Broadening the Research Agenda for Computational Social Choice: Multiple Preference Profiles and Multiple Solutions

Blue Sky Ideas Track

Niclas Boehmer
Algorithmics and Computational Complexity
TU Berlin
niclas.boehmer@tu-berlin.de

Rolf Niedermeier
Algorithmics and Computational Complexity
TU Berlin
rolf.niedermeier@tu-berlin.de

ABSTRACT

The area of computational social choice (COMSOC) analyzes collective decision problems from an algorithmic perspective. So far, the main focus in this area lied on analyzing problems where a single preference relation for each agent is given and a single solution reflecting all agents' preferences needs to be found. However, this modeling is often not rich enough to capture the changing and ambivalent nature of real-world problems. We will argue that one possibility to incorporate such aspects is to allow for multiple preference profiles in the input and multiple solutions in the output. We systematically review different types of arising settings, point out how classical problems and solution concepts can be generalized, and identify several research challenges.

KEYWORDS

Collective decision making; dynamic problems; multimodality; repeated decision making; temporal aspects

ACM Reference Format:
Niclas Boehmer and Rolf Niedermeier. 2021. Broadening the Research Agenda for Computational Social Choice: Multiple Preference Profiles and Multiple Solutions: Blue Sky Ideas Track. In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), Online, May 3–7, 2021*, IFAAMAS, 5 pages.

1 INTRODUCTION

Abstractly speaking, in a collective decision problem, given a set of agents and a profile containing the preferences of each agent over alternatives, the goal is to aggregate the preferences into a compromise solution. Collective decision problems range from voting over coalition formation to fair allocation. Applications are multifaceted such as fairly distributing tasks, matching students to colleges, or selecting which results to display to a user querying an online database or search engine.

In the majority of publications dealing with collective decision problems, it is assumed that each agent has a single and fixed preference relation, and the goal is to find a single solution. However, this modeling disregards important aspects of many real-world problems: Agents, in particular humans, typically have multi-criteria preferences over alternatives that may change over time, and agents may repeatedly participate in the same collective decision problem. To be able to model such aspects, one possibility is to relax the classical paradigm with only a single preference profile as input for which a single solution needs to be found.[1] Thereby, we allow for multiple preference profiles (profiles for short) in the input and/or for multiple separate solutions in the output. This makes it, among others, possible to incorporate two important aspects of real-world collective decision problems: time and multimodality. For instance, if multiple preference profiles are given, then they may have some time-based ordering or may reflect the preferences of agents with respect to different evaluation criteria.

We propose a taxonomy of collective decision problems with multiple profiles in the input and/or multiple solutions in the output. Hence, three natural *settings* arise:

- **(MO)** Given **m**ultiple profiles, the goal is to find **o**ne solution for all profiles;
- **(OM)** given **o**ne profile, the goal is to find **m**ultiple solutions for this profile; and
- **(MM)** given **m**ultiple profiles, the goal is to find **m**ultiple solutions (one solution for each profile).

For each of these three settings, one can distinguish whether the given profiles and computed solutions have some natural (e.g., time-based) ordering or not. As this distinction is conceptually important, we split each setting into an ordered and an unordered version, and refer to the six resulting cases as *subsettings*. An overview of our three settings and six subsettings is shown in Figure 1, which also provides a roadmap for the remainder of the paper.

Running problem. We consider the task of fairly allocating a set of resources to a set of agents who have preferences over subsets (bundles) of resources. A notion of fairness here is envy-freeness: An assignment of resources to agents is *envy-free* if no agent prefers another agent's bundle to her own. Finding an envy-free allocation of resources to agents will be our running problem in the following. However, we want to emphasize that most of our ideas and concepts are applicable to a variety of collective decision problems.

Outline. In Section 2, we start by presenting some ideas how classical solution concepts such as envy-freeness can be generalized if multiple profiles are given and multiple solutions need to be found. Subsequently, we present the six subsettings (see Figure 1) in Sections 3 to 5. For each of them, we start with a pictogram visualizing the situation. On the top of each pictogram, the structure of the solutions to be computed is visualized (each puzzle piece represents one solution). On the bottom, the structure of the given

Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), U. Endriss, A. Nowé, F. Dignum, A. Lomuscio (eds.), May 3–7, 2021, Online.

[1] There also exist alternative approaches to model dynamic aspects of real-world problems that we do not cover here. For instance, it is possible to assume that an instance arrives part by part and decisions need to be made on the fly (see, e.g., [2, 28]).

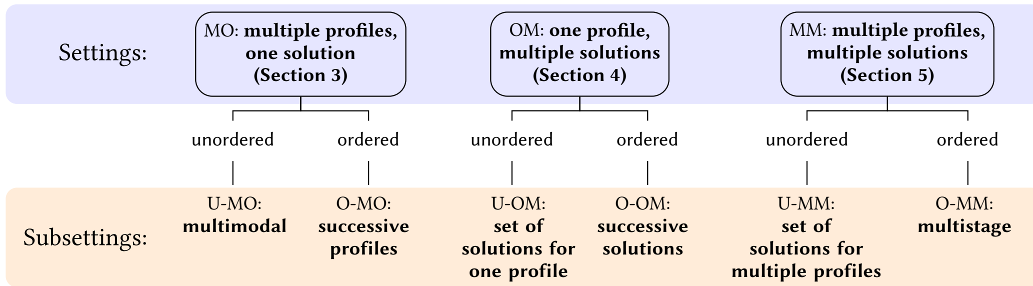

Figure 1: Overview of the three settings and the six subsettings we consider.

profiles is visualized (each group of persons represents one profile). The pictogram is accompanied by a brief description under what circumstances this subsetting can arise in our running problem. Subsequently, we briefly describe related work and point to the specific and most interesting features of this subsetting and give some ideas how to approach them.

2 ON GENERALIZING CLASSICAL SOLUTION CONCEPTS

A straightforward possibility to apply a classical solution concept like envy-freeness to multiple profiles or multiple solutions is to simply require that each solution has to fulfill the concept in all profiles for which the solution is implemented. We refer to the resulting concepts as *local solution concepts*. In contrast to this, it is also possible to apply solution concepts globally, that is, to assess all (or at least some) computed solutions and given profiles together. Global concepts open the possibility to model long-time fairness aspects or to relax classical concepts. On a high level, we distinguish between two families of *global solution concepts*:

Profile-level concepts Solutions are evaluated with respect to profiles, i.e., solutions satisfy certain properties in all profiles from some specific subset of profiles.

Agent-level concepts Solutions are evaluated with respect to each agent separately, i.e., solutions have to satisfy certain properties from the perspective of each agent.

For instance, a profile-level generalization of envy-freeness in the subsetting "successive profiles" (a sequence of profiles is given for which a single solution should be found) is to enforce that the computed solution is envy-free at least once every k profiles. In contrast, an agent-level generalization is to require that each agent is envy-free at least once every k profiles.

3 MULTIPLE PROFILES, ONE SOLUTION (MO)

Problems where multiple profiles are given and one solution needs to be found arise naturally if the agents have multi-criteria preferences or the preferences of agents change over time.

3.1 *Unordered:* Multimodal (U-MO)

Running problem: Agents might assess resources based on different, independent evaluation criteria (such as appearance, usefulness, sales value) and may cast separate preferences for each criterion. Alternatively, in decisions under uncertainty, each preference profile

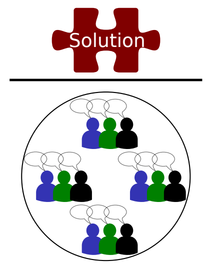

may represent the preferences of agents in some precomputed scenario for the future.

Multimodal preferences have been recently considered in COMSOC by Chen et al. [14] in the context of Stable Marriage. Subsequently, Steindl and Zehavi [44] studied the problem of allocating houses and Jain and Talmon [27] analyzed committee elections in the presence of multimodal preferences.

One goal in this subsetting is to find a solution that satisfies certain properties (such as envy-freeness) in a given number k of the profiles, which results in a profile-level solution concept. Instead, it is also possible to impose such conditions on the agent-level and only require that a solution makes every agent "happy" in at least k of the given profiles or at least that an agent is not "unhappy" in more than k profiles for the same reason [14]. In the context of finding an envy-free allocation, this would mean that each agent does not envy another agent in more than k of the profiles or does not envy the same agent in more than k profiles.

A fundamentally different approach is to assume that agents only care about the best/worst/average quality of the solution in the preference relations of the agent [27]. Consequently, the goal could be to search for a solution maximizing the minimum/average of the resulting agent's valuations of the solution.

3.2 *Ordered:* Successive Profiles (O-MO)

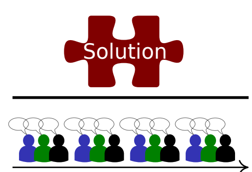

Running problem: The demand of an agent might be predictable but can still be quite different for different months. However, some kinds of resources (such as manpower) cannot be easily reallocated. Thus, a single assignment for multiple months needs to be found.

The problem of finding a single solution for a sequence of two or more profiles has, to the best of our knowledge, not been considered in COMSOC so far. Thus, novel solution concepts need to be developed. One promising approach here is to use a sliding-window technique and only require that the solution fulfills certain properties in each size-k subset of succeeding profiles.[2] In the context of

[2]A similar temporal relaxation technique has been already applied in the context of temporal graphs (graphs with a changing edge set) to classical graph problems such as Clique [45], Coloring [37], and Vertex Cover [1].

finding an envy-free allocation, one may want to find an allocation that is envy-free in a majority of profiles from each size-k subset of succeeding profiles. As in the multimodal subsetting (U-MO), it is also possible to impose such global conditions on the agent-level, that is, to find an allocation such that each agent is envy-free in a majority of profiles from each size-k subset of succeeding profiles or an allocation where each agent envies each other agent only in a minority of profiles from each size-k subset of succeeding profiles. A natural motivation for this relaxation technique is that agents and, in particular, humans tend to like the status-quo and are rather reluctant to change, that is, they need to be dissatisfied for k consecutive steps before a change may be requested.

4 ONE PROFILE, MULTIPLE SOLUTIONS (OM)

Sometimes the same set of agents with fixed preferences has to make a collective decision repeatedly. Note that if multiple solutions are considered, it is possible to impose restrictions on the usage of objects (e.g., resources or candidates) in the solutions. For instance, each object may only be part of a given number of solutions.

4.1 *Unordered:* Set of Solutions for One Profile (U-OM)

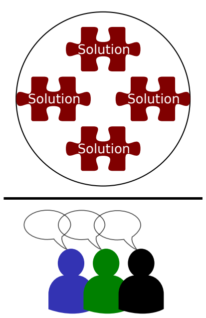

Running problem: In most situations, it is nearly impossible for decision makers to evaluate all solutions. Instead, they would like to be presented a set of high-quality solutions, reflecting all solutions in order to pick one solution from it which will be implemented.

Finding a set of solutions for a single profile is useful in different contexts with implications on which solutions one wants to put in the set. For instance, if the goal is to pick a single solution to be implemented from the set, a popular general approach is to find a diverse set of solutions, that is, a set of optimal or close to optimal solutions with maximum summed pairwise difference. This problem has been already considered in the context of graph algorithms [4] and constraint programming [42] but not in COMSOC. An interesting open (problem-specific) question here is how diversity is measured.

Alternatively, one may want to find a set of solutions to be all implemented simultaneously. In this case, one goal could be to make everyone “happy” to a “fair” extent. That is, the set of solutions proportionally represents the agents or, simpler, every agent is “happy” in at least a given number of solutions. The problem of finding a set of proportionally representative winners of a single-winner election has been already (extensively) studied in the context of multi-winner elections [13, 17, 32]. Recently, voting rules from approval-based multi-winner voting were used to find a proportionally representative set of matchings of agents having approval preferences over each other [7]. More generally, the quite well understood concepts of proportional representation and diversity developed in the context of multi-winner voting form a natural starting point for finding a set of solutions also for other collective decision problems.

4.2 *Ordered:* Successive Solutions (O-OM)

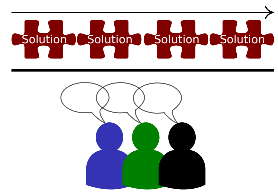

Running problem: Under certain circumstances, the preferences of agents do not change over multiple months. However, to treat agents as fairly as possible, different assignments in each month may be selected.

Again, this subsetting is applicable in different contexts. For instance, one may want to find a sequence of solutions to be implemented one after each other. In COMSOC, the only work in this direction we are aware of is due to Bredereck et al. [12] who studied finding a sequence of committees for a multi-winner election. It is possible to explore different solution concepts. If, for instance, no single solution satisfying everyone exists, then one can instead search for a sequence of solutions such that everyone is satisfied at least once every k steps (agent-level concept). It is also natural to combine such solution concepts with availability constraints, for instance, restricting that each resource can only be included in some given number of solutions or that resources need some recovery time before being included in a solution again. Additionally, depending on the application, it might also be important to impose that successive solutions need to fulfill some properties, for instance, overlap or differ to a certain extent. If such constraints are imposed, then it is also natural to consider profile-level solution concepts, for example, enforcing that the sequence of solutions has to fulfill some given property once every k steps, or local solution concepts, for example, imposing that every solution makes half of the agents “happy”.

An alternative interpretation of the computed ordered sequence of solutions, which calls for fundamentally different solution concepts, is that they provide a ranking of solutions. Already studied by Kemeny [29] and Mallows [35], this problem has a comparably long history in COMSOC. Again, different applications requiring different solution concepts need to be distinguished. For instance, Skowron et al. [43] studied finding proportional rankings in contexts where initial segments of the ranking of different lengths are relevant. Further, Lu and Boutilier [33] studied finding a consensus ranking to cope with situations where only a single candidate is selected in the end but candidates may become unavailable.

5 MULTIPLE PROFILES, MULTIPLE SOLUTIONS (MM)

The settings in Sections 3 and 4 (MO and OM) can be considered as special cases of the setting of this section when either the given profiles or the computed solutions are identical.

5.1 *Unordered:* Set of Solutions for Multiple Profiles (U-MM)

Running problem: Assuming that each agent represents different groups of individuals with contradicting preferences (multiple profiles; one for each group), in certain applications decision makers may want to assign resources specifically to each of the groups (multiple solutions; one for each group).

Here, we are given a set of profiles and the goal is to find a separate solution for each of them, thereby reflecting desirable

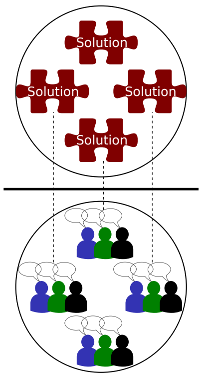

overall constraints. To the best of our knowledge, this subsetting has not yet been studied in COMSOC. A (loosely) related work is a recent paper by Boehmer et al. [6] dealing with finding solutions to multiple single-winner elections where every candidate is only allowed to win one of the elections. This corresponds to the presented subsetting if every object (e.g., resources or candidates) is only allowed to be included in one solution.

From a conceptual point of view, this subsetting is closely related to the subsetting "set of solutions for one profile" (U-OM). However, because agents may have different preferences in the different profiles here, computational and axiomatic questions are a lot harder to tackle. It is nevertheless interesting to come up with proportionality and fairness axioms that can be applied. A starting point could be to search for a set of solutions such that each agent is "happy" in at least k of them in the corresponding profile.

5.2 *Ordered:* Multistage (O-MM)

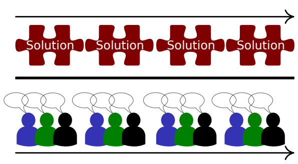

Running problem: The demand of an agents may be quite different each month. To cope with this, decision makers decide on a new distribution every month.

The multistage view on combinatorial optimization problems has been introduced by Eisenstat et al. [16] and Gupta et al. [24] and has attracted quite some follow-up work, especially in the graph algorithms community [3, 20, 25]. Bredereck et al. [11] introduced the multistage view to COMSOC by studying a series of Plurality multi-winner elections. Their goal was to find a sequence of committees, each committee being approved by x voters in the corresponding election such that two consecutive committees overlap in at least/most ℓ candidates. As this subsetting can be understood as a direct generalization of the subsetting "successive profiles" (O-MO) and the subsetting "successive solutions" (O-OM), a natural starting point here are ideas and methods described for the other two subsettings.

This subsetting (in contrast to the other presented subsettings) also admits a natural and meaningful online variant, which has already been addressed in the literature before: Not all profiles may be initially known and it may be necessary to already fix the solutions for some profiles before other profiles arrive. In the resulting online version of this subsetting, the goal is typically to guarantee some long-term fairness or maximize some other objective which takes all profiles and solutions together into account [21, 31].

6 CONCLUSION AND OUTLOOK

Over the last twenty years, COMSOC has become an established research area and the community has made tremendous progress on understanding (algorithmic aspects of) various collective decision problems [9, 19]. However, so far, apart from some notable exceptions (e.g., [5, 22, 39, 40]), the impact of most of the theoretical work on the real world remained rather limited. That is why the future success of the field may significantly depend on the applicability of the considered problems and mechanisms [23]. One particular challenge here is that real-world problems often need a complex modeling (agents may have ambivalent and changing preferences and multiple solutions fulfilling demanding properties may need to be found). Such facets and considerations can often not be captured by traditional models. To allow for modeling them, we propose to generalize classical collective decision problems to settings where multiple profiles are given for which multiple solutions need to be found.

We presented a taxonomy of subsettings naturally arising in this context. This taxonomy offers a unified view which has several advantages. First of all, it highlights the (close) connections of the different subsettings, for instance, the symmetry of the two subsettings "successive profiles" (O-MO) and "successive solutions" (O-OM). These connections are interesting because they can help us to better understand the different subsettings and their multifaceted relationships. In addition to that, our unified view also underlines that some general considerations and ideas are relevant in multiple subsettings, for instance, availability constraints, enforcing that succeeding solutions are similar/different, or distinguishing between employing solution concepts on the profile- or agent-level. This allows one to study the same ideas not only in different collective decision problems but also in the context of different subsettings. Such a broader unified study may contribute to a true and wider applicable understanding of the power and properties of such ideas.

Finally, our taxonomy allows a clean separation, relation, and classification of different models. This can help to identify open and overlooked problems and subsettings, to draw inspiration from previous work in other subsettings, and also to model a real-world problem. Hopefully, at some point, works in the different subsettings will have created a rich and easily configurable toolbox of well-understood solution concepts, axiomatic properties, and algorithmic techniques to choose from in the quest for dealing with complex aspects of real-world collective decision problems.

A similar development towards studying more expressive models has also already taken place in the algorithmic graph community, for instance, in the context of studying multilayer [30, 34] and temporal graphs [26, 38]. General techniques and concepts developed there might be also applicable to collective decision problems.

Lastly, we want to highlight two specific challenges that might arise in the study of settings and subsettings from our taxonomy. Firstly, we expect most classical algorithmic problems to become computationally worst-case intractable if generalized to multiple profiles and/or solutions. Thus, it is important to consider different paths to tractability such as preference restrictions [18] or more advanced algorithmic techniques such as approximation [36] and parameterized algorithmics [10, 15]. Secondly, in the presence of multiple preference profiles, there exists a variety of possibilities how to model the agents' preferences, which also impact the applicability and tractability of the resulting models. Some first steps in this direction have already been made by Boutilier and Procaccia [8] and Parkes and Procaccia [41] by modeling preferences as Markov decision processes.

ACKNOWLEDGMENTS

NB was supported by the DFG project MaMu (NI 369/19).

REFERENCES

[1] Eleni C. Akrida, George B. Mertzios, Paul G. Spirakis, and Viktor Zamaraev. 2020. Temporal vertex cover with a sliding time window. *J. Comput. Syst. Sci.* 107 (2020), 108–123.
[2] Martin Aleksandrov and Toby Walsh. 2020. Online Fair Division: A Survey. In *Proc. of AAAI '20*. AAAI Press, 13557–13562.
[3] Evripidis Bampis, Bruno Escoffier, Michael Lampis, and Vangelis Th. Paschos. 2018. Multistage Matchings. In *Proc. of SWAT '18*. Schloss Dagstuhl, 7:1–7:13.
[4] Julien Baste, Michael R. Fellows, Lars Jaffke, Tomás Masarík, Mateus de Oliveira Oliveira, Geevarghese Philip, and Frances A. Rosamond. 2020. Diversity of Solutions: An Exploration Through the Lens of Fixed-Parameter Tractability Theory. In *Proc. of IJCAI '20*. ijcai.org, 1119–1125.
[5] Péter Biró. 2017. Applications of Matching Models under Preferences. See [19], 345–373.
[6] Niclas Boehmer, Robert Bredereck, Piotr Faliszewski, Andrzej Kaczmarczyk, and Rolf Niedermeier. 2020. Line-Up Elections: Parallel Voting with Shared Candidate Pool. In *Proc. of SAGT '20*. Springer, 275–290.
[7] Niclas Boehmer, Markus Brill, and Ulrike Schmidt-Kraepelin. 2021. Selecting Matchings via Multiwinner Voting: How Structure Defeats a Large Candidate Space. *CoRR* abs/2102.07441 (2021). arXiv:2102.07441 https://arxiv.org/abs/2102.07441
[8] Craig Boutilier and Ariel D. Procaccia. 2012. A Dynamic Rationalization of Distance Rationalizability. In *Proc. of AAAI '12*. AAAI Press, 1278–1284.
[9] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). 2016. *Handbook of Computational Social Choice*. Cambridge University Press.
[10] Robert Bredereck, Jiehua Chen, Piotr Faliszewski, Jiong Guo, Rolf Niedermeier, and Gerhard J. Woeginger. 2014. Parameterized algorithmics for computational social choice: Nine research challenges. *Tsinghua Science and Technology* 19, 4 (2014), 358–373.
[11] Robert Bredereck, Till Fluschnik, and Andrzej Kaczmarczyk. 2020. Multistage Committee Election. *CoRR* abs/2005.02300 (2020). arXiv:2005.02300 https://arxiv.org/abs/2005.02300
[12] Robert Bredereck, Andrzej Kaczmarczyk, and Rolf Niedermeier. 2020. Electing Successive Committees: Complexity and Algorithms. In *Proc. of AAAI '20*. AAAI Press, 1846–1853.
[13] Markus Brill, Paul Gölz, Dominik Peters, Ulrike Schmidt-Kraepelin, and Kai Wilker. 2020. Approval-Based Apportionment. In *Proc. of AAAI '20*. AAAI Press, 1854–1861.
[14] Jiehua Chen, Rolf Niedermeier, and Piotr Skowron. 2018. Stable Marriage with Multi-Modal Preferences. In *Proc. of EC '18*. ACM, 269–286.
[15] Britta Dorn and Ildikó Schlotter. 2017. Having a Hard Time? Explore Parameterized Complexity! See [19], 209–230.
[16] David Eisenstat, Claire Mathieu, and Nicolas Schabanel. 2014. Facility Location in Evolving Metrics. In *Proc. of ICALP '14*. Springer, 459–470.
[17] Edith Elkind, Piotr Faliszewski, Piotr Skowron, and Arkadii Slinko. 2017. Properties of multiwinner voting rules. *Soc. Choice Welf.* 48, 3 (2017), 599–632.
[18] Edith Elkind, Martin Lackner, and Dominik Peters. 2017. Structured Preferences. See [19], 187–207.
[19] Ulle Endriss (Ed.). 2017. *Trends in Computational Social Choice*. AI Access Foundations.
[20] Till Fluschnik, Rolf Niedermeier, Valentin Rohm, and Philipp Zschoche. 2019. Multistage Vertex Cover. In *Proc. of IPEC '19*. Schloss Dagstuhl, 14:1–14:14.
[21] Rupert Freeman, Seyed Majid Zahedi, and Vincent Conitzer. 2017. Fair and Efficient Social Choice in Dynamic Settings. In *Proc. of IJCAI '17*. ijcai.org, 4580–4587.
[22] Jonathan R. Goldman and Ariel D. Procaccia. 2014. Spliddit: unleashing fair division algorithms. *SIGecom Exch.* 13, 2 (2014), 41–46.
[23] Umberto Grandi, Stefan Napel, Rolf Niedermeier, and Kristen Brent Venable. 2020. Application-oriented computational social choice (Dagstuhl seminar 19381). *Dagstuhl Reports* 9, 9 (2020), 45–65.
[24] Anupam Gupta, Kunal Talwar, and Udi Wieder. 2014. Changing Bases: Multistage Optimization for Matroids and Matchings. In *Proc. of ICALP '14*. Springer, 563–575.
[25] Klaus Heeger, Anne-Sophie Himmel, Frank Kammer, Rolf Niedermeier, Malte Renken, and Andrej Sajenko. 2019. Multistage Problems on a Global Budget. *CoRR* abs/1912.04392 (2019). arXiv:1912.04392 http://arxiv.org/abs/1912.04392
[26] Petter Holme and Jari Saramäki (Eds.). 2019. *Temporal Network Theory*. Springer.
[27] Pallavi Jain and Nimrod Talmon. 2020. Committee Selection with Multimodal Preferences. In *Proc. of ECAI '20*. IOS Press, 123–130.
[28] Ian A. Kash, Ariel D. Procaccia, and Nisarg Shah. 2014. No agent left behind: dynamic fair division of multiple resources. *J. Artif. Intell. Res.* 51 (2014), 579–603.
[29] John G Kemeny. 1959. Mathematics without numbers. *Daedalus* 88, 4 (1959), 577–591.
[30] Mikko Kivelä, Alex Arenas, Marc Barthelemy, James P. Gleeson, Yamir Moreno, and Mason A. Porter. 2014. Multilayer networks. *J. Complex Networks* 2, 3 (2014), 203–271.
[31] Martin Lackner. 2020. Perpetual Voting: Fairness in Long-Term Decision Making. In *Proc. of AAAI '20*. AAAI Press, 2103–2110.
[32] Martin Lackner and Piotr Skowron. 2020. Approval-Based Committee Voting: Axioms, Algorithms, and Applications. *CoRR* abs/2007.01795 (2020). arXiv:2007.01795 https://arxiv.org/abs/2007.01795
[33] Tyler Lu and Craig Boutilier. 2010. The unavailable candidate model: a decision-theoretic view of social choice. In *Proc. of EC '10*. ACM, 263–274.
[34] Matteo Magnani and Luca Rossi. 2011. The ML-Model for Multi-layer Social Networks. In *Proc. of ASONAM '11*. IEEE Computer Society, 5–12.
[35] Colin L Mallows. 1957. Non-null ranking models. *Biometrika* 44 (1957), 114–130.
[36] Evangelos Markakis. 2017. Approximation Algorithms and Hardness Results for Fair Division with Indivisible Goods. See [19], 231–247.
[37] George B. Mertzios, Hendrik Molter, and Viktor Zamaraev. 2019. Sliding Window Temporal Graph Coloring. In *Proc. of AAAI '19*. AAAI Press, 7667–7674.
[38] Othon Michail. 2016. An introduction to temporal graphs: an algorithmic perspective. *Internet Math.* 12, 4 (2016), 239–280.
[39] Website of RoboVote. 2016. http://www.robovote.org/. Accessed: 20 November 2020.
[40] Website of the Matching in Practice network. 2014. http://www.matching-in-practice.eu/. Accessed: 20 November 2020.
[41] David C. Parkes and Ariel D. Procaccia. 2013. Dynamic Social Choice with Evolving Preferences. In *Proc. of AAAI '13*. AAAI Press, 767–773.
[42] Thierry Petit and Andrew C. Trapp. 2015. Finding Diverse Solutions of High Quality to Constraint Optimization Problems. In *Proc. of IJCAI '15*. AAAI Press, 260–267.
[43] Piotr Skowron, Martin Lackner, Markus Brill, Dominik Peters, and Edith Elkind. 2017. Proportional Rankings. In *Proc. of IJCAI '17*. ijcai.org, 409–415.
[44] Barak Steindl and Meirav Zehavi. 2020. Parameterized Analysis of Assignment Under Multiple Preferences. *CoRR* abs/2004.00655 (2020). arXiv:2004.00655 https://arxiv.org/abs/2004.00655
[45] Tiphaine Viard, Matthieu Latapy, and Clémence Magnien. 2016. Computing maximal cliques in link streams. *Theor. Comput. Sci.* 609 (2016), 245–252.