can rewrite the sum of marginal contributions of alternatives in W as

$$\sum_{c \in W} \sum_{i \in [n]} m(c) = \sum_{i \in [n]} \sum_{c \in W} (u_i(W) - u_i(W \setminus \{c\})). \quad (1)$$

Now, for each voter i with $A_i \cap W = \emptyset$ (so, in particular, for each voter in V) her contribution to the sum in (1) is 0. On the other hand, if $|A_i \cap W| = s > 0$ then for each $c \in A_i \cap W$ we have $u_i(W) - u_i(W \setminus \{c\}) = \frac{1}{s}$ and for each $c \in W \setminus A_i$ we have $u_i(W) - u_i(W \setminus \{c\}) = 0$. Thus, in this case i contributes $s \cdot \frac{1}{s} = 1$ to the sum in (1). As there are at most $n - \frac{n}{k}$ voters who make a non-zero contribution, the sum in (1) can be bounded as $n - \frac{n}{k} < n$, and therefore there is an alternative $c \in W$ with $m(c) < \frac{n}{k}$. Now, consider the committee $(W \setminus \{c\}) \cup \{a\}$. By removing c , we reduce the PAV score by less than $\frac{n}{k}$, and by adding a we increase the PAV score by at least $\frac{n}{k}$, as each voter in V now contributes 1 to the PAV score. This proves our claim.

Now, consider SeqCCAV. Suppose that after k steps it outputs a committee W such that for some set of voters V with $|V| \geq \frac{n}{k}$ and $\cap_{i \in V} A_i \neq \emptyset$ we have $A_i \cap W = \emptyset$ for all $i \in V$. Again, let a be an alternative in $\cap_{i \in V} A_i$. Note that at each iteration we have $\#a \geq \frac{n}{k}$ because of the voters in V . Note that a was in $A \setminus W$ at each step of the procedure, but was never selected. Thus, each of the alternatives added to W was approved by at least $\frac{n}{k}$ previously unrepresented voters, i.e., at each step the number of voters i with $A_i \cap W \neq \emptyset$ grew by at least $\frac{n}{k}$. Thus, after k steps no unrepresented voters should have remained, a contradiction. ◀

By now, the reader may have an impression that essentially any reasonable approval-based committee selection rule satisfies JR. However, this is not quite the case: there are well-established rules, such as, e.g., Satisfaction Approval Voting (SAV) and Minimax Approval Voting (MAV) that fail this axiom; see, e.g., [11] for the definitions of these rules, and [1] for respective counterexamples. Indeed, the original workshop paper of Aziz and Walsh only identified a single voting rule that satisfied EJRR, namely, SeqCCAV. While the authors had considered PAV, they left it as an open problem to show that this rule satisfied JR. This open problem was resolved independently by two groups of researchers (Brill, Conitzer, Freeman at Duke and Elkind at Oxford), resulting in the joint paper [1].

It is perhaps interesting that for Phragmén’s rule and Monroe’s rule moving from the optimization-based version of the rule to its sequential variant preserves JR, but for PAV this is not the case; moreover, compliance with JR cannot be restored by tweaking the weights.

Now, let us revisit our example. Each of the rules that provide JR guarantees that there will be at least one formal methods talk. But we also want to ensure that at least two speakers representing the algorithms and complexity community are selected, and some of our rules do not have this property. For instance, depending on the distribution of approvals, SeqCCAV may select A , F , and two of the speakers supported by one logician each; this is not precluded by the JR axiom. Thus, perhaps we need a stronger axiom, which says that very large groups of voters with shared interests should get not one, but several representatives, in proportion to their size and cohesiveness.

To formalize this idea, it will be convenient to introduce additional terminology: for each $\ell \geq 1$ we say that a group of voters $V \subseteq N$ is ℓ -large if $|V| \geq \ell \cdot \frac{n}{k}$; V is ℓ -cohesive if $|\bigcap_{i \in V} A_i| \geq \ell$. We will now formulate two axioms that have been proposed in the literature; both axioms aim to capture the idea that large, cohesive groups deserve several representatives.

► **Definition 3. Extended justified representation (EJR).** A committee W provides *extended justified representation (EJR)* if for every $\ell \in [k]$ and every ℓ -large, ℓ -cohesive group of voters V there exists a voter $i \in V$ who approves at least ℓ members of W , i.e., $|A_i \cap W| \geq \ell$.

Proportional Justified Representation (PJR). A committee W provides *proportional justified representation (PJR)* if for every $\ell \in [k]$ and every ℓ -large, ℓ -cohesive group of voters V there are at least ℓ members of W who are approved by a voter from V , i.e., $|W \cap (\cup_{i \in V} A_i)| \geq \ell$.

We say that a committee selection rule satisfies EJR (respectively, PJR) if all committees in the output of this rule provide EJR (respectively, PJR).

The EJR axiom has been proposed by Aziz et al. [1]; the PJR axiom was subsequently suggested in a workshop paper by Sánchez Fernández et al. [17], which was later extended to a AAAI-17 paper [15]. Both axioms are stronger than JR; indeed, JR corresponds to setting $\ell = 1$ in either of these axioms. Both EJR and PJR say that large, cohesive groups should get several representatives, but they interpret this requirement differently, with EJR being more demanding: if one voter gets at least ℓ representatives then clearly collectively the voters in V get at least ℓ representatives. An attractive feature of PJR is that it is compatible with perfect representation axiom: in essence, this axiom says that if there is a committee that perfectly represents the electorate, in the sense that each committee member is approved by exactly $\frac{n}{k}$ voters, then some such committee should be selected. In contrast, EJR is not compatible with this axiom [17, 15].

It is not immediately obvious whether a committee satisfying EJR or PJR always exists. However, it turns out that this is indeed the case, and, in fact, some of our voting rules always produce such committees.

Of course, we cannot expect AV or SeqPAV to satisfy these axioms, as they fail the weaker JR axiom. The other rules on our list satisfy JR; but do they satisfy PJR or EJR? The first attempt to identify rules that satisfy PJR was not fully successful: Sánchez Fernández et al. [15] argue that MonAV and GrMonAV satisfy PJR as long as k divides n , but show that these rules may violate PJR if this condition is not satisfied. Moreover, these rules fail EJR even if k divides n [1]. The variants of Phragmén’s rule (i.e., Phragmen, SeqPhragmen and ODH) all satisfy PJR, but not EJR [7, 16]; this is also the case for the EAR rule [5].

In contrast, the PAV rule satisfies not just PJR, but EJR [1]. The argument is not difficult: just as in the proof of Theorem 2 we consider a committee that fails EJR and identify a swap that increases the total PAV score. Even more surprisingly, both PJR and EJR characterize PAV within the class of w -AV rules: every w -AV rule where w is not equivalent to the harmonic weight sequence fails PJR (and hence EJR) [1, 15]. Thus, the PJR/EJR axioms justify the choice of harmonic weights in the definition of this function, which was made by Thiele back in 1895. Finally, we can now resolve the mystery concerning EJR-Exact: this class of rules was designed with the explicit purpose of obtaining committees that provide EJR, and all rules in this class satisfy EJR (we note that Sánchez Fernández et al. [9] actually describe a broader family of rules; while all rules in this family satisfy PJR, only the EJR-Exact rules provide EJR).

4 Algorithms

So far, we have not paid much attention to the computational complexity of identifying representative committees. However, in practice algorithmic complexity may be an important consideration, especially in large elections where the size of the committee to be elected is large as well. Thus, it would be desirable to have a voting rule that satisfies a powerful justified representation axiom (ideally, EJR) and is polynomial-time computable. Another relevant question is deciding whether a given committee provides (proportional/extended) justified representation.

We start by considering the second question. One may think that, to verify whether a committee provides JR, one has to look at all groups of voters of size at least $\frac{n}{k}$, which is obviously computationally expensive. However, one can obtain an answer much more efficiently by focusing on *alternatives*: for each alternative not in the committee, we can check if there is a ‘large’ group of voters who approve that alternative, but do not approve anyone in the committee. This check can be performed in polynomial time. Regrettably, this approach does not extend to verifying PJR/EJR, as we would have to consider *groups* of alternatives. Indeed, it has been proved that deciding whether a given committee provides PJR or EJR is coNP-complete [1, 3].

Now, finding an efficiently computable rule that satisfies JR is not difficult: for instance, SeqCCAV fits the bill, and so do many other sequential rules discussed in this paper. For PJR the task is somewhat more challenging: the original workshop paper [17] that introduced the concept of PJR could not identify a single polynomial-time computable rule that satisfies this axiom, and the GrMonAV rule, while efficiently computable, is only guaranteed to satisfy PJR when the committee size k divides the number of voters n . However, essentially at the same time it was established that SeqPhragmen satisfies PJR [7], and so does ODH [16]. Subsequently, it was proved that several rules that are based on fractional vote transfers are efficiently computable and provide PJR; this includes, in particular, the EAR rule of Aziz and Lee [5] and the rules of Sánchez Fernández et al. [9]. Thus, there are many different approaches that allow us to find committees providing PJR quickly and efficiently.

In contrast, designing efficient procedures for finding committees that provide EJR turned out to be more challenging. Indeed, even sophisticated sequential rules such as SeqPhragmen and ODH were shown to fail EJR. However, in the beginning of 2017 two groups of researchers [4, 19] independently made the same observation: the proof that committees with the maximum PAV score satisfy EJR also applies to committees with *almost* optimal score. This observation suggests a simple local search algorithm, which we call LS-PAV: we start with an arbitrary committee W , and look for a committee member

$c \in W$ and a non-member $a \in C \setminus W$ such that swapping c with a increases W 's PAV score by a significant amount. The condition that the change should be significant guarantees that we will converge in polynomial time, yet the resulting committee still satisfies EJR. Thus, committees providing EJR can be computed in polynomial time. However, while this swap-based algorithm is very attractive from a theoretical perspective, convincing voters to use it may be difficult: they would see LS-PAV as trying to optimize a certain quantity, but then stopping before the optimal value is reached. Many voters would find that unappealing, especially if they can identify a committee that has a (slightly) higher PAV score than the selected committee.

On the other hand, the EJR-Exact rules of Sánchez Fernández et al. [9], which were proposed at around the same time as both of the local swap algorithms, also admit a polynomial-time winner determination algorithm and provide EJR, while operating similarly to traditional voting rules, such as STV. However, their weight transfer mechanism is not easy to explain, and the proof that these rules satisfy EJR is far from straightforward. The results of all three teams were merged into a paper that was accepted to AAAI-18 [?].

While both LS-PAV and EJR-Exact rules are probably too complex to be adopted in a real-life committee selection setting, they nevertheless provide a useful proof of concept: it is possible to find committees that provide EJR in polynomial time! Recall that checking if a given committee provides EJR/PJR is computationally hard: that is, somewhat counterintuitively, constructing a committee that provides EJR from scratch is easier than deciding if a given committee has this property.

5 Conclusions

While the basic questions concerning JR, PJR and EJR have been answered, further questions remain. For instance, Aziz et al. [1] propose a notion of justified representation that is even stronger than EJR and is based on core stability in an associated cooperative game. At the moment, it remains open whether every input admits a committee that provides this form of justified representation; it is known that PAV may fail to output a committee with this property even when it exists. Another open question is whether EJR is compatible with other important axioms, such as, in particular, committee monotonicity (this axiom says that if a is in a winning committee of size k , it should be in a winning committee of size $k + 1$). Indeed, this research can be seen as a part of the broader agenda of axiomatizing approval-based committee selection rules (see, e.g., [12]) and understanding their computational complexity.

References

- 1 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.
- 2 H. Aziz, S. Gaspers, J. Gudmundsson, S. Mackenzie, N. Mattei, and T. Walsh. Computational aspects of multi-winner approval voting. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2015)*, pages 107–115, 2015.
- 3 Haris Aziz and Shenwei Huang. Computational complexity of testing proportional justified representation. Technical Report arXiv:1612.06476, arXiv.org, December 2016.
- 4 Haris Aziz and Shenwei Huang. A polynomial-time algorithm to achieve extended justified representation. Technical Report arXiv:1703.10415, arXiv.org, March 2017.

- 5 Haris Aziz and Barton Lee. Achieving proportional representation via voting. Technical Report arXiv:1708.07580, arXiv.org, August 2017.
- 6 Haris Aziz and Toby Walsh. Justified representation in approval-based committee voting. In *Proceedings of the 8th Multidisciplinary Workshop on Advances in Preference Handling (MPREF-2014)*, 2014.
- 7 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 (AAAI-2017)*, pages 406–413, 2017.
- 8 Edith Elkind, Piotr Faliszewski, Piotr Skowron, and Arkadii Slinko. Properties of multi-winner voting rules. *Social Choice and Welfare*, 48(3):599–632, 2017.
- 9 Luis Sánchez Fernández, Edith Elkind, and Martin Lackner. Committees providing EJR can be computed efficiently. Technical Report arXiv:1704.00356, arXiv.org, April 2017.
- 10 Svante Janson. Phragmén’s and Thiele’s election methods. Technical Report arXiv:1611.08826, arXiv.org, November 2016.
- 11 D. Kilgour. Approval balloting for multi-winner elections. In J. Laslier and R. Sanver, editors, *Handbook on Approval Voting*, pages 105–124. Springer, 2010.
- 12 Martin Lackner and Piotr Skowron. Consistent approval-based multi-winner rules. Technical Report arXiv:1704.02453, arXiv.org, April 2017.
- 13 E. Phragmén. Sur une méthode nouvelle pour réaliser, dans les élections, la représentation proportionnelle des partis. *Öfversigt af Kongliga Vetenskaps-Akademiens Förhandlingar*, 51(3):133–137, 1894.
- 14 A. Procaccia, J. Rosenschein, and A. Zohar. On the complexity of achieving proportional representation. *Social Choice and Welfare*, 30(3):353–362, April 2008.
- 15 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 (AAAI-2017)*, pages 670–676, 2017.
- 16 L. Sánchez-Fernández, N. Fernández, and J. A. Fisteus. Fully open extensions to the D’Hondt method. Technical Report arXiv:1609.05370, arXiv.org, September 2016.
- 17 L. Sánchez-Fernández, N. Fernández, J. A. Fisteus, and P. Basanta Val. Some notes on justified representation. In *Proceedings of the 10th Multidisciplinary Workshop on Advances in Preference Handling (MPREF-2016)*, 2016.
- 18 P. Skowron, P. Faliszewski, and J. Lang. Finding a collective set of items: From proportional multirepresentation to group recommendation. *Artificial Intelligence*, 241:191–216, 2016.
- 19 Piotr Skowron, Martin Lackner, Edith Elkind, and Luis Sánchez Fernández. Optimal average satisfaction and extended justified representation in polynomial time. Technical Report arXiv:1704.00293, arXiv.org, April 2017.
- 20 T. N. Thiele. Om flerfoldsvalg. In *Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandlingar*, pages 415–441. 1895.
- 21 Nicolaus Tideman. The single transferable vote. *The Journal of Economic Perspectives*, 9(1):27–38, 1995.