Artificial Intelligence Research for Fighting Political Polarisation: A Research Agenda

Andrea Loreggia[1][0000−0002−9846−0157], Nicholas Mattei[2][0000−0002−3569−4335], and Stefano Quintarelli[3]⋆

[1] European University Institute andrea.loreggia@gmail.com
[2] Tulane University nsmattei@tulane.edu
[3] Copernicani ifdad.org@quintarelli.it

Abstract. Political polarization is a growing phenomenon in numerous countries around the world. In many cases, this polarization is leading to a divergence between the positions of elected decision makers, creating a loss of middle-ground and compromise, which is in turn causing the democratic process to grind to a halt. The causes of this polarization are both contentious and numerous, with many sociologists, political scientists, economists, and historians weighing in on the matter. Although computer science cannot comment on the causes, we may be able to develop group decision making algorithms and tools to help ameliorate the situation. In this paper we outline some general research proposals on the possibility of combining machine learning techniques with results from computational social choice to mitigate the polarization. Specifically, we focus on the design of multi-winner voting rules that are both efficiently computable and have provable guarantees on desirable properties such as proportionality and representation. We hope that research advances in these areas can have concrete impact on fighting political polarization.

Keywords: Multi-winner voting rule · Clustering · Political polarization

1 Introduction

One only needs to look at the front page of the local paper to hear about how democracy is experiencing a global crisis caused, in part, by increasing polarization. Political polarization usually refers to the emergence of an extreme divergence in how citizens in a society think, especially when two dominant views emerge driving people apart. This phenomenon is observed and discussed in the United States, Europe, and numerous other other countries [7, 11, 27]). One of the principle arguments for the benefits of a pluralistic democracy relies on the variety of voices and differences in a given society coming together to make collective choices [40]. However, scholars have noted that increasingly citizens are perceiving and describing society along a single dimension, in terms of *us* versus *them* [36].

Finding the root causes that lead to strong polarization within the democratic process is not easy and beyond the scope of research in computer science and computational social choice [9]. The reasons for and consequences of strong polarization are the subjects of extensive research in sociology, political science, economics, and history. Although we cannot cure the causes, we can design and propose tools to tackle and mitigate polarization using techniques from across computer science research, including techniques from artificial intelligence, machine learning, computational social choice [9], preference reasoning [39], and principles of human-centered AI (HCAI) [35, 41]. In order to design and implement technological solutions to assist with the decision making process, we should also guarantee that these tools adhere to ethical and moral principles [23, 38, 28]. One domain where research from the field of computational social choice (COMSOC) may be able to help is creating new tools and techniques to enable the democratic decision making process. In COMSOC, researchers study the computational aspects of classical social choice problems [40] including voting and matching [9]. Typically, in an election, a set of agents submit their preferences over a set of candidates, and then a voting rule, i.e., aggregation mechanism, selects a single or multiple winners to be shared by that group [45]. A representative election process should, through normative proofs and empirical experiment, guarantee that candidates who are elected are both representative of and proportional to the diverse parts of a given society as expressed in their voting preferences. Unfortunately, for many voting rules a long list of strategic behaviors can be exploited on the part of the voters and/or candidates in order to arrive at an outcome that is not intended or has been *manipulated* [18]. There exists an extensive literature on the strategic behavior of agents in voting systems, including manipulation [21], bribery [16, 15] and control [29, 32].

The current COVID-19 pandemic and the ensuing worldwide crisis are highlighting the brittleness of many of the decision making processes in our society at the local, national, and international stages. At the same time, the pandemic has underlined the importance of technology in our everyday life. From remote working to distance learning, from social networks to shopping online, technology plays a key role in maintaining some sense of normalcy by shifting once in-person activities online, leveraging technological solutions. There has been significant recent work in the social choice and COMSOC fields to promote novel methods of electing representatives and delegating votes in what is called *interactive democracy* [1, 10, 22]. We see the development of new, robust, and efficiently computable voting rules as an important step in this direction.

1.1 Our Contribution

In this paper we discuss existing literature in the area of multi-winner elections and propose concrete research directions for combining aspects of artificial intelligence, machine learning, and computational social choice to improve the quality of elected bodies in terms of representation and proportionality. We hope that by addressing these concerns we may mitigate some aspects of political polarization.

In recent years a novel voting rule, called Majority Judgement (MJ) [5, 6], has been proposed for single winner elections. The proponents of MJ for single winner elections tout its superiority as it does not fall victim to several issues in classical social choice including the Condorcet Paradox and Arrow's Paradox, and it is resistant to strategic manipulation on the part of the voters [4, 40]. Informally, our goal is to develop a multi-winner rule that prevents elected committees from being polarized even if the voters themselves are polarized, while still maintaining proportionality and representation. We propose a plan of research to extend majority judgement to multi-winner settings, this advance would: provide a new multi-winner voting rule that may provide guarantees on proportionality and representativeness; scale up algorithms for multi-winner voting that suffer from high computational cost; design a method which could be helpful in fighting strategic behavior connected to manipulation or gerrymandering. We next turn to a background on multi-winner voting and some motivating examples. We then survey some of recent work from the computational social choice research community focused on multi-winner voting, we outline a method to extend majority judgement to multi-winner elections, and end with some research challenges.

2 Social Choice and Motivating Examples

Mapping agents' preferences into a subset of common choices among all the available ones is a well studied problem in many different disciplines. From economic theory to computational social choice many researchers have studied which axiomatic properties are desired and which are not [9, 40]. In a classical social choice setting, a group of agents, i.e., the voters, are asked to express their preferences over a set of alternatives, i.e., the candidates. Voters do this by comparing candidates and reporting, depending on the aggregation mechanism to be used, a ranking that could be strict or contain ties, or possibly approvals over candidates the voters would be willing to accept. These reports are then fed into an aggregation mechanism, i.e., a voting rule, and this voting rule returns an outcome. In a classic social choice setting this is a single-winner, but could also be a ranking over the candidates or a subset of alternatives in a multi-winner procedure, i.e., committee selection. In this paper we focus on the multi-winner setting [17].

Formally, a multi-winner election (V, C, F, k) is defined by a set of voters V expressing preferences over a number of candidates C, and then a voting rule F returns a subset of size k *winning candidates* [17]. Preferences can be specified in many different ways, for instance rankings, approvals, judgements or a set of binary comparisons. The first and more common approach represents preferences as linear orders, i.e., transitive, anti-symmetric and complete binary relations over a set of discrete alternatives C. A profile of preferences $P = (>_1, ..., >_n)$ is defined by a preference relation $>_i$ for each of the n individual voters. We write $a >_i b$ to denote that agent i prefers item b to item a in profile P. A (non-resolute) voting rule F associates with every profile P a non-empty subset of winning candidates. A large number of voting rules have been proposed in the literature along with a number of methods to compare between them [9, 45].

We focus on the use of Majority Judgement [5] and we assume that users express opinions or judgements over a subset of candidates, that is called an opinion profile P. Opinions are natural language tags that are chosen *a priori*. They depict different grades of support for a candidate, i.e., Excellent, Very Good, Good, Passable, Insufficient. Such finite, ordered set of evaluations Λ can be linearly ordered by $\succ$. An electorate's opinion profile on a candidates is the set of her, his, or its grades $\alpha = (\alpha_1, ..., \alpha_n)$, where $\alpha_j \in \Lambda$ is voter j's evaluation of the candidate. Elections are unweighted, or anonymous, meaning only the grades count, not which voter gave what grade. Voting rules are typically judged according to a number of normative properties or axioms. Informally, some of the properties we are interested in are:

Anonymity: A voting rule F is said to be *anonymous* if it treats individuals symmetrically, i.e. switching two individual's preferences does not change the outcome, $F(P_1, \ldots, P_n)$=$F(P_{\pi(1)}, \ldots, P_{\pi(n)})$ where $\pi : N \rightarrow N$ is a permutation.

Neutrality: A voting rule F is said to be *neutral* if it treats alternatives symmetrically.

Representation: A voting rule F provides representation if, for each partition of voters $V_i \in V$ (with $|V_i| \geq \lfloor \frac{n}{k} \rfloor$), it assigns to the committee at least one candidate elected by the partition V_i.

Proportionality: A voting rule F provides proportionality if, for each partition of voters $V_i \in V$, with $|V_i| \geq \lfloor \frac{n}{k} \rfloor$, the number of candidates selected for the committee is proportional to the size of the partition.

2.1 Motivating Example

We now turn to the problem of political polarization as evidenced through the selection of a representative committee. A group of 30 individuals should elect a representative committee of 3 candidates. The set of candidates is $C = \{A, B, C, D, E\}$ and each voter is required to express an approval ballot over the candidates as seen in Table 1.

Table 1. Approvals of the voters in the example election.

# Voters	Approve	Disapprove
20	A,B,C	D,E
10	E	A, B,C,D

There exists a strong polarization in the community: a group of 20 people approve A, B, C, which is the subset disapproved by the other group. If we use the basic approach of *approval voting* [8] for electing the committee and count the number of approvals then A, B, C would get 20 points, E would get 10 points and D would get no points. Thus, the final committee would be formed by $[A, B, C]$. Unfortunately, this solution fully represents the group of 20 people, as their whole approved set is elected, but at the same time it penalizes the smaller group of 10 people, since they disapprove the winning candidates. The committee formed by $[A, B, C]$ may reinforce the strong polarization driven by the group of 20 people, possibly exasperating the political situation.

According to the definition of proportionality, for $k = 3$ every candidate in the committee should be supported by at least 33.33% of the voters. Alternatively, to respect representation, each group with at least $\lfloor \frac{30}{3} \rfloor = 10$ voters should elect at least one candidate. In order to respect proportionality, the number of candidates to be elected should depend on the size of the group. Having the first group of people $20/30 = 66.66\%$ of the whole, it should be represented by at least 2 candidates in the committee. The second group of people has $10/30 = 33.33\%$ of the voters and should be represented with at least one candidate. Instead, we could use a different method to select our committee which would guarantee a proportional outcome. In this example we use Proportional Approval Voting (PAV) [44], which is a well-know multiwinner voting rule, which unfortunately suffers from a high computational cost [3]. Informally under PAV, each voter derives a utility of $\sum_{i=1}^{j} \frac{1}{i}$ from a committee that contains exactly j of her approved candidates. The goal is to maximize the sum of the voters' utilities. Under PAV, the final committee could be chosen among three possible committees, as shown in Table 2: $[A, B, E]$, $[A, C, E]$ and $[B, C, E]$.

Table 2. Total satisfaction for voters in the election

# Voters approving at least	ABC	ABD	ABE	ACE	BCD	BCE
1 candidate	20	20	30	30	20	30
2 candidates	20	20	20	20	20	20
3 candidates	20	0	0	0	0	0
Total Satisfaction	36.66	30	40	40	30	40

Instead of simply approvals, if the voters were to express judgements over the candidates, we could have a judgement profile as reported in Table 3.

Table 3. Distribution of preferences for voters in the election

# Voters	Excellent	Good	Reject
20	A,B	C	D,E
10	E		A, B,C,D

In our multi-winner case, switching to the Majority Judgement voting rule does not resolve the problem. In fact, the final ranking returned from Majority Judgement (using a lexicographic tie-breaking rule) on the whole set of preferences is $A > B > C > E > D$. Using this ranking and choosing the first 3 candidates, the committee would be $[A, B, C]$ which is once again the solution which over-represents one group of voters and under-represents the other. We'll return to this example in the next section.

2.2 Related Work

Multi-winner elections are an interesting area of research that has recently attracted great attention from artificial intelligence and COMSOC researchers. In [14] authors identify two natural classes of selection rules for committees proving that many of the existing rules belong to one or both of these classes. Moreover, a set of desirable properties are derived and formalized. Unfortunately, some of the existing voting rules suffer from high computational cost and thus for these voting rules approximation algorithms are proposed [43]. In [13] the authors use various multiwinner voting rules with preferences that lie in a 2D space to examine what points in the space are picked by various voting rules. Traditionally, research on multiwinner voting rules has focused on properties related to proportionality and representation, including the notion of justified representation [2], but there are other reasons we pick a committee or a set of multiple winners. Group representation and recommendation [42] is an approach studied in [25]: participants rank dishes from the catering menu, then the organizers use a multiwinner voting rule to aggregate these preferences and form the catering order. In such a way at the dinner, participants would be able to choose freely among the catered items. Notice that for some of these approaches there exist good empirical studies [26]. Multi-winner voting rules can be applied to situations as diverse as elections and product recommendations, and that rules satisfy important properties which should be selected based on application domain. In our proposal for extending MJ, we plan to use clustering algorithms from the machine learning literature, k-means [33], k-medoids [24], to discover partitions of voters according to their preferences or other outside features. These could include locality for which we can use a number of different distance metrics between voters including distances on more structured preferences such as CP-nets [30, 31]. Our proposal to leverage clustering is similar to the classic Monroe rule [37, 20] where each candidate has a capacity of voters that they can represent, i.e., a group of voters are "assigned" to some candidate. In [43] authors propose to cluster voters following different heuristics and a satisfaction function in order to approximate Monroe and Chamberlin-Courant rules. Finally, in [12] authors look at the agreement between various multiwinner voting rules and single winner rules to see if there are agreements between the rules, and we plan a similar experiment for our extension.

3 Multi-winner Majority Judgement

The main idea of our multi-winner extension of Majority Judgement is to split voters in groups (or clusters) based on the similarity of their preferences and to allow these groups of voters to elect one or more candidates based on their size, i.e. the number of voters in the group. This is similar to districting or partitioning voters in normal elections, but here we allow this to happen based off the reported preferences of the voters. Informally, given k the number of candidates to be selected for the committee, initially, the algorithm tries to split the set of voters in at most k partitions or clusters. Each group of voters elects its winners using the Majority Judgement voting rule which returns a ranking over the set of candidates. The same candidate cannot be added twice to the list of winners, because each winner can seat only in one place of the committee. At the end of any round, if we have not selected a size k committee, then the procedure is repeated with a new clustering but with a smaller number of clusters, i.e,, given that there are some number of seats k' left over, the *same* algorithm is repeated with k' clusters. If this is not possible then the procedure is repeated with a smaller number of clusters. This allows for new elections with higher number of voters in resulting clusters. It is interesting to notice that the support for these candidates increases with the number of voters in the cluster. This process of reducing the number of groups with respect to the previous step helps to ensure that the candidates selected are representative of larger and larger groups of voters. The algorithm always finishes in a finite number of steps.

3.1 Returning to Our Running Example

Let's now briefly use our proposal that combines Majority Judgement and a clustering algorithm to elect the committee. The clustering algorithm naturally groups voters based on the similarity of their preferences. Remember that the bigger group gets to elect 2 candidates out of 3, while the other group only has sufficient size to choose 1 candidate out of 3. A MJ election over the first group would output the ranking $A > B > C > D > E$ from which the algorithm select the first 2 candidates. A MJ election over the second group would output the ranking $E > A > B > C > D$ from which the algorithm select the first candidate. Thus the final committee would be $[A, B, E]$ that in this case is the only committee which should be elected. At a first glimpse, we can notice that candidate C, even if it has a strong support from the group of 20 people, should not be elected because people's opinions are slightly less supportive of it than A and B.

4 Discussion and Research Challenges

In this paper we have discussed some ways in which research in artificial intelligence, including machine learning and computational social choice, may give us tools to combat political polarization. While we have discussed an extension of Majority Judgement to the multi-winner setting, there are still concrete research challenges that we need to address.

Challenge 1: How do the properties of MJ change when we use clustering to bring down the computational cost? When used in a single winner election, Majority Judgement is particularly appealing because it does not suffer from either Condorcet's or Arrow's Paradox, and it also resists strategic manipulation [4]. Do these properties persist in our proposed approximation?

Challenge 2 Which metrics should be used to measure satisfaction of an electorate? Skowron et al. [43] propose to group voters in order to maximize a satisfaction functions which connected to the preferences of the individual voters. Elections which employ Majority Judgement use judgements to represent users' preferences over candidate. Satisfaction functions should represent how the committee is appreciated by the community.

Challenge 3 What properties should rules have in order to be appealing for policy makers? Axiomatization of voting rules is a broad area of study in computational social choice which tries to characterize algorithms using standard properties [39]. Policy makers may prefer some properties in some scenario while adopting other properties for different applications. When leveraging machine learning, some properties cannot be assured anymore due to the probabilistic approach the algorithms.

Challenge 4 Can randomization be a deterrent to malicious agents? Strategic behaviour requires agents to posses knowledge of the domain in order to change the outcome of the election. Our proposed machine learning techniques are not deterministic. This means that running the same algorithm on the same set of data may result in slightly different outcomes. This may make it unfeasible for an attacker to undermine the reliability of the information gathered and used in the election.

Challenge 5 How do we handle selecting non-polarizing representatives? If the voters are polarized then the selected candidates may be as well. We want to prevent polarization without alienating minority groups by taking clusters according to properties of the voters, and then selecting candidates from each cluster that are most central to that cluster. While this ideally preserves representation and proportionality, it may not be the best way to address polarization. We must experiment with other selection techniques, e.g., selecting the representatives from each cluster that are central w.r.t. all voters.

Challenge 6 How should multi-winner MJ and other voting rules be validated? PrefLib [34] and other libraries and synthetic datasets [19] make election data publicly available on the Web. But it is hard to find a ground truth to compare results about multi winner elections.

References

1. Abramowitz, B., Mattei, N.: Flexible representative democracy: An introduction with binary issues. arXiv preprint arXiv:11811.02921 (2018), http://arxiv.org/abs/1811.02921
2. Aziz, H., Brill, M., Conitzer, V., Elkind, E., Freeman, R., Walsh, T.: Justified representation in approval-based committee voting. Social Choice and Welfare **48**(2), 461–485 (2017)
3. Aziz, H., Gaspers, S., Gudmundsson, J., Mackenzie, S., Mattei, N., Walsh, T.: Computational aspects of multi-winner approval voting. In: Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems. pp. 107–115 (2015)
4. Balinski, M., Laraki, R.: Election by majority judgement: experimental evidence (2007)

5. Balinski, M., Laraki, R.: Majority judgment: measuring, ranking, and electing. MIT press (2011)
6. Balinski, M., Laraki, R.: Majority judgment vs. majority rule. Social Choice and Welfare **54**(2), 429–461 (2020)
7. Barber, M., McCarty, N., Mansbridge, J., Martin, C.J.: Causes and consequences of polarization. Political negotiation: A handbook **37**, 39–43 (2015)
8. Brams, S.J., Fishburn, P.C.: Approval Voting. 2nd edn. (2007)
9. Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A.D. (eds.): Handbook of Computational Social Choice. Cambridge University Press (2016)
10. Brill, M.: Interactive democracy. In: Proceedings of the17thInternational Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS). pp. 1183–1187 (2018)
11. Center, B.: Ungovernable america? the causes and consequences of polarized democracy. In: Jorde Symposium. Accessed February. vol. 14, p. 2015 (2010)
12. Diss, M., Kamwa, E., Tlidi, A.: The chamberlin-courant rule and the k-scoring rules: Agreement and condorcet committee consistency (2018)
13. Elkind, E., Faliszewski, P., Laslier, J., Skowron, P., Slinko, A., Talmon, N.: What do multiwinner voting rules do? an experiment over the two-dimensional euclidean domain. In: Proceedings of AAAI Conference on Artificial Intelligence. pp. 494–501. AAAI Press (2017)
14. Elkind, E., Faliszewski, P., Skowron, P., Slinko, A.: Properties of multiwinner voting rules. Social Choice and Welfare **48**(3), 599–632 (2017)
15. Faliszewski, P., Hemaspaandra, E., Hemaspaandra, L.A.: How hard is bribery in elections? J. Artif. Intell. Res. **35**, 485–532 (2009)
16. Faliszewski, P., Rothe, J.: Control and bribery in voting. In: Handbook of Computational Social Choice (2016)
17. Faliszewski, P., Skowron, P., Slinko, A., Talmon, N.: Multiwinner voting: A new challenge for social choice theory. In: Endriss, U. (ed.) Trends in Computational Social Choice, chap. 2 (2017)
18. Faliszewski, P., Procaccia, A.D.: AI's war on manipulation: Are we winning? AI Mag. **31**(4), 53–64 (2010)
19. Faliszewski, P., Skowron, P., Slinko, A., Szufa, S., Talmon, N.: How similar are two elections? In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 33, pp. 1909–1916 (2019)
20. Faliszewski, P., Slinko, A., Stahl, K., Talmon, N.: Achieving fully proportional representation by clustering voters. Journal of Heuristics **24**(5), 725–756 (2018)
21. Gibbard, A.: Manipulation of voting schemes: a general result. Econometrica: journal of the Econometric Society pp. 587–601 (1973)
22. Green-Armytage, J.: Direct voting and proxy voting. Constitutional Political Economy **26**(2), 190–220 (2015)
23. Hleg, A.: Ethics guidelines for trustworthy AI. B-1049 Brussels (2019)
24. Kaufmann, L.: Clustering by means of medoids. In: Proc. Statistical Data Analysis Based on the L1 Norm Conference, Neuchatel, 1987. pp. 405–416 (1987)
25. Kocot, M., Kolonko, A., Elkind, E., Faliszewski, P., Talmon, N.: Multigoal committee selection. In: Kraus, S. (ed.) Proceedings of the 28th International Joint Conference on Artificial Intelligence, IJCAI 2019. pp. 385–391. ijcai.org (2019)
26. Lackner, M., Skowron, P.: A quantitative analysis of multi-winner rules. In: Kraus, S. (ed.) Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019. pp. 407–413. ijcai.org (2019)
27. Lemennicier, B.: Political polarization and the french rejection of the european constitution. European Journal of Political Economy **21**(4), 1077–1084 (2005)

28. Loreggia, A., Mattei, N., Rossi, F., Venable, K.B.: Preferences and ethical principles in decision making. In: Proceedings of the 1st AAAI/ACM Conference on AI, Ethics, and Society (AIES) (2018)
29. Loreggia, A., Narodytska, N., Rossi, F., Venable, K., Walsh, T.: Controlling elections by replacing candidates: Theoretical and experimental results. vol. WS-14-10, pp. 61–66 (2014)
30. Loreggia, A., Mattei, N., Rossi, F., Venable, K.B.: On the distance between cp-nets. In: Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems. pp. 955–963. International Foundation for Autonomous Agents and Multiagent Systems (2018)
31. Loreggia, A., Mattei, N., Rossi, F., Venable, K.B.: CPMetric: Deep siamese networks for metric learning on structured preferences. In: Artificial Intelligence. IJCAI 2019 International Workshops. pp. 217–234. Springer (2020)
32. Loreggia, A., Narodytska, N., Rossi, F., Venable, K.B., Walsh, T.: Controlling elections by replacing candidates or votes. In: Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems. pp. 1737–1738 (2015)
33. MacQueen, J., et al.: Some methods for classification and analysis of multivariate observations. In: Proceedings of the fifth Berkeley symposium on mathematical statistics and probability. vol. 1, pp. 281–297. Oakland, CA, USA (1967)
34. Mattei, N., Walsh, T.: PrefLib: A library for preferences, HTTP://WWW.PREFLIB.ORG. In: Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT) (2013)
35. Mattei, N.: Closing the loop: Bringing humans into empirical computational social choice and preference reasoning. In: Bessiere, C. (ed.) Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020. pp. 5169–5173. ijcai.org (2020). https://doi.org/10.24963/ijcai.2020/729
36. McCoy, J., Rahman, T., Somer, M.: Polarization and the global crisis of democracy: Common patterns, dynamics, and pernicious consequences for democratic polities. American Behavioral Scientist **62**(1), 16–42 (2018)
37. Monroe, B.L.: Fully proportional representation. American Political Science Review pp. 925–940 (1995)
38. Quintarelli, S., Corea, F., Fossa, F., Loreggia, A., Sapienza, S.: AI: profili etici. una prospettiva etica sull'intelligenza artificiale: princìpi, diritti e raccomandazioni. BioLaw Journal-Rivista di BioDiritto **18**(3), 183–204 (2019)
39. Rossi, F., Venable, K., Walsh, T.: A Short Introduction to Preferences: Between Artificial Intelligence and Social Choice. Morgan and Claypool (2011)
40. Sen, A.: Collective Choice and Social Welfare. Harvard University Press (2018)
41. Shneiderman, B.: Human-centered artificial intelligence: Reliable, safe & trustworthy. International Journal of Human–Computer Interaction **36**(6), 495–504 (2020)
42. Skowron, P., Faliszewski, P., Lang, J.: Finding a collective set of items: From proportional multirepresentation to group recommendation. Artif. Intell. **241**, 191–216 (2016). https://doi.org/10.1016/j.artint.2016.09.003
43. Skowron, P., Faliszewski, P., Slinko, A.M.: Achieving fully proportional representation: Approximability results. Artif. Intell. **222**, 67–103 (2015). https://doi.org/10.1016/j.artint.2015.01.003
44. Thiele, T.N.: Om flerfoldsvalg. Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandlinger **1895**, 415–441 (1895)
45. Zwicker, W.S.: Introduction to the theory of voting. In: Handbook of Computational Social Choice, pp. 23–56. Cambridge University Press (2016). https://doi.org/10.1017/CBO9781107446984.003