Towards Reality: Smoothed Analysis in Computational Social Choice

Blue Sky Ideas Track

Dorothea Baumeister, Tobias Hogrebe, Jörg Rothe
Institut für Informatik
Heinrich-Heine-Universität Düsseldorf
{d.baumeister,tobias.hogrebe,rothe}@uni-duesseldorf.de

ABSTRACT

Hemaspaandra [22] celebrated the quite close relationship between computational social choice and computational complexity as a two-way street from which both areas benefited in the past, and expressed his hope that the areas become best friends forever. Later on, Rothe [38] celebrated the prominent Borda voting rule and surveyed recent advances on the complexity of problems related to the three most fundamental models of tampering with elections—namely, via manipulation, control, and bribery—and even related to using Borda beyond voting: in fair division and coalition formation in hedonic games. But now the party is over: no more celebration! Instead, we present a common criticism regarding computational social choice persistently making use of worst-case complexity. To overcome this shortcoming, we propose our blue sky idea of applying to problems from computational social choice the method of smoothed analysis due to Spielman and Teng [43, 44] and also used by Bläser and Manthey [7], as some sort of a middle ground between the worst-case and the average-case analysis of algorithms.

KEYWORDS

social choice; voting; computational complexity; smoothed analysis

ACM Reference Format:

Dorothea Baumeister, Tobias Hogrebe, Jörg Rothe. 2020. Towards Reality: Smoothed Analysis in Computational Social Choice. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), Auckland, New Zealand, May 9–13, 2020,* IFAAMAS, 5 pages.

1 INTRODUCTION

The two-way interaction between *computational complexity* and *computational social choice* (a.k.a. *COMSOC*) has been emphasized by Hemaspaandra [22, p. 7971], who points out that COMSOC's "existence as a recognized, distinctive research area within AI/multi-agent-systems is quite recent" and "its growth as an area during these recent years has been explosive." He further makes a strong case for a synergy between COMSOC and computational complexity being at least "partially responsible for that growth." Approvingly, Rothe [38, p. 9835] calls COMSOC "a true success story within AI," showcasing—by results for the prominent Borda voting rule—recent advances on the complexity of problems related to the three most fundamental models of how to influence the outcome of an election: manipulation, control, and bribery, each coming in a variety of specific scenarios motivated by real life. There are tons of literature studying the computational complexity of the related problems, elaborately surveyed for manipulation by Conitzer and Walsh [10], for control and bribery by Faliszewski and Rothe [19], and for manipulation, control, and bribery by Baumeister and Rothe [5]. We here focus on these three basic scenarios (manipulation, control, and bribery) as they are at the heart of COMSOC, widely known within the community, and responsible for the great success of COMSOC in recent years. Of course, in the long run, one might consider our approach also in a broader setting, for example regarding multiwinner elections, judgment aggregation scenarios, and so on. Indeed, the combination of social choice and computational complexity is only one aspect in the broad research area of COMSOC that covers a wide range of topics from fair division over opinion diffusion and matching to coalition formation. An overview is given in the textbooks by Brandt et al. [8], Endriss [14], and Rothe [37]. However, here we will focus on the computational complexity of various interference problems in elections.

A common criticism regarding the current orientation of computational social choice is the persistent use of worst-case complexity. Often, worst-case complexity is implicitly or explicitly stated as an argument as to why voting systems are resistant to malicious interferences such as manipulation, control, and bribery.[1] This approach is tempting, due to results from social choice theory such as the Gibbard–Satterthwaite theorem [20, 41], which roughly says that essentially every reasonable voting rule is susceptible to manipulation. Therefore, one may hope that they still resist manipulative attacks in practice simply due to the computational burden being too high. This argument was first made by Bartholdi et al. [2] and was fundamental to the emergence of computational social choice.

Also, experimental evaluations of real-world or artificial data show that "hard" problems can often be solved efficiently in practice. For example, Walsh [46] showed that checking the manipulability of an STV (single transferable vote) election with a single manipulator, a problem that is resistant in terms of being NP-hard [1], can in fact be solved with low effort on both real-world data and artificial data. Walsh [46] made similar observations about the veto rule. This issue has long been known and has often been discussed for manipulative attacks on elections, for example, by Faliszewski and Procaccia [17], Walsh [45], and Rothe and Schend [40]. Beyond voting, Hemaspaandra and Williams [23] have given an excellent, though atypical, treatise of typical-case heuristic algorithms.

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand.

[1] Such attacks are not always considered malicious: Bribery can also be seen positively as "campaign management" [18, 19, 42] and regarding manipulation Rothe [38, p. 9832] writes, "After all, every voter—human or software agent—has the right to think strategically about which vote to cast; not doing so would not be smart."

While empirical studies using real-world data and artificially generated data give us an idea of what the practical complexity of such problems may look like, their validity and reliability is often debatable: Suitable real-world data may be difficult to obtain and not scalable, which means that we observe only a fraction of the large-scale behavior. Although generated data is scalable, it can be inappropriate due to the models: Naive models, such as impartial culture (i.e., assuming the uniform distribution over all possible votes and profiles), can be easily distinguished from real-world data and are therefore inappropriate, as, e.g., Regenwetter et al. [36] point out. Richer models, such as the model of Mallows [31] or a spatial model [33], must first be trained with or fitted to real-world data to produce realistic results, making a theoretical investigation difficult. Especially if the choice of parameters relies on real-world data, scalability may be a problem. Either way, the choice of a distribution restricts the results and makes them less reliable.

To avoid explicit assumptions about the distribution of votes, Procaccia and Rosenschein [35] define the class of *junta distributions*. The idea is to capture a class of distributions that concentrate on hard instances in a way that if a manipulation problem can be solved efficiently for just those distributions, it can also be solved efficiently for other reasonable distributions. Erdélyi et al. [15], however, show that the three core properties of junta distributions (which, thus restricted, they call "basic junta") allow to construct a basic junta for which NP-hard problems such as SAT can be solved in deterministic heuristic polynomial time with high probability.

Another approach is to study the problems in terms of their parameterized complexity. This multivariate complexity analysis may yield fixed-parameter tractable algorithms if some input parameter can be bounded. Parameterized complexity results in COMSOC have been surveyed by Bredereck et al. [9], Dorn and Schlotter [12], and Lindner and Rothe [28]. However, as this again relies on worst-case complexity, the same issues arise. In some applications, it may be better to use fast approximation algorithms instead of slow optimal algorithms. Markakis [32] and Nguyen et al. [34] surveyed recent results on approximation algorithms in fair division.

We propose an alternative approach that allows theoretical investigations of realistic complexity without having to make such strong assumptions: *smoothed analysis*, some sort of a middle ground between the worst-case and the average-case analysis of algorithms. This approach seems to have been overlooked in the quest for more realistic complexity results protecting elections against manipulation, control, and bribery. In a nutshell, instead of considering either the worst complexity of a single instance or the expected complexity regarding a distribution of all instances at a given length, we now consider *the worst expected complexity of an algorithm with respect to a random perturbation of a given instance*. Smoothed analysis was very successfully introduced by Spielman and Teng [43] to explain the practical efficiency of the simplex algorithm, despite its exponential worst-case complexity. For developing this novel approach and its great practical relevance, Spielman and Teng were awarded the 2008 Gödel Prize by the ACM and the EATCS.

2 SMOOTHED ANALYSIS

As mentioned earlier, smoothed analysis measures the complexity of an algorithm with respect to the worst expected complexity of an instance, taking into account a random perturbation of the instance. The idea behind this is that even if in principle instances leading to an exponential worst-case complexity of an algorithm can occur, it is in practice highly unlikely to hit exactly those hard instances. The expected complexity of an algorithm would thus be much better a measure when only some hard instances are like islands rising steeply over an ocean of easy instances with a much lower complexity. Almost all real-world data is subject to such perturbations. For example, data that is automatically measured by sensors is subject to permanent fluctuations due to changes and measurement inaccuracies. Further, data can be perturbed to limit disclosure of sensitive information or to protect confidential data [27, 29]. And, finally, elections and surveys are also subject to permanent fluctuations by changes in the voters' minds (even if their number remains constant), i.e., in the individuals' opinions due to psychological, emotional, and behavioral effects [36, 39]. In particular, for a sufficiently large number of voters, the expected profile is subject to massive uncertainty, even if good predictions exist. In fact, this is the reason why the application of smoothed analysis to voting problems in computational social choice does make sense and could be fruitful.

We now formally define the notion of smoothed analysis and will then present our corresponding models regarding elections. Since we here consider election profiles originating from discrete domains, yet the original model of smoothed analysis by Spielman and Teng [44] is mainly designed for the perturbation of continuous domains, we present the definitions of Bläser and Manthey [7] instead.

In the model of Bläser and Manthey [7], perturbation is modeled as a family of distributions $\mathcal{D} = (D_{n,x,\phi})$, which includes a distribution $D_{n,x,\phi}$ for each possible instance x as the mean of the distribution, ϕ as a parameter that limits the maximum density of the distribution, and, only for clarity, n as the length of the instance x. By definition, they demand that $D_{n,x,\phi}(y) \leq \phi$ holds for each combination of n, x, ϕ, and y. The idea of choosing the maximum density ϕ as a parameter of the distributions is attributed to Beier and Vöcking [6] by Bläser and Manthey [7].

Additionally, they demand that the binary encoding of the support $\{y \mid D_{n,x,\phi}(y) > 0\}$ of each distribution $D_{n,x,\phi}$ is contained in $\{0,1\}^{\mathrm{poly}(n)}$, thus bounding the length of each "neighbor" y of instance x with nonzero probability by a polynomial in n, the length of x. Moreover, the possible values of ϕ should be discretized in a way that they are representable in $\{0,1\}^{\mathrm{poly}(n)}$ and should also be contained in $[1/N_{n,x}, 1]$ with $N_{n,x} = |\{y \mid D_{n,x,\phi}(y) > 0 \text{ for some } \phi\}|$.

For $\phi = 1$, we allow the distribution to collapse to x with probability 1, giving us the worst-case complexity. For $\phi = 1/N_{n,x}$, we enforce the uniform distribution, which gives us the classical average-case complexity with respect to the uniform distribution.

Now, an algorithm A admits a *smoothed polynomial complexity* regarding distribution family $\mathcal{D}$ if there exists an $\epsilon > 0$ such that $\mathbb{E}_{y \sim D_{n,x,\phi}}(\mathrm{time}_A(y, n, \phi)^{\epsilon}) = O(n \cdot N_{n,x} \cdot \phi)$ for all combinations of n, x, and ϕ. Note that this definition is analogous to the definition of average-case polynomial-time complexity by Levin [26]. Levin did not use the naive definition of the expected running time being bounded by a polynomial as it would not be robust, e.g., against polynomial slowdowns of the algorithm due to changing the computational model (as noticed, e.g., by Goldreich [21]).

For the same reason, Bläser and Manthey [7] also chose to define the expected running time raised to some power ϵ. A crucial difference between worst-case complexity and smoothed analysis is that the latter focuses on specific algorithms, whereas the former makes a statement on hard instances irrespective of the used algorithm.

As an example, Bläser and Manthey [7] show that the graph coloring problem admits a smoothed polynomial complexity. Given a graph $G = (V, E)$, the perturbation model is defined through φ, which means that in the perturbed graph G' each edge is flipped independently with a probability of φ while the set of vertices remains unchanged. A graph with n vertices can then be represented as a binary string of length $\binom{n}{2}$, and it holds that $N_{n,G} = 2^{\binom{n}{2}}$. The parameter ϕ for the smoothed analysis is given through φ by choosing $\varphi \leq \frac{1}{2}$ such that $(1-\varphi)^{\binom{n}{2}} = \phi$. Note that this model includes, on one side of the spectrum, the fully random graph by choosing $\phi = 2^{-\binom{n}{2}}$ and $\varphi = \frac{1}{2}$ and, on its other side, the worst case with $\phi = 1$ and thus $\varphi = 0$.

The problem k-Coloring asks whether the vertices of a graph can be colored with at most k different colors such that no pair of adjacent vertices is colored the same. To show that this problem admits a smoothed polynomial complexity, it suffices to analyze the simple algorithm which checks whether there is a clique of size $k + 1$ in the input graph that is perturbed according to the model described above. If so, the answer is no; otherwise, an exhaustive search is performed. It can be shown that the expected running time raised to some constant power is polynomial.

Smoothed analysis has also been applied in algorithmic game theory. For example, Deng et al. [11] introduce the notion of smoothed approximation ratio for the performance comparison between optimal and truthful mechanisms, under the assumption of perturbed inputs, and Huang and Teng [24] analyze the smoothed complexity of solving the approximate Leontief market exchange problem through Scarf's general fixed-point algorithm.

As demonstrated, for instance, by Bläser and Manthey [7], there are also concepts to show hardness under smoothed analysis. Here, however, we focus on smoothed efficiency, i.e., on smoothed polynomial complexity.

3 PERTURBATION MODELS FOR ELECTIONS

In order to perform a smoothed analysis of problems in computational social choice, appropriate perturbation models are needed. One of the main research fields within this area is voting, and in particular various ways of how to tamper with the outcome of elections.

An election consists of a set C of candidates and a preference profile V (i.e., a list of the voters' preferences) over C. Then a voting rule $\mathcal{R}$ is applied in order to determine the winning candidate(s). Typical problems studied in computational social choice include winner determination, manipulation, bribery, and control. Let us describe these problems in a nutshell. Let $\mathcal{R}$ be a voting rule.

• In the *winner determination* problem, we ask whether some designated candidate is an $\mathcal{R}$ winner of a given election.

• The (most basic variant of the) *manipulation* problem for $\mathcal{R}$ asks, given an election with a distinguished candidate, whether it is possible for a manipulator (joining the given election) to cast a strategic vote such that the distinguished candidate is an $\mathcal{R}$ winner of the resulting election. This problem was introduced by Bartholdi et al. [2], and there are many more variants of manipulation depending, for instance, on whether the election is weighted or unweighted; a destructive variant aiming at preventing someone's victory complements the above constructive variant; there may be a *coalition* of manipulators instead of just a single manipulator; and so on. For more information, we refer to the book chapters by Conitzer and Walsh [10] and Baumeister and Rothe [5].

• In (the most basic variant of) *bribery*, given an election with a designated candidate, a budget, and a collection of price functions (one per voter). An external agent bribes the voters to change their votes in the limit of the briber's budget. The question is whether there is an alternative preference profile resulting from such a bribery action such that the designated candidate is an $\mathcal{R}$ winner of the resulting election. This basic variant was first studied by Faliszewski et al. [16].

• Finally, for, e.g., the problem *constructive control by deleting voters* under $\mathcal{R}$, we are given an election with a distinguished candidate and some nonnegative integer bound k, and we ask whether it is possible for the election chair (who has the power to change the structure of elections) to make the distinguished candidate an $\mathcal{R}$ winner by deleting up to k votes. The study of control in elections was initiated by Bartholdi et al. [3].

Again, there are many more variants of bribery and control problems; see, e.g., the book chapters by Faliszewski and Rothe [19] and Baumeister and Rothe [5].

In the following, we present suggestions regarding models for the perturbation of instances in the context of elections. Note that here we define perturbations for election profiles themselves, in other words, for lists of votes. While there are election problems, such as the winner determination and the manipulation problem, whose instances consist of only one preference profile (including a distinguished candidate), there also exist election problems, such as bribery and control problems, which require additional inputs such as the briber's budget or an integer limiting the number of deleted voters or candidates in control. For the latter, one can either accept the budget or deletion limit as commonly known (meaning that it is not subject to perturbation), or one must extend the perturbation to the budget or deletion limit. The application of smoothed analysis in computational social choice is of course not limited to the analysis of different ways to tamper with the outcome of an election. We present a general perturbation model for preference profiles that are central to many problems in elections and hence may be extended for the study of other problems like gerrymandering, committee elections, and many more. In addition to the analysis of the computational complexity these models can be used to determine the practical relevance of properties such as the occurrence of voting paradoxes or ties without having a general understanding of the distribution of the profiles.

We distinguish between the two best known vote types: *complete linear orders* and *approval vectors*. For complete linear orders, we propose an extension of the classical *model of Mallows* [31] to complete preference profiles as our perturbation model.[2] In this model, the profiles are weighted exponentially decreasing according to

[2] As mentioned in the introduction, Mallows is used here as a perturbation model and not as a model for generating artificial data.

their pairwise comparisons deviating from the original profile. The strength of the probability decrease is indicated by the dispersion factor $\varphi \in [0, 1]$. Note that the size of the alternative (i.e., perturbed) profiles is fixed to the size of the original profile.

For a set C of m candidates, an original profile $V = (v_1, \dots, v_n)$, and an alternative profile $V' = (v'_1, \dots, v'_n)$, the model is then given by the following probability distribution for V': $\Pr(V' \mid V, \varphi) = \varphi^{\sum_{i=1}^{n} \mathrm{inv}(v_i, v'_i)} / (Z_{m,\varphi})^n$ with $\mathrm{inv}(v, v') = |\{(a, b) \in C \times C \mid a >_v b \wedge b >_{v'} a\}|$ the number of inversions. The normalization factor $Z_{m,\varphi}$ can be expressed as $1 \cdot (1 + \varphi) \cdot (1 + \varphi + \varphi^2) \cdots (1 + \varphi + \cdots + \varphi^{m-1})$, as proposed, e.g., by Lu and Boutilier [30]. Note that the presented probability distribution is the product of the vote-wise application of the classical Mallows model. For $\varphi = 0$, the distribution collapses to V with probability 1. For $\varphi = 1$, we obtain the uniform distribution over all profiles. To make it applicable to smoothed analysis, we obtain the following parameters analogously to the analysis of the graph coloring problem due to Bläser and Manthey [7]. The length of the instance is $n \cdot m$ and the number of preference profiles with nonzero probability is $N_{n \cdot m, V} = (m!)^n$.

The parameter ϕ is given by φ and vice versa, since $\phi = \frac{1}{(Z_{m,\varphi})^n}$ and the maximum probability is always given by the probability of the original profile itself. For proofs, it may be advisable to argue about the choice of φ and the thus resulting ϕ the other way around, since deriving the associated φ from ϕ seems to be problematic.

For approval vectors in $\{0, 1\}^m$, we propose a variant of the Mallows model based on the *Hamming distance* H: $\Pr(V' \mid V, \varphi) = \varphi^{\sum_{i=1}^{n} H(v_i, v'_i)} / (Z_{m,\varphi})^n$. Again, the normalization $Z_{m,\varphi}$ can be expressed efficiently by $\sum_{k=0}^{m} \binom{m}{k} \cdot \varphi^k$. While this can be used well as a perturbation model and we can also efficiently sample from it, efficiently generating more realistic data through a related mixture model is problematic, as Irurozki et al. [25] point out. Similar models can be defined for variants with a fixed number of approvals.

While the two presented models are compelling models for investigating a variety of problems like winner determination, manipulation, bribery, and control, there are also other suitable models that may depend on the problems themselves, the motivation, and the assumptions about the uncertainty regarding the instances of the problems. For example, in addition to perturbing actual votes, the question of whether or not a voter (or even a candidate) will participate in an election may itself be subject to uncertainty. This could be expressed by a model in which voters (or candidates) have a certain probability of not participating in the election. Similar models have already been used in the literature, e.g., by Wojtas and Faliszewski [47]. If the profiles are given in succinct representation (i.e., as a list containing each occurring vote with the corresponding multiplicity) or if the voters are weighted, one can also use models in which the respective numbers or weights are uncertain.

Another type of preference profile that is also often used are *tournament graphs* in which the pairwise comparison between two candidates is represented as the orientation of an edge in a complete graph whose vertices are the candidates (or participants when it is not an election). Such tournaments may result from pairwise elections. A suitable model for this would be that—analogously to the graph model presented above in which the existence of an edge is flipped with a certain probability—the *orientation* of each edge is now reversed with a certain probability. This models, for example, the situation in sports tournaments where the outcome of a game is unpredictable to a certain degree. Majority graphs that represent the pairwise majority relation of a given preference profile are also used; however, a different perturbation model incorporating the additional information from the preferences would be needed here.

4 CONCLUSIONS AND OUTLOOK

Using smoothed complexity is an important step towards reality in analyzing problems of computational social choice from a new perspective when implementing suitable decision-making systems. In times of digitized decision making, one of COMSOC's main goals should be to provide practical algorithms to solve the many hard (in terms of worst-case complexity) problems that arise. Smoothed analysis can be a tool to prove with theoretical means the practical efficiency of algorithms for those hard problems. By analyzing the practical complexity of the problems, not only the opportunities but also the risks of digital democracy from a complexity point of view can be estimated. The two perturbation models we have proposed may be fundamental to analyze algorithms for election problems involving preference profiles based on linear orders or approval vectors. Challenges for future research are now to develop and analyze algorithms for such election problems that can be used in practical applications. This involves a deeper understanding of the structure of the problems, and identifying components that are responsible for their high worst-case complexity. Here, existing analyses from a parameterized point of view may be useful, since they also explore the structure of the given instances.

While we have focused primarily on elections, there are other areas in COMSOC or related to COMSOC where it could be fruitful to apply smoothed analysis. These include, for example, *fair allocation*, *matching*, *hedonic games*, and *judgment aggregation*. Another example is *abstract argumentation*, which builds on the seminal work of Dung [13]. Argumentations are described as graphs, where the vertices correspond to the arguments and the edges model the attack relation between them. Different semantics are then used to identify acceptable sets of arguments with certain properties. In particular, the application and automatic processing of argumentations in the context of online participation becomes more and more important and algorithms with a fast practical running time are desired. Many related problems, however, are computationally hard (again, in the worst case) for abstract argumentation frameworks, especially so when uncertainty is involved [4]. Therefore, smoothed analysis would again be useful in this context.

So far, there are only a few papers using smoothed analysis, and some of the analyzed problems are rather artificial. We suspect that computational social choice is an area where results on smoothed analysis may be obtained for *natural* and *practically important* problems, as it was the case in the worst-case analysis for the higher levels of the polynomial hierarchy [22].

ACKNOWLEDGMENTS

We thank the AAMAS'20 reviewers for helpful comments and Scott Alfeld for pointing us to the literature on smoothed complexity. This work was supported by DFG grants BA 6270/1-1 and RO 1202/14-2.

REFERENCES

[1] J. Bartholdi and J. Orlin. 1991. Single Transferable Vote Resists Strategic Voting. *Social Choice and Welfare* 8, 4 (1991), 341–354.

[2] J. Bartholdi, C. Tovey, and M. Trick. 1989. The Computational Difficulty of Manipulating an Election. *Social Choice and Welfare* 6, 3 (1989), 227–241.

[3] J. Bartholdi, C. Tovey, and M. Trick. 1992. How Hard is it to Control an Election? *Mathematical Comput. Modelling* 16, 8/9 (1992), 27–40.

[4] D. Baumeister, D. Neugebauer, J. Rothe, and H. Schadrack. 2018. Verification in Incomplete Argumentation Frameworks. *Artificial Intelligence* 264 (2018), 1–26.

[5] D. Baumeister and J. Rothe. 2015. Preference Aggregation by Voting. In *Economics and Computation*, J. Rothe (Ed.). Springer, 197–325.

[6] R. Beier and B. Vöcking. 2004. Random Knapsack in Expected Polynomial Time. *J. Comput. System Sci.* 69, 3 (2004), 306–329.

[7] M. Bläser and B. Manthey. 2012. Smoothed Complexity Theory. In *Proc. MFCS'12*. Springer *LNCS #7464*, 198–209.

[8] F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia (Eds.). 2016. *Handbook of Computational Social Choice*. Cambridge University Press.

[9] R. Bredereck, J. Chen, P. Faliszewski, J. Guo, R. Niedermeier, and G. Woeginger. 2014. Parameterized Algorithmics for Computational Social Choice: Nine Research Challenges. *Tsinghua Science and Technology* 19, 4 (2014), 358–373.

[10] V. Conitzer and T. Walsh. 2016. Barriers to Manipulation in Voting. In *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia (Eds.). Cambridge University Press, 127–145.

[11] X. Deng, Y. Gao, and J. Zhang. 2017. Smoothed and Average-Case Approximation Ratios of Mechanisms: Beyond the Worst-Case Analysis. In *Proceedings of the 42th International Symposium on Mathematical Foundations of Computer Science*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, LIPIcs, 16:1–16:15.

[12] B. Dorn and I. Schlotter. 2017. Having a Hard Time? Explore Parameterized Complexity. In *Trends in Computational Social Choice*, U. Endriss (Ed.). AI Access Foundation, 209–230.

[13] P. Dung. 1995. On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-Person Games. *Artificial Intelligence* 77, 2 (1995), 321–357.

[14] U. Endriss (Ed.). 2017. *Trends in Computational Social Choice*. AI Access Foundation.

[15] G. Erdélyi, L. Hemaspaandra, J. Rothe, and H. Spakowski. 2009. Generalized Juntas and NP-Hard Sets. *Theoretical Computer Science* 410, 38–40 (2009), 3995–4000.

[16] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. 2009. How Hard Is Bribery in Elections? *Journal of Artificial Intelligence Research* 35 (2009), 485–532.

[17] P. Faliszewski and A. Procaccia. 2010. AI's War on Manipulation: Are We Winning? *AI Magazine* 31, 4 (2010), 53–64.

[18] P. Faliszewski, Y. Reisch, J. Rothe, and L. Schend. 2015. Complexity of Manipulation, Bribery, and Campaign Management in Bucklin and Fallback Voting. *Autonomous Agents and Multi-Agent Systems* 29, 6 (2015), 1091–1124.

[19] P. Faliszewski and J. Rothe. 2016. Control and Bribery in Voting. In *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia (Eds.). Cambridge University Press, 146–168.

[20] A. Gibbard. 1973. Manipulation of Voting Schemes: A General Result. *Econometrica* 41, 4 (1973), 587–601.

[21] O. Goldreich. 1997. *Notes on Levin's Theory of Average-Case Complexity*. Technical Report TR97-058. Electronic Colloquium on Computational Complexity.

[22] L. Hemaspaandra. 2018. Computational Social Choice and Computational Complexity: BFFs? In *Proc. AAAI'18*. AAAI Press, 7971–7977.

[23] L. Hemaspaandra and R. Williams. 2012. An Atypical Survey of Typical-Case Heuristic Algorithms. *SIGACT News* 43, 4 (Dec. 2012), 71–89.

[24] L. Huang and S. Teng. 2007. On the Approximation and Smoothed Complexity of Leontief Market Equilibria. In *Proceedings of the 1st Annual International Workshop of Frontiers in Algorithmics*. Springer *LNCS #4613*, 96–107.

[25] E. Irurozki, B. Calvo, and J. Lozano. 2019. Mallows and Generalized Mallows Model for Matchings. *Bernoulli* 25, 2 (2019), 1160–1188.

[26] L. Levin. 1986. Average Case Complete Problems. *SIAM J. Comput.* 15, 1 (1986), 285–286.

[27] X. Li and S. Sarkar. 2006. Privacy Protection in Data Mining: A Pertubation Approach to Catagorical Data. *Information Systems Research* 17, 3 (2006), 254–270.

[28] C. Lindner and J. Rothe. 2008. Fixed-Parameter Tractability and Parameterized Complexity, Applied to Problems From Computational Social Choice. In *Mathematical Programming Glossary*, A. Holder (Ed.). INFORMS Computing Society.

[29] L. Liu, M. Kantarcioglu, and B. Thuraisingham. 2008. The Applicability of the Perturbation Based Privacy Preserving Data Mining for Real-World Data. *Data & Knowledge Engineering* 65, 1 (2008), 5–21.

[30] T. Lu and C. Boutilier. 2014. Effective Sampling and Learning for Mallows Models with Pairwise-Preference Data. *The Journal of Machine Learning Research* 15, 1 (2014), 3783–3829.

[31] C. Mallows. 1957. Non-Null Ranking Models. *Biometrika* 44 (1957), 114–130.

[32] E. Markakis. 2017. Approximation Algorithms and Hardness Results for Fair Division. In *Trends in Computational Social Choice*, U. Endriss (Ed.). AI Access Foundation, 231–248.

[33] R. McKelvey and P. Ordeshook. 1978. A Decade of Experimental Research on Spatial Models of Elections and Committees. In *Advances in the Spatial Theory of Voting*, J. Enelow and M. Hinich (Eds.). Cambridge University Press, 99–144.

[34] T. Nguyen, M. Roos, and J. Rothe. 2013. A Survey of Approximability and Inapproximability Results for Social Welfare Optimization in Multiagent Resource Allocation. *Annals of Mathematics and Artificial Intelligence* 68, 1–3 (2013), 65–90.

[35] A. Procaccia and J. Rosenschein. 2007. Junta Distributions and the Average-Case Complexity of Manipulating Elections. *Journal of Artificial Intelligence Research* 28 (2007), 157–181.

[36] M. Regenwetter, B. Grofman, I. Tsetlin, and A. Marley. 2006. *Behavioral Social Choice: Probabilistic Models, Statistical Inference, and Applications*. Cambridge University Press.

[37] J. Rothe (Ed.). 2015. *Economics and Computation. An Introduction to Algorithmic Game Theory, Computational Social Choice, and Fair Division*. Springer-Verlag.

[38] J. Rothe. 2019. Borda Count in Collective Decision Making: A Summary of Recent Results. In *Proc. AAAI'19*. AAAI Press, 9830–9836.

[39] J. Rothe. 2019. How Can We Model Emotional and Behavioral Dynamics in Collective Decision Making? In *The Future of Economic Design*, J. Laslier, H. Moulin, R. Sanver, and W. Zwicker (Eds.). Springer, 245–251.

[40] J. Rothe and L. Schend. 2013. Challenges to Complexity Shields That Are Supposed to Protect Elections against Manipulation and Control: A Survey. *Annals of Mathematics and Artificial Intelligence* 68, 1–3 (2013), 161–193.

[41] M. 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.

[42] I. Schlotter, P. Faliszewski, and E. Elkind. 2011. Campaign Management under Approval-Driven Voting Rules. In *Proc. AAAI'11*. AAAI Press, 726–731.

[43] D. Spielman and S. Teng. 2004. Smoothed Analysis of Algorithms: Why the Simplex Algorithm Usually Takes Polynomial Time. *J. ACM* 51, 3 (2004), 385–463.

[44] D. Spielman and S. Teng. 2009. Smoothed Analysis: An Attempt to Explain the Behavior of Algorithms in Practice. *Commun. ACM* 52, 10 (2009), 76–84.

[45] T. Walsh. 2011. Is Computational Complexity a Barrier to Manipulation? *Annals of Mathematics and Artificial Intelligence* 62, 1–2 (2011), 7–26.

[46] T. Walsh. 2011. Where Are the Hard Manipulation Problems? *Journal of Artificial Intelligence Research* 42 (2011), 1–29.

[47] K. Wojtas and P. Faliszewski. 2012. Possible Winners in Noisy Elections. In *Proc. AAAI'12*. AAAI Press, 1499–1505.