

Preferences Single-Peaked on a Circle

Dominik Peters and Martin Lackner

Department of Computer Science
University of Oxford, UK
{dominik.peters, martin.lackner}@cs.ox.ac.uk

Abstract

We introduce the domain of preferences that are single-peaked on a circle, which is a generalization of the well-studied single-peaked domain. This preference restriction is useful, e.g., for scheduling decisions, and for one-dimensional decisions in the presence of extremist preferences. We give a fast recognition algorithm of this domain, provide a characterisation by finitely many forbidden subprofiles, and show that many popular single- and multi-winner voting rules are polynomial-time computable on this domain. In contrast, Kemeny’s rule remains hard to evaluate, and several impossibility results from social choice theory can be proved using only profiles that are single-peaked on a circle.

1 Introduction

A central problem in the study of multi-agent systems is the aggregation of agents’ preferences in order to make group decisions. Impossibility theorems and computational hardness result make this problem a hard one to solve. However, a successful line of research going back to Black’s (1948) seminal article has managed to circumvent many problems in (computational) social choice for the special case when agents’ preferences are *single-peaked*. Under this preference restriction, we assume that agents have preferences over the possible values of a one-dimensional quantity such as the timing of a deadline, a tax rate, a thermostat setting, or the price of a new product. A preference ordering is *single-peaked* if an agent has a certain most-preferred value of the quantity and derives less and less satisfaction from values that are further away from the subjective optimum. Another popular application of this setting is in political elections, where it is often held that candidates can be ordered on a left-to-right spectrum making the voters’ preferences single-peaked.

Preference profiles that consist solely of single-peaked preference orderings have attractive properties, both algorithmically and in terms of their social choice behaviour (Elkind, Lackner, and Peters 2016). For example, winner determination problems that are computationally hard in the general case tend to be easy when restricted to single-peaked profiles (Brandt et al. 2015; Betzler, Slinko, and Uhlmann 2013), and the single-peaked domain admits a strategyproof voting rule




Figure 1: Example of preferences single-peaked on a circle.


and guarantees the existence of Condorcet winners as well as transitivity of the majority relation (Moulin 1991).

The usefulness of results of this type is limited by the extent to which profiles in practice actually happen to be single-peaked. One way of dealing with this is to consider less restrictive generalisations of single-peakedness. Maybe the structure of the alternative space is not quite one-dimensional, and in this case it might be useful to consider preferences that are single-peaked on a *tree* (Demange 1982). This domain is notably larger, yet still retains many desirable properties in social choice terms; however, its algorithmic usefulness is more mixed (Yu, Chan, and Elkind 2013; Peters and Elkind 2016).

In this paper, we identify a new preference restriction: being single-peaked on a *circle*. Here we assume that alternatives can be placed on a circle, with agents’ preferences again being decreasing on both sides of their peaks. See Figure 1 for some example shapes that ‘preference curves’ might have in this setting; higher points are more preferred. Note that the circle wraps around, and so h and a are adjacent. Intuitively, a preference profile is single-peaked on a circle if, for every agent, we can ‘cut’ the cycle once so that the agent’s preferences are single-peaked on the resulting line. Crucially, the location of the cutting point may differ for each agent.

The aim of this paper is to explore this new preference domain in detail. We will find that this domain is algorithmically useful (it often allows for efficient winner determination), but it performs less convincingly in terms of axiomatic properties (since voting paradoxes still occur and impossibility results can still be proven). Interestingly, this is precisely opposite to how the results turned out for single-peakedness on trees.

Motivating Examples. There are many practical scenarios where we might expect preferences to be single-peaked on a circle. This is even the case when, on first sight, there seems to be no circle anywhere. Indeed, suppose that alternatives are naturally ordered on a line; we may pretend this line is a circle by joining up its endpoints. Of course, every order that is single-peaked on the line is also single-peaked on the circle. But crucially, the *reverse* of such an order, now single-caved on the line, is still single-peaked on the same circle. Thus, our new preference restriction allows us to *combine* single-peaked and single-caved votes (as shown on the right). One interpretation is that this move allows us to accommodate “extremists”. For example, while most people have a sweet spot somewhere on the left-right political axis, some people might dislike centrist options and prefer the extremes. When planning a vacation, some might have an optimal climate in mind, while others like it both very cold (skiing) and very hot (beaches), but dislike compromises (England).



Other examples of alternative spaces are more explicitly cyclic. Consider, for example, finding a good time for a daily event (such as a day or night shift, or a meeting, or the timing of backups) where possibilities are arranged in a 24 hour cycle. A similar structure exists when scheduling an international phone call; here, different time zones are arranged along the equator, and lead to cyclic preferences.

But perhaps the most appealing example of preferences that can be expected to be single-peaked on a circle come from problems inspired by facility location. Rather many structures have a boundary that is (roughly) isomorphic to a cycle, including most cities and countries. The problem of deciding where to locate a new airport for a city would be one example, since airports are usually positioned on the boundary. Similarly, where should a company build new office space? To which coastal region should a family move? Where do we want to sit in a football stadium? Another plausible application could be inspired by security concerns, if we consider the placement of border security checkpoints.

Contributions. We formally define single-peakedness on circles and immediately extend this definition to preferences with ties, and to dichotomous (approval) preferences.

We show that it is possible to efficiently recognise whether a given preference profile is single-peaked on some circle, and if so return a suitable circle. For the case of preferences without ties, we give a recognition algorithm that runs in linear time, matching the performance in the case for the line.

We characterise the domain of preferences single-peaked on a circle through a list of finitely many *forbidden subprofiles* with 2 voters and 5 alternatives, and with 3 voters and 4 alternatives. The proof of this characterisation implies that our linear-time recognition algorithm can *certify* its negative answers by exhibiting a forbidden subprofile.

While single-peakedness on a line serves as a way to circumvent many impossibility results in social choice, we show that such impossibilities (including the Gibbard–Satterthwaite theorem) can still be proven when preferences

are allowed to be single-peaked on a circle.

We then study the algorithmic properties of our new preference extension. We show that Young’s voting rule (and also Young *scores*) can be efficiently computed if preferences are single-peaked on a circle; this algorithm also improves upon the state-of-the-art when it comes to preferences single-peaked on a line. We also show that several multi-winner voting rules are efficiently computable in our restricted case, including the Chamberlin–Courant rule and Proportional Approval Voting (PAV). These results rely on a recent technique using total unimodularity and integer programming (Peters 2016). On the other hand, we show that McGarvey’s theorem (McGarvey 1953) can be proven using only preferences single-peaked on a circle; thus, Kemeny’s rule remains NP-hard to evaluate even in this setting.

2 Definition

Let A be a finite set of *alternatives* (or *candidates*). A *weak order* (or *preference relation*) is a binary relation \succsim over A which is complete and transitive. A *linear order* is a weak order that is antisymmetric, and so does not allow preference ties; a *strict linear order* \succ is the irreflexive part of a linear order. A *profile* $P = (v_1, \dots, v_n)$ over A is a list of weak orders over A . The elements of $N = \{1, \dots, n\}$ are called *voters*, and we associate voter $i \in N$ with the order v_i , which we call the *vote* of voter i . For convenience, we write $a \succsim_i b$ whenever $(a, b) \in v_i$, i.e., when voter i weakly prefers alternative a to alternative b . We also use \succ_i and \sim_i for the strict and indifference parts of \succsim_i . We will always write m for the number of alternatives and n for the number of voters. If v_i is a linear order, we write $\text{top}(v_i)$ for i ’s most-preferred alternative.

An *axis* \triangleleft is a strict linear order of the alternatives. We usually think of an axis as describing the underlying one-dimensional structure of the alternative space. A linear order v_i is *single-peaked* with respect to the axis \triangleleft if for each pair of alternatives $a, b \in A$ with $\text{top}(v_i) \triangleleft b \triangleleft a$ or $a \triangleleft b \triangleleft \text{top}(v_i)$ it holds that $b \succ_i a$. Let us also give another, equivalent definition. An *interval* $I \subseteq A$ of an axis \triangleleft is any set such that for all $a, b, c \in A$, if we have $a, c \in I$ and $a \triangleleft b \triangleleft c$, then $b \in I$. Then a vote v_i is single-peaked with respect to the axis \triangleleft if and only if for every $c \in A$, the top-initial segment $\{a \in A : a \succ_i c\}$ is an interval of \triangleleft . This definition in terms of intervals immediately gives a definition of the single-peaked property for *weak* orders as well. There are several possible definitions of single-peakedness for weak orders; the one we use here is often referred to as *possible single-peakedness* (Lackner 2014), since it is equivalent to saying that there exists a linearisation of the weak order which is single-peaked.

We say that two axes \triangleleft and \triangleleft' are *cyclically equivalent* if there is $l \in [m]$ such that we can write $a_1 \triangleleft a_2 \triangleleft a_3 \triangleleft \dots \triangleleft a_m$ and $a_l \triangleleft' a_{l+1} \triangleleft' \dots \triangleleft' a_m \triangleleft' a_1 \triangleleft' \dots \triangleleft' a_{l-1}$. For an axis \triangleleft , we then define the *circle* $C(\triangleleft)$ of \triangleleft to be the set of axes cyclically equivalent to \triangleleft . Any set C of axes that can be written as $C = C(\triangleleft)$ for some \triangleleft we call a *circle*. For example, $C = \{a \triangleleft b \triangleleft c, b \triangleleft' c \triangleleft' a, c \triangleleft'' a \triangleleft'' b\}$ is a circle. Note that “cutting” a circle C at a point yields an axis $\triangleleft \in C$. We say that \triangleleft *starts in* $a \in A$ if $a \triangleleft b$ for all $b \in A \setminus \{a\}$.

Definition. Let C be a circle. A vote v_i is single-peaked on C if there is an axis $\triangleleft \in C$ such that v_i is single-peaked with respect to \triangleleft . A preference profile P is single-peaked on a circle (SPOC) if there exists a circle C such that every vote $v_i \in P$ is single-peaked on C .

Intuitively, a vote v_i is single-peaked on C if C can be cut so that v_i is single-peaked on the resulting line.

Again let us state another equivalent definition. An interval $I \subseteq A$ of a circle C is a set that is an interval of one of the axes $\triangleleft \in C$ of the circle. Then a vote is single-peaked on a circle C if and only if each top-initial segment $\{a \in A : a \succ_i c\}$ is an interval of C . Note that the complement $A \setminus I$ of an interval I of C is again an interval. Thus, a weak order \succ is single-peaked on C if and only if its reverse $\succ^{-1} = \{(b, a) : (a, b) \in \succ\}$ is also single-peaked on C .

A vote is *single-caved* if its reverse is single-peaked. It follows, then, that mixtures of single-peaked and single-caved orders (on the same axis) are SPOC. However, not all SPOC profiles have this form; one such example is the profile shown in Figure 1, where the circle cannot be cut so as to make every preference curve either single-peaked or single-caved.

3 Recognition Algorithms

In this section we design algorithms that decide whether a given profile is single-peaked on some circle, and if so, return a suitable circle C .

A matrix $M = (a_{ij})$ with $a_{ij} \in \{0, 1\}$ has the *consecutive ones property* if the columns of M can be put into a linear order \triangleleft so that for every row of M , the columns with 1-entries form an interval of \triangleleft . The matrix M has the *circular ones property* if its columns can be arranged in a circle C so that the 1-entries of each row form an interval of C . Given our definitions in terms of intervals above, it is straightforward to translate a profile P of weak orders into an $mn \times m$ matrix M so that P is single-peaked [single-peaked on a circle] if and only if M has the consecutive [circular] ones property (Bartholdi III and Trick 1986): Take one column for each alternative, and one row for every top-initial segment of every voter in P ; the row is the incidence vector of the segment. Since it is possible to check in linear time whether a matrix A has the consecutive or circular ones property (Booth and Lueker 1976), this gives us an $O(m^2n)$ algorithm to recognise profiles that are single-peaked on a circle.

In the remainder of this section, we design a more explicit algorithm that runs in time $O(mn)$ when the input profile consists of *linear orders*.¹

Suppose $P = (v_1, \dots, v_n)$ is a profile of linear orders over A , and fix some alternative $z \in A$. We will build another profile $P' = (v_1^u, v_1^l, \dots, v_n^u, v_n^l)$ of $2n$ weak orders by *slicing* each vote v_i at z into an *upper* part v_i^u and a *lower* part v_i^l . The upper part v_i^u ranks all alternatives a such that $a \succ_i z$ in order of \succ_i , and puts all remaining alternatives into a least-preferred indifference class. The lower part v_i^l ranks all alternatives a such that $z \succ_i a$ in reverse order of \succ_i , and

again puts all remaining alternatives into a least-preferred indifference class.

Example. Slicing the order $a \succ b \succ c \succ z \succ d \succ e \succ f$ at z yields the upper part $a \succ^u b \succ^u c \succ^u z \sim^u d \sim^u e \sim^u f$ and the lower part $f \succ^l e \succ^l d \succ^l z \sim^l a \sim^l b \sim^l c$.

The notion of slicing reduces SPOC to single-peakedness:

Proposition 1. Suppose a profile P' of weak orders is obtained by slicing each vote of a profile P of linear orders at some fixed $z \in A$. Then P is SPOC if and only if the profile P' is single-peaked.

Proof. Suppose P is SPOC on C , and let $\triangleleft \in C$ be an axis starting in z . Since z is least-preferred by all voters in P' , z is not contained in any top-initial segment of any voter in P' . However, all top-initial segments of votes in P' are intervals of C . Since they do not contain z , they must also be intervals of \triangleleft . Thus, P' is single-peaked with respect to \triangleleft .

Suppose P' is single-peaked with respect to \triangleleft . We show that P is SPOC on $C = C(\triangleleft)$. Take a top-initial segment S of a vote v_i in P ; we prove that S is an interval of C . If $z \notin S$, then S is a top-initial segment of v_i^u in P' . Thus, S is an interval of \triangleleft and so an interval of C . If however $z \in S$, then the complement $A \setminus S$ is a top-initial segment of v_i^l in P' , hence an interval of \triangleleft , and so $A \setminus S$ is an interval of C . But the complement of an interval of a circle is again an interval, and so S is an interval of C . Hence P is SPOC. \square

Thus, we can use an algorithm that decides whether a profile of weak orders is single-peaked to decide whether a profile of linear orders is SPOC. Next, note that if we select $z \in A$ to be the alternative that is ranked last by v_1 (say), then the profile P' obtained by slicing P at z contains a linear order (namely the upper part of v_1). Lackner (2014) has given a $O(mn)$ time algorithm that decides whether a profile of weak orders containing at least one linear order is single-peaked. Since P' can be constructed from P in time $O(mn)$, by Proposition 1, we obtain the following.

Theorem 2. There is an $O(mn)$ time algorithm that decides whether a profile of linear orders is single-peaked on a circle.

4 Characterisation by Forbidden Subprofiles

Ballester and Haeringer (2011) have characterised the domain of single-peaked profiles of linear orders by a finite collection of forbidden subprofiles. More precisely, they gave forbidden profiles with 3 voters and 3 alternatives, and with 2 voters and 4 alternatives such that a profile P is *not* single-peaked if and only if it is possible to obtain one of their forbidden profiles from P by deleting and reordering voters, and deleting and renaming alternatives. A similar characterisation exists for *single-crossing* profiles (Bredereck, Chen, and Woeginger 2013), but no finite characterisation exists for *d-Euclidean* profiles (Chen, Pruhs, and Woeginger 2015; Peters 2017).

Here, we prove that a profile is SPOC unless it contains certain forbidden subprofiles with 2 voters and 5 alternatives and with 3 voters and 4 alternatives. Let us write $B \succ_i C$ to mean that $b \succ_i c$ for all $b \in B$ and $c \in C$.

Theorem 3. A profile P of linear orders on A is not SPOC if and only if one of the following three cases occurs.

¹ Actually, the algorithm works whenever P contains at least one linear order.

1. There are distinct alternatives $a, b, c, d, e \in A$ and voters v_i and v_j in P such that

$$\begin{aligned} \{a, b\} &\succ_i \{c\} \succ_i \{d, e\}, \\ \{a, e\} &\succ_j \{c\} \succ_j \{d, b\}. \end{aligned}$$

2. There are distinct alternatives $a, b, c, d \in A$ and voters v_i, v_j , and v_k in P such that

$$\begin{aligned} \{a, b\} &\succ_i \{c, d\}, \\ \{a, c\} &\succ_j \{b, d\}, \\ \{a, d\} &\succ_k \{b, c\}. \end{aligned}$$

3. There are distinct alternatives $a, b, c, d \in A$ and voters v_i, v_j , and v_k in P such that

$$\begin{aligned} \{a, b\} &\succ_i \{c, d\}, \\ \{b, c\} &\succ_j \{a, d\}, \\ \{c, a\} &\succ_k \{b, d\}. \end{aligned}$$

Proof. Sufficiency. We prove that if one of the three cases occurs, then P is not SPOC. Since SPOC is closed under alternative deletion, in each case, we may assume wlog that P only involves alternatives mentioned in the forbidden condition. Suppose P was single-peaked on the circle C .

1. Considering top-initial segments of size 2, we see that a must have neighbours b and e in C . Considering top-initial segments of size 3, and taking complements, we see that d must have neighbours b and e in C . But this uniquely determines a circle with $a \triangleleft b \triangleleft d \triangleleft e \triangleleft a$; yet this circle does not include c , a contradiction.
2. Considering top-initial segments of size 2, we see that a must have neighbours b, c , and d in C . But no vertex of a circle has three neighbours, a contradiction.
3. Since P is single-peaked on C , so is P with every order reversed. But after reversing every order, we are again in case 2: d must have three neighbours.

Necessity. This direction is much more involved, and we leave the details to the full version of this paper. The proof strategy is as follows: if P is not SPOC, then the recognition algorithm of Theorem 2, run on input P , will return a negative answer. We analyse every way the algorithm of Lackner (2014) could answer negatively, and in each case construct a witnessing forbidden structure from among those identified in the theorem statement. \square

For the benefit of future research, let us briefly describe how we obtained Theorem 3. We first implemented a recognition algorithm for SPOC profiles (Theorem 2) and then iterated through all possible profiles of certain sizes (up to isomorphism), checking for each whether they were *minimal counterexamples*: not SPOC, but every profile obtained by deleting a voter or an alternative is SPOC. We analysed the resulting list by hand to come up with the compact representation in Theorem 3. The proof strategy of case analysis of the “no”-conditions of a recognition algorithm is (implicitly) also the approach used in previous characterisations (Ballester and Haeringer 2011; Brederick, Chen, and Woeginger 2013; Cornaz, Galand, and Spanjaard 2012).

5 Impossibility Theorems

One of the major advantages of the traditional single-peaked domain is the existence of a non-manipulable voting rule on this domain: The well-known median voter procedure sorts voters’ most preferred alternatives according to the axis \triangleleft and then returns the median alternative a . This alternative is, in fact, a (weak) Condorcet winner: for any other alternative b , a (weak) majority of voters prefers a to b . One might hope to be able to extend this procedure to circles, but this turns out to be impossible: the Gibbard–Satterthwaite theorem can be proven using only profiles that are single-peaked on a circle.

A *resolute voting rule* f on SPOC profiles is a function assigning a single winning alternative to every SPOC profile of linear orders. The rule f is *non-dictatorial* if there is no fixed voter i such that f always picks i ’s top alternative. The profile obtained from P by replacing vote v_i by v'_i is denoted by (P_{-i}, v'_i) . A voting rule f on SPOC profiles is *strategyproof* if $f(P) \succ_i f(P_{-i}, v'_i)$ for all orders v'_i such that (P_{-i}, v'_i) is still SPOC.

Theorem 4 (Gibbard–Satterthwaite Theorem for SPOC). *There is no resolute voting rule on SPOC profiles that is non-dictatorial, onto, and satisfies strategyproofness.*

Proof. This follows immediately from the results of Kim and Roush (1980) and Sato (2010), who prove this result for an even more restricted domain consisting only of the $2m$ orders which traverse the circle clockwise and counter-clockwise starting from every possible alternative. \square

Note that the SPOC orders used in this proof are ‘un-balanced’, in that the most- and least-preferred alternatives are adjacent on the circle for every agent. Still, a similar dictatorship result can be proved even using orders that are ‘Euclidean’ on a circle, where preferences decrease uniformly in both directions from the peak (Schummer and Vohra 2002). It can also be shown that, with these Euclidean orders, the *random dictatorship* rule is *group-strategyproof* (Alon et al. 2010b), and there is an intriguing randomized mechanism that is strategyproof and provides a $3/2$ -approximation to the egalitarian social welfare (Alon et al. 2010a).

Another desirable axiomatic property is *participation*, which, intuitively, states that no voter can strictly benefit by abstaining from an election. A celebrated result of Moulin (1988) shows that this property is incompatible with Condorcet-consistency, which requires the voting rule to select the Condorcet winner if one exists. This result can also be proven using only SPOC profiles.

Theorem 5 (No-Show Paradox for SPOC). *For $m \geq 4$, there is no resolute voting rule on SPOC profiles that is Condorcet-consistent and satisfies participation.*

Proof. For $m = 4$ alternatives, one can check that the original proof of Moulin (1988) only uses SPOC preference orders; the same holds for the alternative proofs due to Brandt, Geist, and Peters (2016). More care is needed to extend this to $m \geq 5$; we leave this to the full version of this paper. \square

As described in the next section, further impossibilities about tournament-based rules can be deduced from Lemma 6.

6 Kemeny's and Young's Rules

In this section, we will consider the problem of determining an election winner according to two well-known voting rules, Young's rule and Kemeny's rule, that are NP-hard to evaluate in general (Bartholdi III, Tovey, and Trick 1989; Rothe, Spakowski, and Vogel 2003; Hemaspaandra, Spakowski, and Vogel 2005). We will be interested to see whether these problems can be solved in polynomial time for SPOC profiles. We leave the complexity of Dodgson's rule for SPOC profiles for future work.

Kemeny's rule is a *rank aggregation* rule: Given a profile P over A , its aim is to produce a *consensus* ranking over A . Suppose r is a linear order over A . Its *Kemeny score* is $\sum_{i \in N} |v_i \cap r|$, the number of pairwise agreements of r with P . A *Kemeny ranking* is a linear order r with maximum Kemeny score. While it is NP-hard to find a Kemeny ranking (Bartholdi III, Tovey, and Trick 1989), this problem is easy for single-peaked profiles whose transitive majority relation is easily seen to give rise to a Kemeny ranking. For SPOC preferences, the situation is less clear: the Condorcet paradox profile $(x \succ_1 y \succ_1 z, y \succ_2 z \succ_2 x, z \succ_3 x \succ_3 y)$ on 3 alternatives is SPOC, so SPOC does not guarantee a transitive majority relation. In fact, SPOC does not guarantee *anything* at all about the majority relation.

Lemma 6 (McGarvey's theorem for SPOC). *All (weighted) majority tournaments are inducible by SPOC profiles.*

Proof. Fix a circle C with $x_1 \triangleleft x_2 \triangleleft \dots \triangleleft x_m$. For any arc (x_i, x_j) of the target tournament, consider the following two votes which are single-peaked on C : (subscripts modulo m)

$$\begin{aligned} & x_{i+1} \succ \dots \succ x_{j-1} \succ x_i \succ x_j \succ x_{j+1} \succ \dots \succ x_{i-1} \\ & x_{x-1} \succ \dots \succ x_{j+1} \succ x_i \succ x_j \succ x_{j-1} \succ \dots \succ x_{i+1} \end{aligned}$$

These two votes induce a majority arc $x_i \rightarrow x_j$ with weight 2, but all other arcs have weight 0. By combining pairs of such votes, any tournament can be obtained. If odd edge weights are desired, start with an arbitrary single order, and then use pairs as above to adjust the weights as needed. \square

Recall that Kemeny scores only depend on the weighted majority relation of a profile. Since the profiles in the proof of McGarvey's theorem above can be produced in polynomial time, the hardness of Kemeny in the general case carries over.

Theorem 7. *Finding a Kemeny ranking is NP-hard, even for SPOC preferences.*

Indeed, by the same argument essentially all negative (axiomatic or computational) results in the sphere of voting rules based on (weighted) tournaments (see Brandt, Brill, and Harrenstein 2016; Fischer, Hudry, and Niedermeier 2016) still hold restricted to SPOC preferences.

Young's rule. Given a profile P over A , an alternative $c \in A$ is a *Condorcet winner* if for every $b \in A \setminus \{c\}$, a majority of voters in P strictly prefers c to b . The *Young score* of an alternative $c \in A$ is the minimum number of voters that have to be deleted from P so that c becomes a Condorcet winner. Then, Young's rule selects all alternatives with minimum Young score as winners. It is known that Young winners can

be found in polynomial time for single-peaked preferences (Brandt et al. 2015), since in this case Condorcet winners always exists when the number of voters n is odd; and the case with n even is also handled easily.

Because SPOC does not guarantee the existence of a Condorcet winner, a different approach is needed. We will use the interpretation of SPOC in terms of intervals of the underlying circle to give a polynomial-time algorithm that calculates the Young score of every alternative; clearly this algorithm can then be run repeatedly to find a Young winner. Of course, our algorithm also works for preferences single-peaked on a line; while the algorithm of Brandt et al. (2015) returns only a Young winner, our algorithm can find the Young score of *any* alternative. Note that precise definitions of Young scores differ slightly: sometimes it is only required that an alternative be made a *weak* Condorcet winner through voter deletion; our algorithm can be easily adapted for this alternative definition.

Theorem 8. *For SPOC profiles, the Young score of an alternative can be computed in $O(mn^2)$ time.*

Proof. We fix an axis $\triangleleft \in C$ that starts with the alternative a whose Young score we want to compute; let $a \triangleleft b \triangleleft \dots \triangleleft c$ (b is the candidate right of a , c is the rightmost candidate). We partition voters into two sets: $N_1 = \{i \in N : b \succ_i a\}$ and $N_2 = N \setminus N_1$. Since P is SPOC, for any voter i , the set $I_i := \{d \in A : d \succ_i a\}$ forms an interval of \triangleleft . Voters in N_1 correspond to intervals containing b ; voters in N_2 correspond to intervals containing c but not b , and to empty intervals.

The idea behind our algorithm is that if there are voters i and j with $I_i \subseteq I_j$, then it is at least as profitable (for purposes of making a the Condorcet winner) to remove voter j as to remove voter i . Now note that the intervals I_i of voters in N_1 are nested by set inclusion, and similarly for voters in N_2 . Thus, we let N_1^{-r} and N_2^{-s} denote the subsets of N_1 and N_2 obtained by deleting, respectively, the r and s largest (with respect to set inclusion) intervals I_i . Because of the nesting property, if there is a way of deleting r and s voters from N_1 and N_2 that makes a the Condorcet winner, then the deletions giving N_1^{-r} and N_2^{-s} also make a the Condorcet winner.

These observations suggest the following simple algorithm: For every pair (r, s) with $0 \leq r \leq |N_1|$ and $0 \leq s \leq |N_2|$, we check whether a is the Condorcet winner in $N_1^{-r} \cup N_2^{-s}$. We return a pair (r^*, s^*) with $r^* + s^*$ minimum for which this is the case. Then the Young score of a is $r^* + s^*$. If no such pair exists, the Young score of a is infinite.

To see that this algorithm can be run in $O(mn^2)$ time, we show how to check in $O(m)$ time whether a is the Condorcet winner in $N_1^{-r} \cup N_2^{-s}$. To do so, we precompute for every $x \in A \setminus \{c\}$, $0 \leq r \leq |N_1|$, and $0 \leq s \leq |N_2|$ the numbers

$$\begin{aligned} d_r^1(x) &= |\{i \in N_1^{-r} : a \succ_i x\}| - |\{i \in N_1^{-r} : x \succ_i a\}|, \\ d_s^2(x) &= |\{i \in N_2^{-s} : a \succ_i x\}| - |\{i \in N_2^{-s} : x \succ_i a\}|. \end{aligned}$$

Note that a is a Condorcet winner in $N_1^{-r} \cup N_2^{-s}$ if and only if for all $x \in A \setminus \{c\}$ it holds that $d_r^1(x) + d_s^2(x) > 0$. The quantities $d_r^1(x)$ and $d_s^2(x)$ can be precomputed in $O(mn^2)$ time. Verifying whether $d_r^1(x) + d_s^2(x) > 0$ requires constant time and hence $O(m)$ time for every $x \in A \setminus \{c\}$. \square

7 Multi-Winner Elections

Much recent work has focussed on voting rules that select a committee $W \subseteq A$ of candidates, where $|W| = k$ has some desired size k . A particularly influential rule for this task is due to Chamberlin and Courant (1983). According to their rule, each voter i is represented by their favourite (highest-ranked) alternative in W , say $c_i \in W$. Voter i then derives her Borda score (i.e., position counting from the bottom) of c_i in her ranking as her utility. The Chamberlin–Courant rule selects a committee of size k that maximises total utility. By replacing Borda scores by other scoring vectors, we obtain a whole family of rules. Finding a winning committee under the Chamberlin–Courant rule is known to be NP-hard for Borda scores (Lu and Boutilier 2011); however, this problem can be solved by dynamic programming in polynomial time if the input profile is single-peaked (Betzler, Slinko, and Uhlmann 2013). An alternative method for single-peaked preferences proceeds via integer programming with totally unimodular constraints (Peters 2016). This method can be extended to also work for SPOC using a technique described in a useful survey by Dom (2009, Sec 4.1.4).

Theorem 9. *For SPOC profiles, a winning committee for the Chamberlin–Courant rule can be found in polynomial time for any scoring vector.*

Proof. Let P be a SPOC preference profile, and k the target committee size. Our algorithm adapts the approach of Peters (2016), where the following integer-programming formulation for calculating an optimal Chamberlin–Courant committee is introduced:

$$\begin{aligned}
 & \text{maximise } \sum_{i \in N} \sum_{r \in [m]} w_r \cdot x_{i,r} & (\text{CC-IP}) \\
 & \text{subject to } \sum_{c \in C} y_c = k & (S) \\
 & x_{i,r} \leq \sum_{c : \text{rank}_i(c) \geq r} y_c & \text{for } i \in N, r \in [m] \quad (P_{i,r}) \\
 & x_{i,r} \in \{0, 1\} & \text{for } i \in N, r \in [m] \\
 & y_c \in \{0, 1\} & \text{for } c \in C
 \end{aligned}$$

Here, (S) and $(P_{i,r})$ are labels for the constraints. In the IP, the binary variables y_c encode the k committee members, and in optimum we have $x_{i,r} = 1$ if and only if there is a committee member c who is ranked in position r or higher by voter i . The coefficients w_r in the objective function determine the scoring vector used; choosing $w_r = 1$ for all r yields Borda scores.

Using the algorithms from Section 3, find a circle C such that P is single-peaked on C , and take some $\triangleleft \in C$ arbitrarily. We now split constraints of form $(P_{i,r})$ into two types. Set $T_{i,r} = \{c : \text{rank}_i(c) \geq r\}$ to be the r th top-initial segment of i . Define

$$\begin{aligned}
 R_1 &= \{(P_{i,r}) : T_{i,r} \text{ is an interval of } \triangleleft\}, \\
 R_2 &= \{(P_{i,r}) : C \setminus T_{i,r} \text{ is an interval of } \triangleleft\} \setminus R_1.
 \end{aligned}$$

Since P is single-peaked on C , every $T_{i,r}$ is an interval of C . Thus, either $T_{i,r}$ or $C \setminus T_{i,r}$ is an interval of \triangleleft . It follows that

R_1 and R_2 partition the set of all the constraints of type $(P_{i,r})$. We now modify (CC-IP) by subtracting the constraint (S) (in the form $0 = \sum_{c \in C} y_c - k$) from each constraint $(P_{i,r}) \in R_2$. In this way, we obtain the following constraints:

$$\begin{aligned}
 x_{i,r} &\leq \sum_{c \in T_{i,r}} y_c & \text{for } (P_{i,r}) \in R_1 & \quad (P_{i,r}^*) \\
 x_{i,r} &\leq -\sum_{c \notin T_{i,r}} y_c + k & \text{for } (P_{i,r}) \in R_2 & \quad (P_{i,r}^*)
 \end{aligned}$$

We write (CC-IP*) for the program obtained by replacing $(P_{i,r})$ by $(P_{i,r}^*)$ in (CC-IP). It is easy to see that a solution vector (\mathbf{x}, \mathbf{y}) is feasible in (CC-IP) if and only if it is feasible in (CC-IP*), since both systems include the constraint $\sum_{c \in C} y_c = k$. Hence the two formulations are equivalent.

We now claim that the constraint matrix of (CC-IP*) is totally unimodular, and thus this integer program can be solved in polynomial time by solving its LP relaxation (Conforti et al. 2014, p. 183). To see this, consider first the constraint matrix $A_{\mathbf{y}}$ of only the variables $\{y_c : c \in C\}$. If we order its columns in the order given by \triangleleft , then a row corresponding to a constraint in R_1 consists of a block of consecutive 1s surrounded by 0s, and a constraint from R_2 consists of a block of consecutive -1 s surrounded by 0s. Constraint (S) is a row consisting only of 1s. Hence $A_{\mathbf{y}}$ is an ‘interval matrix’ and thus totally unimodular (Schrijver 1998, Sec 19.3–4). To obtain the full constraint matrix $A_{\mathbf{x}, \mathbf{y}}$ of (CC-IP*), we need to add columns to $A_{\mathbf{y}}$ corresponding to $x_{i,r}$ -variables. But each such variable occurs in exactly 1 constraint, and thus each column to be added contains just a single -1 , and 0s otherwise. Since totally unimodular matrices are closed under appending (positive or negative) unit columns (Schrijver 1998, Sec 19.4), we see that $A_{\mathbf{x}, \mathbf{y}}$ is also totally unimodular. \square

The other algorithms via total unimodularity discussed by Peters (2016) can be adapted using the same kind of argument; in particular this gives a polynomial-time algorithm for *egalitarian* Chamberlin–Courant, for Proportional Approval Voting (PAV), and certain OWA-based rules.

8 Discussion and Open Problems

Our results show that restricted preference domains that behave unfavourably in terms of axiomatic properties might still be very useful for algorithmic purposes. Indeed, our algorithm for Young’s rule demonstrates that it is possible to move to a larger class than single-peaked preferences while maintaining polynomial-time runtime bounds. Thus, our findings can be seen as a challenge to established algorithmic results based on restricted preferences: to which degree can their application domain be extended without resorting to super-polynomial algorithms? A specific open problem: what is the complexity of computing Dodgson scores for SPOC profiles? This seems to be open even for single-peaked profiles.

Our definition of SPOC is not the only sensible definition. One alternative definition that would fit into the generalised notion of single-peakedness introduced by Nehring and Puppe (2007) is based on *shortest paths*: it requires that for every voter i and for every alternative x , there is a *shortest*

path between $\text{top}(i)$ and x along which i 's preferences are decreasing. (Without the word “shortest” this is equivalent to SPOC.) The impact of this alternative definition is that every voter's least-preferred alternative needs to be *antipodal* to the voter's peak; this is strictly more restrictive than SPOC. It would be interesting to see whether this smaller domain allows for a wider range of positive results than SPOC.

Another direction for future work would be to extend the SPOC concept to two (and more) dimensions – preferences single-peaked on a *sphere* – but this may be difficult since little is known even about extensions of single-peakedness on a *line* to two or more dimensions (see Sui et al. 2013).

Acknowledgements We thank the anonymous reviewers for useful suggestions and pointing out a mistake in one of our algorithms. We thank Edith Elkind for helpful discussions. This work was supported by the European Research Council (ERC) under grant number 639945 (ACCORD). Dominik Peters is supported by EPSRC.

References

- Alon, N.; Feldman, M.; Procaccia, A. D.; and Tennenholtz, M. 2010a. Strategyproof approximation of the minimax on networks. *Mathematics of Operations Research* 35(3):513–526.
- Alon, N.; Feldman, M.; Procaccia, A. D.; and Tennenholtz, M. 2010b. Walking in circles. *Discrete Mathematics* 310(23):3432–3435.
- Ballester, M. A., and Haeringer, G. 2011. A characterization of the single-peaked domain. *Social Choice and Welfare* 36(2):305–322.
- Bartholdi III, J., and Trick, M. A. 1986. Stable matching with preferences derived from a psychological model. *Operation Research Letters* 5(4):165–169.
- Bartholdi III, J.; Tovey, C. A.; and Trick, M. A. 1989. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare* 6(2):157–165.
- Betzler, N.; Slinko, A.; and Uhlmann, J. 2013. On the computation of fully proportional representation. *Journal of Artificial Intelligence Research* 47(1):475–519.
- Black, D. 1948. On the rationale of group decision-making. *The Journal of Political Economy* 56(1):23–34.
- Booth, K. S., and Lueker, G. S. 1976. 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.
- Brandt, F.; Brill, M.; Hemaspaandra, E.; and Hemaspaandra, L. A. 2015. Bypassing combinatorial protections: Polynomial-time algorithms for single-peaked electorates. *Journal of Artificial Intelligence Research* 53:439–496.
- Brandt, F.; Brill, M.; and Harrenstein, P. 2016. Tournament solutions. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A., eds., *Handbook of Computational Social Choice*. Cambridge University Press.
- Brandt, F.; Geist, C.; and Peters, D. 2016. Optimal bounds for the no-show paradox via SAT solving. In *AAMAS '16*, 314–322.
- Bredereck, R.; Chen, J.; and Woeginger, G. J. 2013. A characterization of the single-crossing domain. *Social Choice and Welfare* 41(4):989–998.
- Chamberlin, B., and Courant, P. 1983. Representative deliberations and representative decisions: Proportional representation and the Borda rule. *American Political Science Review* 77(3):718–733.
- Chen, J.; Pruhs, K.; and Woeginger, G. J. 2015. The one-dimensional Euclidean domain: Finitely many obstructions are not enough. *arXiv preprint arXiv:1506.03838*.
- Conforti, M.; Cornuéjols, G.; and Zambelli, G. 2014. *Integer programming*, volume 271. Springer.
- Cornaz, D.; Galand, L.; and Spanjaard, O. 2012. Bounded single-peaked width and proportional representation. In *ECAI'12*, 270–275.
- Demange, G. 1982. Single-peaked orders on a tree. *Mathematical Social Sciences* 3(4):389–396.
- Dom, M. 2009. Algorithmic aspects of the consecutive-ones property. *Bulletin of the European Association for Theoretical Computer Science* 98:27–59.
- Elkind, E.; Lackner, M.; and Peters, D. 2016. Preference restrictions in computational social choice: Recent progress. In *IJCAI '16*, 4062–4065.
- Fischer, F.; Hudry, O.; and Niedermeier, R. 2016. Weighted tournament solutions. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A., eds., *Handbook of Computational Social Choice*. Cambridge University Press.
- Hemaspaandra, E.; Spakowski, H.; and Vogel, J. 2005. The complexity of Kemeny elections. *Theoretical Computer Science* 349(3):382–391.
- Kim, K. H., and Roush, F. W. 1980. Special domains and nonmanipulability. *Mathematical Social Sciences* 1(1):85–92.
- Lackner, M. 2014. Incomplete preferences in single-peaked electorates. In *AAAI '14*, 742–748.
- Lu, T., and Boutilier, C. 2011. Budgeted social choice: From consensus to personalized decision making. In *IJCAI'11*, 280–286.
- McGarvey, D. C. 1953. A theorem on the construction of voting paradoxes. *Econometrica* 21(4):608–610.
- Moulin, H. 1988. Condorcet's principle implies the no show paradox. *Journal of Economic Theory* 45(1):53–64.
- Moulin, H. 1991. *Axioms of Cooperative Decision Making*. Cambridge University Press.
- Nehring, K., and Puppe, C. 2007. The structure of strategy-proof social choice, Part I: General characterization and possibility results on median spaces. *Journal of Economic Theory* 135(1):269–305.
- Peters, D., and Elkind, E. 2016. Preferences single-peaked on nice trees. In *AAAI '16*, 594–600.
- Peters, D. 2016. Single-peakedness and total unimodularity: Efficiently solve voting problems without even trying. *arXiv preprint arXiv:1609.03537*.
- Peters, D. 2017. Recognising multidimensional Euclidean preferences. In *AAAI '17*. arXiv preprint arXiv:1602.08109.
- Rothe, J.; Spakowski, H.; and Vogel, J. 2003. Exact complexity of the winner problem for Young elections. *Theory of Computing Systems* 36(4):375–386.
- Sato, S. 2010. Circular domains. *Review of Economic Design* 14(3–4):331–342.
- Schrijver, A. 1998. *Theory of linear and integer programming*. John Wiley & Sons.
- Schummer, J., and Vohra, R. V. 2002. Strategy-proof location on a network. *Journal of Economic Theory* 104(2):405–428.
- Sui, X.; Francois-Nienaber, A.; and Boutilier, C. 2013. Multi-dimensional single-peaked consistency and its approximations. In *IJCAI '13*, 375–382.
- Yu, L.; Chan, H.; and Elkind, E. 2013. Multiwinner elections under preferences that are single-peaked on a tree. In *IJCAI'13*, 425–431.