Interactive Democracy

Blue Sky Ideas Track

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

ABSTRACT

Interactive Democracy is an umbrella term that encompasses a variety of approaches to make collective decision making processes more engaging and responsive. A common goal of these approaches is to utilize modern information technology—in particular, the Internet—in order to enable more interactive decision making processes. An integral part of many interactive democracy proposals are online decision platforms that provide much more flexibility and interaction possibilities than traditional democratic systems. This is achieved by embracing the novel paradigm of delegative voting, often referred to as *liquid democracy*, which aims to reconcile the idealistic appeal of direct democracy with the practicality of representative democracy. The successful design of interactive democracy systems presents a multidisciplinary research challenge; one important aspect concerns the elicitation and aggregation of *preferences*. However, existing proposals are mostly disconnected from the vast body of scientific literature on preference aggregation and related topics. In this article, I argue that tools and techniques developed in the multiagent systems literature should be employed to aid the design of online decision platforms and other interactive democracy systems. Insights from *computational social choice*, an emerging research area at the intersection of computer science and economics, will be particularly relevant for this endeavor.

KEYWORDS

preference aggregation; group decision making; liquid democracy; e-democracy; preference handling; computational social choice; algorithmic decision theory; delegative voting; proportional representation; crowdsourcing; participatory budgeting

ACM Reference Format:
Markus Brill. 2018. Interactive Democracy. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018,* IFAAMAS, 5 pages.

1 INTRODUCTION

In her 2014 TED talk *How to upgrade democracy for the Internet era*, Pia Mancini poignantly states that "we are 21st-century citizens, doing our very, very best to interact with 19th century-designed institutions that are based on an information technology of the 15th century" [47]. Mancini goes on to observe that the way democratic societies make collective decisions is highly outdated. This leads to the question: "If Internet is the new printing press, then what is democracy for the Internet era?" Mancini and her collaborators approached this question by developing an app, *DemocracyOS* [48], that allows users to propose, debate, and vote on issues. *DemocracyOS* is only one example of a quickly growing number of approaches that aim to reconcile established democratic processes with the desire of citizens to participate in political decision making.[1] Another example is the software *LiquidFeedback* [6], which is developed by the *Association for Interactive Democracy.*[2] Currently, these tools are mainly used for decision making within progressive political parties [9, p. 162] or in the context of community engagement platforms such as *WeGovNow* [10]. A common goal of these approaches, often summarized under the umbrella term *Interactive Democracy*[3] (henceforth ID), is to utilize modern information technology—in particular, the Internet—in order to enable more interactive decision making processes.

When designing a platform for interactive collective decision making, there are lots of design decisions to be made, regarding, for example, *issue selection* (which issues are considered?), *option generation* (which options are on the ballot?), *interaction opportunities* (how is deliberation and delegation organized?), *ballot structure* (in which format can participants express their preferences?), and *aggregation methods* (which method is used to tally the votes?). There is no shortage of concrete suggestions of how ID platforms could be implemented (see the blog post by Ford [34] for an overview). Most of these suggestions, however, are rather ad hoc in nature and little attention is devoted to a principled comparison and evaluation of methods. This increases the risk of employing methods with unintended flaws.

In this article, I argue that concepts and techniques from the multiagent systems literature—particularly those dealing with *preferences* and their aggregation—should be employed to aid the design of online decision platforms and other ID tools. *Computational social choice* (COMSOC), an interdisciplinary subfield at the intersection of economics and computer science, seems to be particularly relevant in this endeavor.[4] Even though research in COMSOC has made tremendous progress in recent years [15, 31], the practical impact of the field has remained rather limited.[5] This is partly due to the fact that many of the rather sophisticated preference handling and preference aggregation mechanisms that are routinely studied

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden.

[1] *DemocracyOS* has since been superseded by *Sovereign*, developed by the *Democracy Earth Foundation* (http://www.democracy.earth).
[2] http://www.interaktive-demokratie.org/index.en.html
[3] The field is lacking a unified terminology. For example, Interactive Democracy is sometimes referred to as *iDemocracy* [18] or *participatory democracy* [2]. The terms *liquid democracy* and *delegative democracy* usually refer to the paradigm of delegative voting (see Section 2). And terms like *e-democracy* [60], *digital democracy* [41], and *Internet democracy* [49] emphasize the role of information technology.
[4] Strongly related and equally relevant fields are *algorithmic decision theory* and *preference handling* [13, 26]. The focus on COMSOC in the current article is due to the author's background and should not be interpreted as a claim to superior relevance.
[5] Notable exceptions are websites like *Spliddit* [37] and *RoboVote* (http://robovote.org).

in the COMSOC literature, though superior in theory, are rarely used in practice. In this article, I argue that the novel application area of interactive democracy has the potential to change that.

In the following, I provide examples of challenges that are encountered when building ID systems, together with pointers to tools and techniques from the COMSOC literature that could be employed to tackle these challenges. Section 2 focusses on issues related to the paradigm of delegative voting (aka *liquid democracy*), and Section 3 describes other challenges. Section 4 concludes.

2 DELEGATIVE VOTING

Participants of online decision platforms can often choose whether they want to vote directly on a particular issue or whether they want to *delegate* their vote to somebody they trust. Delegations can be specified either on an issue-by-issue basis, for whole topic areas, or even globally. Crucially, delegations are transitive and decisions whether to vote directly, to delegate, or to abstain can be changed at any time.[6] This paradigm of delegative voting, which is often referred to as *liquid democracy*, aims to reconcile the idealistic appeal of *direct democracy* (where every voter votes directly on every issue) with the practicality of *representative democracy* (where voters vote for delegates, who then vote on the voters' behalf on all issues) by giving voters the opportunity to have their say on all issues, but not requiring them to get informed on each issue.[7]

Delegative voting gives rise to several novel questions in voting theory. For instance, in order to successfully implement a delegative voting infrastructure, one has to think about potential problems such as

(1) *delegation cycles* (voter 1 delegates her vote to voter 2, who delegates to voter 3, who delegates back to voter 1),
(2) *abstentions* (what if one delegates to somebody who abstains from voting?), and
(3) *inconsistent outcomes* (what if different delegations for different issues lead to a globally incompatible set of decisions?).

The remainder of this section contains approaches to tackle these problems, and a discussion of strategic considerations that delegative voting gives rise to.

2.1 Ranked Delegation

An intuitively appealing approach to address problems (1) and (2) lets voters specify "backup" delegations for the case that their preferred delegation leads to a cycle or an abstaining voter.[8] Taking this idea one step further, we can ask each voter to specify a *ranked list* of delegations. Given a collection of such delegation lists, however, it is not clear how to resolve certain kinds of voting situations. For example, say that voter 1 delegates to voters 2, 3, and 4 (in this order) and that voter 2 delegates to voters 1 and 5 (in this order);

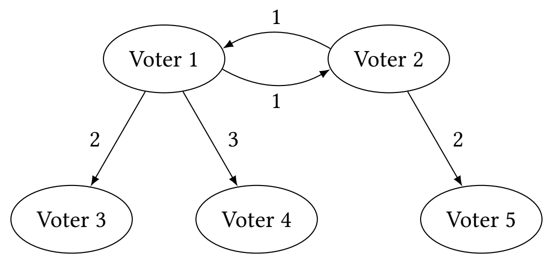

Figure 1: Illustration of a voting situation using ranked delegations. Labels of outgoing edges from a node correspond to ranks in the delegation list of the respective voter.

see Figure 1 for an illustration. Using the first (most preferred) delegation in each ranked list results in a delegation cycle involving voters 1 and 2. Therefore, we have to "go down" in at least one of the lists. Taking the second option in voter 1's list would result in both voters 1 and 2 (by transitivity) delegating to voter 3. On the other hand, taking the second option in voter 2's list would result in both voters 1 and 2 delegating to voter 5. A priori, no solution seems preferable to the other one, suggesting perhaps that one delegation should go to voter 3 and the other one to voter 5. The situation gets more complicated if we consider the possibility that voters abstain. For example, if voter 3 (the second most preferred delegate of voter 1) abstains, should voter 1's delegation then go to voter 4 (the third most preferred delegate of voter 1) or to voter 5 (the second most preferred delegate of voter 2, who is the most preferred delegate of voter 1)?

One possibility to resolve situations like the one just described consists in employing a *Markov chain* approach. Let $\epsilon \in (0, 1)$ and consider a random walk on the delegation graph (e.g., the graph in Figure 1) that starts at some fixed delegating voter and follows delegation edges as follows: With probability proportional to 1, go to the most preferred delegate of the current voter; with probability proportional to ϵ, go to the second most preferred delegate; with probability proportional to ϵ^2, go to the third most preferred delegate; and so on. The idea behind this definition is that we will almost always delegate to the most preferred delegate; if this results in a delegation cycle, however, we will eventually delegate elsewhere to leave the cycle. When the random walk gets stuck at an abstaining voter, it is restarted. When it reaches a voter that actually voted, we have found the voter to which the vote is delegated. This defines a randomized way to resolve delegation cycles that also takes care of abstentions.[9]

Whether this Markov chain approach always yields sensible solutions is not obvious. The axiomatic method, essential in (computational) social choice, will be a valuable guide when analyzing different implementations of delegative voting (such as the one outlined above or alternative proposals [11]). Another important

[6]For details, see the articles by Ford [33], Green-Armytage [40], and Blum and Zuber [9]. Some of the ideas behind delegative voting can be traced back to the works of Dodgson [25], Tullock [64], and Miller [50]. For an historical overview of ideas, see the surveys by Ford [34] and Behrens [5].

[7]The question whether the delegative voting paradigm actually leads to "superior" voting outcomes (as compared to direct and representative democracy) has been addressed from a variety of perspectives [1, 21, 38, 40, 42, 43].

[8]Behrens and Swierczek [7] argue that allowing backup delegations may lead to undesirable properties.

[9]Of course, the outcome—which can be seen as a probability distribution over non-abstaining and non-delegating voters—depends on the parameter ϵ. In order to get rid of this dependence, we could take the limit of this distribution as ϵ approaches zero. In the example in Figure 1, this procedure would result in both voters 1 and 2 delegating to voter 5 (because a delegation from voter 2 to voter 5 is, in the limit, infinitely more likely than a delegation from voter 1 to voter 4; recall that we have assumed that voter 3 abstains, so every time the random walk reaches that voter, it is restarted).

criterion for a delegative voting systems is that the voting weights resulting from the (possibly ranked) delegations can be computed efficiently. Whether this is the case for (variants of) the Markov chain approach outlined above is an interesting open problem. In particular, it would be interesting to construct an algorithm that computes the aggregate distribution of voting power without computing the limit distribution of each delegated vote individually.

2.2 Inconsistent Outcomes

As for problem (3), it has been observed by Blum and Zuber that the flexibility of the delegative voting paradigm, which allows voters to delegate their vote to different representatives depending on the area of the issue, can lead to outcomes that are not consistent on a global level [9, pp. 178–179]. Christoff and Grossi [20], who have studied this problem in the formal context of binary aggregation, suggest to employ techniques from opinion diffusion [24, 39] in order to resolve such inconsistencies. (They also propose a way to address problems (1) and (2), by requiring voters to specify "default" votes that are used to overrule delegation decisions should the latter lead to delegation cycles or abstentions.) And Brill and Talmon [16], who have studied the special case in which each pairwise comparison between two alternatives can be delegated to a different voter, suggest to employ the framework of distance rationalization [30].

In general, there appears to be a tradeoff between flexible and fine-grained delegation possibilities on the one side and increased potential of inconsistent (or underspecified) outcomes on the other side. In this context, it will be interesting to explore generalizations of delegative voting such as *statement voting* [65].

2.3 Strategic Aspects

Voters are said to *strategically manipulate*[10] an election if they achieve a preferable outcome by misstating their preferences. The famous Gibbard-Satterthwaite theorem [35, 58] states that every reasonable voting rule is vulnerable to strategic manipulation. Online decision platforms based on delegative voting give rise to novel issues regarding strategic manipulation.

For instance, some implementations of the delegative voting paradigm give voters the ability to see "where their delegation goes," i.e., which alternative the voter to which they delegated their vote, voted for. (If this voter delegated their vote further, voters will learn that as well.) Crucially, this transparency is often provided before the election takes place, in order to give voters the ability to reconsider their delegation. This way, voters can find out what other voters are going to vote for. (Indeed, this is sometimes considered an important advantage of the delegative voting paradigm [40].) Since delegations can be changed arbitrarily, voters have the ability to learn arbitrarily many votes of other voters (provided that those other voters have already submitted their votes/delegations to the system). Knowing how other voters vote makes it much easier for a voter to strategically manipulate an election. Thus, the desirable property of "delegation transparency" is in direct conflict with the goal of making elections less manipulable.

[10] Strategic manipulation should not be confused with malicious behavior such as deliberately not counting submitted votes, adding fake votes, etc. While ensuring the technological security of online decision platforms is without doubt a very important research area (see, e.g., [54]), it is not the focus of the current article.

A further and rather counterintuitive strategic issue regarding delegative voting was discovered by Schelling (see [52], Puzzle 5). Assuming that preferences are common knowledge, and using the amendment procedure, voters can be *worse off* in a situation where they have greater voting power (i.e., when other voters have delegated their vote to them) as compared to the situation where they just have a single vote (i.e., no other voter delegated their vote to them). In a sense, delegations can be used as commitment devices, changing the equilibria of the voting game [46]. This counterintuitive phenomenon perhaps suggests that *voters should not be forced to accept delegations.* Rather, there needs to be a mechanism by which voters can choose which delegations to accept and which to refuse. How exactly such a mechanism should be implemented, and which voting rules are vulnerable to Schelling's paradox, are important research questions that need to be addressed.

3 OTHER CHALLENGES

Interactive Democracy also gives rise to voting-theoretic problems that are independent from the delegative voting paradigm. This section briefly discusses three examples.

3.1 Proportional Representation

A defining feature of ID systems is that all participants are allowed—and encouraged—to contribute to the decision making process, either directly by participating in discussions and voting on issues, or indirectly by delegating their decision power. In particular, in a context where one (or more) out of several competing options needs to be selected, each participant can propose their own option if they are not satisfied with the existing ones. This may lead to situations where a very large number of alternative options needs to be considered. Since it cannot be expected that every participant looks at all available options before making a decision, the *order* in which competing options are presented (e.g., on a website that facilitates the discussion and voting process) plays a crucial role (see [6], Chapter 4.10). A natural approach is to order competing options by their "support", i.e., by the number of participants that have expressed their approval of the option in question.

This gives rise to what Behrens et al. [6] describe as the "noisy minorities" problem: relatively small groups of very active participants can "flood" the system with their contributions, creating the impression that their opinion is much more popular than it actually is. This is problematic insofar as alternative options (that are potentially much more popular) run the risk of being "buried" and not getting sufficient exposure. In order to prevent this problem, the mechanism ordering competing options needs to ensure that the order adequately reflects the opinions of the participants. The search for orderings that are "representative" in this sense leads to challenging algorithmic problems not unlike those underlying the problem of choosing representative committees (see, e.g., [4, 19, 51]). In a recent paper, Skowron et al. [62] approach the problem by formalizing the notion of *proportional ranking* and experimentally evaluating the representativeness of common ranking methods. Roughly speaking, a representative ranking is one in which the number of top positions allocated to options supported by a particular group is proportional to the size of that group.

3.2 Properties of Aggregation Functions

Social choice theory [3] tells us that the way preferences are aggregated matters, and that there is no perfect voting rule satisfying all desirable properties. Therefore, every voting rule represents a tradeoff between the desirable properties it satisfies and the desirable properties it violates. When selecting a voting rule for a particular application (either choosing an existing rule or designing a new rule), attention should be paid to the question which properties are important or even indispensable for the application at hand. For the development of ID systems it is therefore relevant to identify and study properties that are desirable in the context of online decision platforms.

One example of such a property is *independence of clones*. Two alternatives are said to be *clones* if they are ranked consecutively by all voters. For example, imagine a situation where citizens have to decide between two different proposals A and B for the public budget of a city, and assume that support for the two options is split roughly equally. Now, supporters of proposal A could come up with a third proposal, B', which is very similar to B (say, B and B' only differ in a handful of categories, and even there they differ by very small amounts). In the absence of any coordination among the B supporters, there is the risk that supporters of proposal B split into two camps, one favoring B and one favoring B'. If a voting rule naively counts support for the three proposals (as does, for example, the ubiquitous plurality rule), the introduction of new alternatives (B' in our example) has the potential to completely change the outcome (from a close race between A and B to a comfortable win for A). A voting rule is *independent of clones* if its result cannot be altered by the introduction of clones [63]. Independence of clones is a very desirable property because it eliminates incentives to strategically introduce new alternatives. It is particularly important in the context of online decision platforms, where all participants can nominate alternatives (see Section 3.1) and thus the existence of multiple very similar alternatives is very likely.

Only a handful of voting rules are known to be independent of clones [44, 59, 63]. Somewhat surprisingly, these voting rules do not seem to have much in common. Further work is needed in order to better understand independence of clones and other properties that are particularly relevant in an online voting context.[11] It would also be interesting to analyze how these properties can be exploited *algorithmically*, by building on earlier work that is mostly restricted to particular (classes of) voting rules [14, 22, 29].

3.3 New Forms of Aggregation

In traditional voting theory, it is usually assumed that there is a finite set of alternatives and the preferences of voters are given as rank-orderings over this set. This framework, while very general in theory, is not always practical. Often, the space of alternatives has some combinatorial structure, and exploiting this structure is necessary for both eliciting and aggregating preferences in a meaningful way [45]. Two example scenarios that arise in ID applications are the aggregation of *societal tradeoffs* [23], where the goal is to aggregate numerical tradeoffs between different kinds of socially undesirable activities and the collection of aggregate tradeoffs is required to be consistent, and *participatory budgeting* [17], where the goal is to allocate budgetary spending of a local government based on citizens' preferences, and a combinatorial structure of the solution space is imposed by a budget constraint.

For those and related scenarios, the structure of the solution space makes it impractical to elicit preferences directly. Rather, there are various ways in which users can specify their preferences. Each preference format presents a tradeoff between *expressive power* (does the format allow to express fine-grained preferences?), *succinctness* (can preferences be represented compactly?), and *aggregatability* (does the preference format allow for a computationally efficient and axiomatically desirable aggregation mechanism?), among other things.[12] Techniques developed in COMSOC facilitate a much-needed principled comparison of preference formats and aggregation methods. First steps in that direction have recently been made by Goel et al. [36] and Benade et al. [8].[13]

4 CONCLUSION

The emergence of citizen participation systems in general—and of online voting platforms in particular—appears to be an irreversible development. The question is not if, but rather when, these systems become standard components of the democratic process. A multidisciplinary research program is necessary for making these systems secure, equitable, inclusive, user-friendly, and computationally reliable. In this article, I have argued that insights and tools from computational social choice (COMSOC) are relevant for this important endeavor.

COMSOC is by far not the only relevant research area within multiagent systems that can inform the design of interactive democracy applications. For instance, in the future it is conceivable that *reputation systems* [55] and *recommender systems* [56, 57] are employed to help voters decide where to delegate their vote (or how to vote), *argumentation frameworks* [27, 28] are employed to structure deliberation processes, and logical frameworks like *CP-nets* [12] are employed to give voters more flexibility when expressing their preferences.

We know from experience that even a well-intended design often results in voting systems that exhibit unexpected flaws (e.g., see [32]). Moreover, once a voting system is established, it is very hard to change it due to its incumbent position. The emergence of interactive democracy thus presents a unique opportunity to influence the future of our democracies for the better.

ACKNOWLEDGMENTS

This material is based upon work supported by a Feodor Lynen research fellowship of the Alexander von Humboldt Foundation and by the Deutsche Forschungsgemeinschaft under grant BR 4744/2-1. The author would like to thank Dorothea Baumeister, Jan Behrens, Steven Brams, Vincent Conitzer, Paul Gölz, Umberto Grandi, Davide Grossi, Michel Le Breton, Rolf Niedermeier, Dominik Peters, Marcus Pivato, Ariel Procaccia, Ehud Shapiro, Nimrod Talmon, and Bill Zwicker for helpful comments and the anonymous referees for valuable feedback.

[11] Behrens et al. [6, pp. 89–90] note that independence of clones is necessary, but not sufficient, to avoid the problem of vote-splitting.

[12] Identifying reasonable formats in which voters can express their preferences has similarities to the search for bidding languages in combinatorial auctions [53].

[13] Shapiro and Talmon [61] have proposed an alternative approach towards participatory budgeting based on the Condorcet principle.

REFERENCES

[1] D. Alger. 2006. Voting by proxy. *Public Choice* 126, 1–2 (2006), 1–26.
[2] E. Aragonès and S. Sánchez-Pagés. 2009. A theory of participatory democracy based on the real case of Porto Alegre. *European Economic Review* 53, 1 (2009), 56–72.
[3] K. J. Arrow, A. K. Sen, and K. Suzumura (Eds.). 2002. *Handbook of Social Choice and Welfare*. Vol. 1. North-Holland.
[4] H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. 2017. Justified Representation in Approval-Based Committee Voting. *Social Choice and Welfare* 48, 2 (2017), 461–485.
[5] J. Behrens. 2017. The origins of liquid democracy. *The Liquid Democracy Journal* 5 (2017).
[6] J. Behrens, A. Kistner, A. Nitsche, and B. Swierczek. 2014. *The Principles of LiquidFeedback*.
[7] J. Behrens and B. Swierczek. 2015. Preferential Delegation and the Problem of Negative Voting Weight. *The Liquid Democracy Journal* 5 (2015), 6–34.
[8] G. Benade, S. Nath, A. D. Procaccia, and N. Shah. 2017. Preference Elicitation For Participatory Budgeting. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*. AAAI Press, 376–382.
[9] C. Blum and C. I. Zuber. 2016. Liquid Democracy: Potentials, Problems, and Perspectives. *Journal of Political Philosophy* 24, 2 (2016), 162–182.
[10] G. Boella, L. Francis, E. Grassi, A. Kistner, A. Nitsche, A. Noskov, L. Sanasi, A. Savoca, C. Schifanella, and I. Tsampoulatidis. 2018. WeGovNow: A Map Based Platform to Engage the Local Civic Society. In *Companion Proceedings of the The Web Conference 2018 (WWW '18)*. 1215–1219.
[11] P. Boldi, F. Bonchi, C. Castillo, and S. Vigna. 2011. Viscous democracy for social networks. *Commun. ACM* 54, 6 (2011), 129–137.
[12] C. Boutilier, R. I. Brafman, C. Domshlak, H. H. Hoos, and D. Poole. 2004. CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *Journal of Artificial Intelligence Research* 21 (2004), 135–191.
[13] R. Brafman and C. Domshlak. 2009. Preference Handling—An Introductory Tutorial. *AI Magazine* 30, 1 (2009), 58–86.
[14] F. Brandt, M. Brill, and H. G. Seedig. 2011. On the Fixed-Parameter Tractability of Composition-Consistent Tournament Solutions. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*. AAAI Press, 85–90.
[15] F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia (Eds.). 2016. *Handbook of Computational Social Choice*. Cambridge University Press.
[16] M. Brill and N. Talmon. 2018. Pairwise Liquid Democracy. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence and the 23rd European Conference on Artificial Intelligence (IJCAI-ECAI-18)*. Forthcoming.
[17] Y. Cabannes. 2004. Participatory budgeting: a significant contribution to participatory democracy. *Environment and Urbanization* 16, 1 (2004), 27–46.
[18] D. Carswell. 2012. *The End of Politics and the Birth of iDemocracy*. Biteback Publishing.
[19] J. R. Chamberlin and P. N. Courant. 1983. Representative Deliberations and Representative Decisions: Proportional Representation and the Borda Rule. *The American Political Science Review* 77, 3 (1983), 718–733.
[20] Z. Christoff and D. Grossi. 2017. Binary Voting with Delegable Proxy: An Analysis of Liquid Democracy. In *Proceedings of the 16th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*. 134–150.
[21] G. Cohensius, S. Mannor, R. Meir, E. Meirom, and A. Orda. 2017. Proxy Voting for Better Outcomes. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. IFAAMAS, 858–866.
[22] V. Conitzer. 2006. Computing Slater Rankings Using Similarities among Candidates. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*. AAAI Press, 613–619.
[23] V. Conitzer, M. Brill, and R. Freeman. 2015. Crowdsourcing Societal Tradeoffs. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS) Blue Sky Ideas track*. IFAAMAS, 1213–1217.
[24] M. H. DeGroot. 1974. Reaching a Consensus. *J. Amer. Statist. Assoc.* 69, 345 (1974), 118–121.
[25] C. L. Dodgson. 1884. *The Principles of Parliamentary Representation*. Harrison and Sons.
[26] C. Domshlak, E. Hüllermeier, S. Kaci, and H. Prade. 2011. Preferences in AI: An overview. *Artificial Intelligence* 175, 7–8 (2011), 1037–1052.
[27] P. M. Dung. 1995. On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-Person Games. *Artificial Intelligence* 77 (1995), 321–357.
[28] P. E. Dunne and T. J. M. Bench-Capon. 2007. Argumentation in Artificial Intelligence. *Artificial Intelligence* 171, 10-15 (2007), 619–641.
[29] E. Elkind, P. Faliszewski, and A. Slinko. 2012. Clone Structures in Voters' Preferences. In *Proceedings of the 13th ACM Conference on Electronic Commerce (ACM-EC)*. ACM Press, 496–513.
[30] E. Elkind and A. Slinko. 2016. Rationalizations of Voting Rules. In *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia (Eds.). Cambridge University Press, Chapter 8.
[31] U. Endriss (Ed.). 2017. *Trends in Computational Social Choice*. AI Access.
[32] P. C. Fishburn. 1974. Paradoxes of Voting. *The American Political Science Review* 68, 2 (1974), 537–546.
[33] B. Ford. 2002. Delegative Democracy. (2002). Unpublished manuscript. Available at http://www.brynosaurus.com/deleg/deleg.pdf.
[34] B. Ford. 2014. Delegative Democracy Revisited. Blog post. (2014). Available at http://bford.github.io/2014/11/16/deleg.html.
[35] A. Gibbard. 1973. Manipulation of Voting Schemes: A General Result. *Econometrica* 41, 4 (1973), 587–601.
[36] A. Goel, A. K. Krishnaswamy, S. Sakshuwong, and T. Aitamurto. 2016. Knapsack voting: Voting mechanisms for Participatory Budgeting. (2016). Unpublished manuscript.
[37] J. Goldman and A. D. Procaccia. 2014. Spliddit: Unleashing Fair Division Algorithms. *ACM SIGecom Exchanges* 13, 2 (2014), 41–46.
[38] P. Gölz, A. Kahng, S. Mackenzie, and A. D. Procaccia. 2018. The Fluid Mechanics of Liquid Democracy. (2018). Working paper.
[39] U. Grandi. 2017. Social Choice and Social Networks. In *Trends in Computational Social Choice*, U. Endriss (Ed.). AI Access, Chapter 9, 169–184.
[40] J. Green-Armytage. 2015. Direct voting and proxy voting. *Constitutional Political Economy* 26, 2 (2015), 190–220.
[41] B. N. Hague and B. D. Loader (Eds.). 1999. *Digital Democracy: Discourse and Decision Making in the Information Age*. Routledge.
[42] A. Kahng, S. Mackenzie, and A. D. Procaccia. 2018. Liquid Democracy: An Algorithmic perspective. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*. AAAI Press. Forthcoming.
[43] C. C. Kling, J. Kunegis, H. Hartmann, M. Strohmaier, and S. Staab. 2015. Voting Behaviour and Power in Online Democracy. In *Proceedings of the 9th International AAAI Conference on Web and Social Media (ICWSM)*. 208–217.
[44] G. Laffond, J. Lainé, and J.-F. Laslier. 1996. Composition-Consistent Tournament Solutions and Social Choice Functions. *Social Choice and Welfare* 13, 1 (1996), 75–93.
[45] J. Lang and L. Xia. 2016. Voting in Combinatorial Domains. In *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia (Eds.). Cambridge University Press, Chapter 9.
[46] M. Le Breton. 2016. *The Condorcet Principle Implies the Proxy Voting Paradox*. TSE Working Papers 16-619. Toulouse School of Economics (TSE).
[47] P. Mancini. 2014. How to upgrade democracy for the Internet era [video file]. Retrieved from http://www.ted.com/talks/pia_mancini_how_to_upgrade_democracy_for_the_internet_era. (2014).
[48] P. Mancini. 2015. Why it is time to redesign our political system. *European View* 14, 1 (2015), 69–75.
[49] M. Margolis and G. Moreno-Riaño. 2013. *The Prospect of Internet Democracy*. Routledge.
[50] J. C. Miller. 1969. A program for direct and proxy voting in the legislative process. *Public Choice* 7, 1 (1969), 107–113.
[51] B. L. Monroe. 1995. Fully Proportional Representation. *The American Political Science Review* 89, 4 (1995), 925–940.
[52] B. Nalebuff. 1988. Puzzles: Cider in Your Ear, Continuing Dilemma, The Last Shall Be First, and More. *Journal of Economic Perspectives* 2, 2 (1988), 149–156.
[53] N. Nisan. 2006. Bidding Languages for Combinatorial Auctions. In *Combinatorial Auctions*, P. Cramton, Y. Shoham, and R. Steinberg (Eds.). MIT Press.
[54] T. Peacock, P. Y. A. Ryan, S. Schneider, and Z. Xia. 2013. Verifiable Voting Systems. In *Computer and Information Security Handbook* (2nd ed.), J. R. Vacca (Ed.). Elsevier, Chapter 69.
[55] P. Resnick, K. Kuwabara, R. Zeckhauser, and E. Friedman. 2000. Reputation systems. *Commun. ACM* 43, 12 (2000), 45–48.
[56] P. Resnick and H. R. Varian. 1997. Recommender Systems. *Commun. ACM* 40, 3 (1997), 56–58.
[57] F. Ricci, L. Rokach, B. Shapira, and P. B. Kantor (Eds.). 2011. *Recommender Systems Handbook*. Springer.
[58] M. A. Satterthwaite. 1975. Strategy-Proofness and Arrow's Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions. *Journal of Economic Theory* 10, 2 (1975), 187–217.
[59] M. Schulze. 2011. A new monotonic, clone-independent, reversal symmetric, and Condorcet-consistent single-winner election method. *Social Choice and Welfare* 36, 2 (2011), 267–303.
[60] E. Shapiro. 2018. Foundations of e-Democracy. *Commun. ACM* (2018). Forthcoming.
[61] E. Y. Shapiro and N. Talmon. 2017. *A Democratically-Optimal Budgeting Algorithm*. Technical Report. http://arxiv.org/abs/1709.05839.
[62] P. Skowron, M. Lackner, M. Brill, D. Peters, and E. Elkind. 2017. Proportional Rankings. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*. AAAI Press, 409–415.
[63] T. N. Tideman. 1987. Independence of Clones as a Criterion for Voting Rules. *Social Choice and Welfare* 4, 3 (1987), 185–206.
[64] G. Tullock. 1967. *Towards a Mathematics of Politics*. University of Michigan Press.
[65] B. Zhang and H. Zhou. 2017. Brief Announcement: Statement Voting and Liquid Democracy. In *Proceedings of the 36th ACM Symposium on Principles of Distributed Computing (PODC)*. 359–361.