

Preference Restrictions in Computational Social Choice: Recent Progress

Edith Elkind

University of Oxford
United Kingdom
elkind@cs.ox.ac.uk

Martin Lackner

University of Oxford
United Kingdom
martin.lackner@cs.ox.ac.uk

Dominik Peters

University of Oxford
United Kingdom
dominik.peters@cs.ox.ac.uk

Abstract

The goal of this short paper is to provide an overview of recent progress in understanding and exploiting useful properties of restricted preference domains, such as, e.g., the domains of single-peaked, single-crossing and 1-Euclidean preferences.

1 Introduction

Aggregating preferences of multiple agents by means of voting is often hard. One can identify two primary causes of this phenomenon. The first of them has to do with fundamental properties of collective preferences. Indeed, it was already known to the Marquis de Condorcet that even when individual voters are fully rational, their collective judgment may be irrational: if each voter ranks all alternatives in $\{a, b, c\}$ from best to worst, it may happen that a majority of voters prefer a to b , a majority of voters prefer b to c , yet a majority of voters prefers c to a , i.e., there is a cycle in majority preferences. In a similar spirit—but almost two centuries later—Arrow showed in his seminal work [1950] that when there are at least three alternatives, there is no perfect voting rule: he identified a small set of appealing axioms such that no voting rule for three or more alternatives can satisfy all of them. A closely related result was subsequently proved by Gibbard [1973] and Satterthwaite [1975], who observed that under any ‘reasonable’ voting rule there exists a scenario where some voter benefits from misrepresenting her preferences.

The results listed in the previous paragraph indicate that preference aggregation is hard from a conceptual perspective. However, it is also hard in a precise technical sense: there are many useful preference aggregation procedures whose output is NP-hard to compute. In particular, this is the case for the Kemeny rule, which is arguably the most natural voting rule to aggregate a set of preference rankings into a single ranking, as well as for two popular committee selection rules that provide fully proportional representation, namely, the Chamberlin–Courant rule and the Monroe rule.

Now, social choice theorists have observed that the first source of hardness can be circumvented by focusing on scenarios where voters’ preferences share some common structure. The most famous result of this type dates back to the

important early work of Black [1948] and says that if voters’ preferences are essentially single-dimensional, then there are no cycles in the majority preferences, and there is a voting rule that is strategyproof. The specific domain of preferences considered by Black is that of single-peaked preferences; similar results have been subsequently obtained for other restricted preference domains, such as those of preferences that are single-crossing or single-peaked on a tree (we provide formal definitions of these notions in Section 2).

It is then natural to ask whether the same approach can be used to circumvent computational complexity issues as well. The first foray in this direction was made by Walsh [Walsh, 2007], and since 2007 hardness and easiness results for various preference domains have been obtained for winner determination under a variety of voting rules, as well as for the problems of manipulation, control and bribery.

Interestingly, while purely social choice-theoretic issues (such as manipulability or majority cycles) vanish as soon as we assume that voters’ preferences belong to a suitable restricted domain, many of the algorithms for voting-related problems require the knowledge of the respective structural relationship among voters/alternatives (such as the order of candidates witnessing that the profile is single-peaked). This means that, to make use of these algorithms, one also needs an efficient procedure to discover whether a given preference profile has the required structural property and to find a respective witness. Consequently, the problem of designing such procedures has received a considerable amount of attention, too, resulting in polynomial-time algorithms for recognizing preferences that belong to several prominent restricted domains, as well as hardness results for some further domains.

The goal of this paper is to discuss both of these strands of algorithmic results, as well as to provide pointers to the literature; a longer literature survey that covers this area in further detail is in preparation.

2 Domain restrictions

Preliminaries For every positive integer n , set $[n] = \{1, \dots, n\}$. Let A be a finite set of *alternatives*, or *candidates*, and let $m = |A|$. A *weak order*, or *preference relation*, is a binary relation over A that is complete and transitive. A *linear order* is a weak order that, in addition, is antisymmet-

ric. Given a linear order v over A , we denote the top alternative in v by $\text{top}(v)$.

A *profile* $P = (v_1, \dots, v_n)$ over a set of alternatives A is a list of linear orders over A . We associate $P = (v_1, \dots, v_n)$ with a set of *voters* $N = [n]$; the order v_i is called the *vote* of voter i . For convenience, we write $a \succ_i b$ whenever $(a, b) \in v_i$, i.e., when voter i strictly prefers a to b . Given a profile P over A , we define its *majority relation* \succeq_{maj} as a weak order over A such that

$$a \succeq_{\text{maj}} b \iff |\{i \in N : a \succ_i b\}| \geq |\{i \in N : b \succ_i a\}|.$$

We write $a \succ_{\text{maj}} b$ if $a \succeq_{\text{maj}} b$, but not $b \succeq_{\text{maj}} a$. Alternative a is a *weak Condorcet winner* if $a \succeq_{\text{maj}} b$ for all $b \in A$; it is a *Condorcet winner* if $a \succ_{\text{maj}} b$ for all $b \in A$.

Single-Peaked Preferences The domain of single-peaked preferences was first defined by Black [1948]. It captures settings where there is a natural ordering over the alternatives, and voters' preferences are consistent with this order. Popular examples include voting on tax rates, the military budget, or simply the temperature in the room.

Let \triangleleft be a linear order over the set of alternatives A . A vote v over A is *single-peaked with respect to* \triangleleft if for every pair of candidates $a, b \in A$ with $\text{top}(v) \triangleleft b \triangleleft a$ or $a \triangleleft b \triangleleft \text{top}(v)$ it holds that v ranks b above a . A profile P over A is *single-peaked with respect to* \triangleleft if every vote in P is single-peaked with respect to \triangleleft ; P is *single-peaked* if there exists a linear order \triangleleft over A such that P is single-peaked with respect to \triangleleft . We refer to any such order \triangleleft as an *axis* for P .

The domain of single-peaked profiles has many attractive properties: it admits a family of voting rules that are not susceptible to manipulation (see, e.g., [Moulin, 1991]), and for any profile in this domain the majority relation is transitive (i.e., Condorcet's paradox is circumvented).

Single-Crossing Preferences In contrast to single-peaked profiles, single-crossing profiles are defined in terms of an ordering of the voters. The definition of this domain can be traced back to the work of Mirrlees [1971] and Roberts [1977] on income taxation.

A profile $P = (v_1, \dots, v_n)$ over A is *single-crossing with respect to the given order of voters* if for every pair of candidates $a, b \in C$ both sets $\{i \in [n] : a \succ_i b\}$ and $\{i \in [n] : b \succ_i a\}$ are contiguous subsets of $[n]$; P is *single-crossing* if the votes in P can be permuted so that P becomes single-crossing with respect to the resulting order of voters.

Single-crossing profiles have many of the same properties as single-peaked profiles. In particular, the majority relation of a single-crossing profile is transitive.

Euclidean Preferences Euclidean preferences capture settings where voters and alternatives can be identified with points in the Euclidean space, with voters' preferences driven by distances to alternatives. This domain was considered by Coombs [1950].

Formally, a profile P is *d-Euclidean* (where d is a positive integer) if there exists a map $x : N \cup A \rightarrow \mathbb{R}^d$ such that for all $i \in N$ and all $a, b \in A$ it holds that

$$a \succ_i b \implies \|x(i) - x(a)\| < \|x(i) - x(b)\|.$$

That is, voter i prefers those alternatives that are closer to her according to the embedding x . Here, $\|\cdot\|$ refers to the usual Euclidean ℓ_2 -norm on \mathbb{R}^d .

It is not hard to see that 1-Euclidean preferences are both single-peaked and single-crossing, with the respective orderings of voters/alternatives determined by their positions on the real line under x . Interestingly, the converse is not true, i.e., there exist profiles that are single-peaked and single-crossing, but not 1-Euclidean [Coombs, 1950; Elkind *et al.*, 2014; Chen *et al.*, 2015]. Another simple observation is that every profile is d -Euclidean for sufficiently large d . We also remark that d -Euclidean profiles with $d > 2$ do not necessarily have a Condorcet winner.

Preferences Single-Peaked/Single-Crossing on a Tree De-mange [1982] observes that if we place candidates on a tree rather than a line, and require the voters' preferences to be driven by candidates' positions on that tree, we obtain a large preference domain that nevertheless retains some of the attractive properties of the single-peaked domain.

Formally, let $T = (A, E)$ be a tree with vertex set A . A profile P over A is *single-peaked on T* if $P|_{T'}$ is single-peaked for every path $T' \subseteq T$. A profile P is *single-peaked on a tree* if there is a tree T such that P is single-peaked on T .

It can be shown that if a profile is single-peaked on a tree, it necessarily has a weak Condorcet winner; however, its majority relation need not be transitive.

In a similar manner we can define what it means for a profile to be single-crossing on a tree [Kung, 2015].

3 Recognition Algorithms

In this section, we will consider the problem of efficiently deciding whether a given preference profile belongs to one of the restricted domains listed in Section 2.

Single-Peaked Preferences Bartholdi III and Trick [1986] were the first to show that single-peaked preferences can be recognized in polynomial time. Their argument employed a reduction to the *consecutive ones* problem, which asks whether the columns of a given 0/1-matrix can be reordered so that in every row the 1s occur consecutively, i.e., as a contiguous block. Since the consecutive ones problem admits a linear-time algorithm [Booth and Lueker, 1976], this reduction implies that checking single-peakedness is possible in $O(m^2n)$ time. Later, Doignon and Falmagne [1994] developed a direct algorithm that runs in time $O(mn + n^2)$, and Escoffier *et al.* [2008] discovered an algorithm whose running time is $O(nm)$. The algorithm of Escoffier *et al.* [2008] builds up an underlying axis from the outside in, and is based on the crucial observation that alternatives that are ranked last by some voters must be placed on one of the extreme ends of the axis under construction.

Ballester and Haeringer [2011] show that single-peaked preferences can be characterized in terms of two *forbidden configurations*: they identify two constant-size profiles (one with two voters and four alternatives and one with three voters and three alternatives) such that a profile P is single-peaked if and only if neither of these two profiles is isomorphic to a subprofile of P . This observation provides an alternative polynomial-time algorithm for recognizing single-

peaked preferences: one can simply go over all subprofiles of a given size, and check that none of them is isomorphic to one of the two forbidden profiles.

Single-Crossing Preferences A simple way to recognize single-crossing preferences is to guess the leftmost vote, sort all other votes by their Kendall-tau distance to the first vote, and check if the resulting profile is single-crossing with respect to the resulting order of voters [Doignon and Falmagne, 1994; Elkind *et al.*, 2012]; this can be done in $O((m \log m)n^2)$ time. One can also recognize single-crossing preferences in $O(m^2n)$ time via a reduction to the consecutive ones problem [Bredereck *et al.*, 2013]. This domain also admits a characterization in terms of a small number of forbidden configurations [Bredereck *et al.*, 2013].

Euclidean Preferences Doignon and Falmagne [1994] present the first polynomial-time algorithm for recognizing 1-Euclidean profiles (this algorithm was later rediscovered by Elkind and Faliszewski [2014]). Their algorithm is based on the observation that any 1-Euclidean profile is single-crossing, and the single-crossing order of voters is essentially unique. It checks if the input profile is single-crossing, and if so, uses the respective ordering of voters to determine a possible ordering of the alternatives in \mathbb{R} ; the task of determining the actual positions of voters and alternatives in \mathbb{R} is then delegated to a linear (feasibility) program. Knoblauch [2010] describes a polynomial-time algorithm that is based on the fact that 1-Euclidean preferences are single-peaked; this algorithm, too, first places voters and alternatives on the line and then checks if the associated linear program has a feasible solution.

Interestingly, in contrast with single-peaked and single-crossing preferences, 1-Euclidean preferences cannot be characterized by a constant number of forbidden configurations [Chen *et al.*, 2015; Peters, 2016].

For $d \geq 2$, d -Euclidean preferences become quite unwieldy. The recognition problem is NP-hard for any fixed $d \geq 2$ and is, in fact, equivalent to the existential theory of the reals (ETR) [Peters, 2016]. Further, this problem is unlikely to be contained in NP, since there are d -Euclidean profiles all of whose embeddings need exponentially many bits to specify; the best known complexity upper bound is PSPACE.

Preferences Single-Peaked/Single-Crossing on Trees Trick [1989] describes an efficient algorithm for checking whether a given profile is single-peaked on some tree; interestingly, this algorithm may output a complicated tree even if the input profile is single-peaked on a path. One can then ask whether we can recognize in polynomial time if a given profile is single-peaked on a *specific* tree, or on some tree that has certain desirable properties, such as small diameter, low maximum degree or a constant number of leaves. It turns out that, while answering this question for a given tree is NP-hard [Peters and Elkind, 2016], a ‘nice’ tree can be identified efficiently for many underlying notions of ‘niceness’ [Yu *et al.*, 2013; Peters and Elkind, 2016]. Complementing these results, Kung [2015] provides an efficient algorithm for deciding if a given profile is single-crossing on some tree.

4 Algorithms for Social Choice Problems on Restricted Domains

Many popular single-winner rules, such as, e.g., Plurality, Borda, Copeland, Maximin and ranked pairs are computationally easy; three prominent exceptions are the Dodgson rule, the Kemeny rule, and the Young rule, which all have NP-hard winner determination problems, see [Bartholdi III *et al.*, 1989; Hemaspaandra *et al.*, 1997; Dwork *et al.*, 2001; Rothe *et al.*, 2003; Hemaspaandra *et al.*, 2005]. However, for each of these rules the winners can be computed in polynomial time for all restricted domains considered in this paper as long as the number of voters is odd. Indeed, these rules are Condorcet-consistent, i.e., they output the Condorcet winner when it exists, and for each of our domains it holds that every profile with an odd number of voters has a Condorcet winner. The profiles with an even number of voters are more challenging, since for some of these rules the set of winners may be a proper subset of the set of weak Condorcet winners; Brandt *et al.* [2015] show how to handle this case in polynomial time when voters’ preferences are single-peaked.

Most of the results mentioned in the previous paragraph do not rely on having a witness that the input belongs to the respective domain (such as a single-peaked axis or a single-crossing order of the voters). In contrast, restricted-domain algorithms for two important committee selection rules, namely, the Chamberlin–Courant rule and the Monroe rule, tend to make heavy use of this information. Typically, they proceed by dynamic programming, where the structure of the dynamic program is driven by the respective ordering of voters/alternatives. This is the case, for instance, for the algorithms that determine Chamberlin–Courant winners and egalitarian Monroe winners under single-peaked preferences [Betzler *et al.*, 2013] or Chamberlin–Courant winners under single-crossing preferences [Skowron *et al.*, 2016]. Interestingly, the classic variant of the Monroe rule is surprisingly resistant to domain restrictions: it remains hard for both single-peaked [Betzler *et al.*, 2013] and single-crossing [Skowron *et al.*, 2016] preferences.

Restricting the preferences to be single-peaked on a tree simplifies the winner determination problem for some committee selection rules (such as the egalitarian variant of the Chamberlin–Courant rule); however, to obtain an efficient algorithm for the classic variant of the Chamberlin–Courant rule we need to place additional constraints on the structure of the tree [Yu *et al.*, 2013; Peters and Elkind, 2016]. In contrast, if a profile is single-crossing on a tree, classic Chamberlin–Courant winners can be computed efficiently [Clearwater *et al.*, 2015], irrespective of the shape of that tree.

Finally, many problems concerning strategic behavior in voting (such as manipulation, control and bribery) also become easier when voters’ preferences are single-peaked or single-crossing [Walsh, 2007; Faliszewski *et al.*, 2011; Brandt *et al.*, 2015; Faliszewski *et al.*, 2014; Magiera and Faliszewski, 2014; Erdélyi *et al.*, 2015; Elkind *et al.*, 2016].

References

[Arrow, 1950] K. J. Arrow. A difficulty in the concept of social welfare. *The Journal of Political Economy*, pages 328–346, 1950.

- [Ballester and Haeringer, 2011] M. A. Ballester and G. Haeringer. A characterization of the single-peaked domain. *Social Choice and Welfare*, 36(2):305–322, 2011.
- [Bartholdi III and Trick, 1986] J. Bartholdi III and M. A. Trick. Stable matching with preferences derived from a psychological model. *Operation Research Letters*, 5(4):165–169, 1986.
- [Bartholdi III et al., 1989] J. Bartholdi III, C. A. Tovey, and M. A. Trick. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare*, 6(2):157–165, 1989.
- [Betzler et al., 2013] N. Betzler, A. Slinko, and J. Uhlmann. On the computation of fully proportional representation. *Journal of Artificial Intelligence Research*, 47(1):475–519, 2013.
- [Black, 1948] D. Black. On the rationale of group decision-making. *The Journal of Political Economy*, pages 23–34, 1948.
- [Booth and Lueker, 1976] K. S. Booth and G. S. Lueker. Testing for the consecutive ones property, interval graphs, and graph planarity using PQ-tree algorithms. *Journal of Computer and System Sciences*, 13(3):335–379, 1976.
- [Brandt et al., 2015] F. Brandt, M. Brill, E. Hemaspaandra, and L. A. Hemaspaandra. Bypassing combinatorial protections: Polynomial-time algorithms for single-peaked electorates. *Journal of Artificial Intelligence Research*, 53:439–496, 2015.
- [Bredereck et al., 2013] R. Bredereck, J. Chen, and G. J. Woeginger. A characterization of the single-crossing domain. *Social Choice and Welfare*, 41(4):989–998, 2013.
- [Chen et al., 2015] J. Chen, K. Pruhs, and G. J. Woeginger. The one-dimensional Euclidean domain: Finitely many obstructions are not enough. Technical Report arXiv:1506.03838 [cs.GT], arXiv.org, 2015.
- [Clearwater et al., 2015] A. Clearwater, C. Puppe, and A. Slinko. Generalizing the single-crossing property on lines and trees to intermediate preferences on median graphs. In *IJCAI*, pages 32–38, 2015.
- [Coombs, 1950] C. H. Coombs. Psychological scaling without a unit of measurement. *Psychological review*, 57(3):145, 1950.
- [Demange, 1982] G. Demange. Single-peaked orders on a tree. *Mathematical Social Sciences*, 3(4):389–396, 1982.
- [Doignon and Falmagne, 1994] J.-P. Doignon and J.-C. Falmagne. A polynomial time algorithm for unidimensional unfolding representations. *Journal of Algorithms*, 16(2):218–233, 1994.
- [Dwork et al., 2001] C. Dwork, R. Kumar, M. Naor, and D. Sivakumar. Rank aggregation methods for the web. In *WWW*, pages 613–622, 2001.
- [Elkind and Faliszewski, 2014] E. Elkind and P. Faliszewski. Recognizing 1-Euclidean preferences: An alternative approach. In *SAGT*, pages 146–157, 2014.
- [Elkind et al., 2012] E. Elkind, P. Faliszewski, and A. Slinko. Clone structures in voters’ preferences. In *ACM EC*, pages 496–513, 2012.
- [Elkind et al., 2014] E. Elkind, P. Faliszewski, and P. Skowron. A characterization of the single-peaked single-crossing domain. In *AAAI’14*, pages 654–660, 2014.
- [Elkind et al., 2016] E. Elkind, E. Markakis, S. Obraztsova, and P. Skowron. Complexity of finding equilibria of plurality voting under structured preferences. In *AAMAS*, 2016.
- [Erdélyi et al., 2015] G. Erdélyi, M. Lackner, and A. Pfandler. Manipulation of k-approval in nearly single-peaked electorates. In *ADT*, pages 71–85, 2015.
- [Escoffier et al., 2008] B. Escoffier, J. Lang, and M. Öztürk. Single-peaked consistency and its complexity. In *ECAI*, volume 8, pages 366–370, 2008.
- [Faliszewski et al., 2011] P. Faliszewski, E. Hemaspaandra, L. A. Hemaspaandra, and J. Rothe. The shield that never was: Societies with single-peaked preferences are more open to manipulation and control. *Information and Computation*, 209(2):89–107, 2011.
- [Faliszewski et al., 2014] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. The complexity of manipulative attacks in nearly single-peaked electorates. *Artificial Intelligence Journal*, 207:69–99, 2014.
- [Gibbard, 1973] A. Gibbard. Manipulation of voting schemes: a general result. *Econometrica*, pages 587–601, 1973.
- [Hemaspaandra et al., 1997] E. Hemaspaandra, L. A. Hemaspaandra, and J. Rothe. Exact analysis of Dodgson elections: Lewis Carroll’s 1876 voting system is complete for parallel access to NP. *Journal of the ACM*, 44(6):806–825, 1997.
- [Hemaspaandra et al., 2005] E. Hemaspaandra, H. Spakowski, and J. Vogel. The complexity of Kemeny elections. *Theoretical Computer Science*, 349(3):382–391, 2005.
- [Knoblauch, 2010] V. Knoblauch. Recognizing one-dimensional Euclidean preference profiles. *Journal of Mathematical Economics*, 46(1):1–5, 2010.
- [Kung, 2015] F.-C. Kung. Sorting out single-crossing preferences on networks. *Social Choice and Welfare*, 44(3):663–672, 2015.
- [Magiera and Faliszewski, 2014] K. Magiera and P. Faliszewski. How hard is control in single-crossing elections? In *ECAI*, pages 579–584, 2014.
- [Mirrlees, 1971] J. A. Mirrlees. An exploration in the theory of optimum income taxation. *The Review of Economic Studies*, 38(2):175–208, 1971.
- [Moulin, 1991] H. Moulin. *Axioms of Cooperative Decision Making*. Cambridge University Press, 1991.
- [Peters and Elkind, 2016] D. Peters and E. Elkind. Preferences single-peaked on nice trees. In *AAAI*, 2016.
- [Peters, 2016] D. Peters. Recognising multidimensional Euclidean preferences. *arXiv preprint arXiv:1602.08109*, 2016.
- [Roberts, 1977] K. W. S. Roberts. Voting over income tax schedules. *Journal of Public Economics*, 8(3):329–340, 1977.
- [Rothe et al., 2003] J. Rothe, H. Spakowski, and J. Vogel. Exact complexity of the winner problem for young elections. *Theory of Computing Systems*, 36(4):375–386, 2003.
- [Satterthwaite, 1975] M. A. Satterthwaite. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, 1975.
- [Skowron et al., 2016] P. Skowron, L. Yu, P. Faliszewski, and E. Elkind. The complexity of fully proportional representation for single-crossing electorates. *Theoretical Computer Science*, 569:43–57, 2016.
- [Trick, 1989] M. A. Trick. Recognizing single-peaked preferences on a tree. *Mathematical Social Sciences*, 17(3):329–334, 1989.
- [Walsh, 2007] T. Walsh. Uncertainty in preference elicitation and aggregation. In *AAAI*, volume 7, pages 3–8, 2007.
- [Yu et al., 2013] L. Yu, H. Chan, and E. Elkind. Multiwinner elections under preferences that are single-peaked on a tree. In *IJCAI*, pages 425–431, 2013.