On Recognising Nearly Single-Crossing Preferences

Florian Jaeckle, Dominik Peters, Edith Elkind

Department of Computer Science
University of Oxford, UK
florian.jaeckle@eng.ox.ac.uk, {dominik.peters, edith.elkind}@cs.ox.ac.uk

Abstract

If voters' preferences are one-dimensional, many hard problems in computational social choice become tractable. A preference profile can be classified as one-dimensional if it has the single-crossing property, which requires that the voters can be ordered from left to right so that their preferences are consistent with this order. In practice, preferences may exhibit some one-dimensional structure, despite not being single-crossing in the formal sense. Hence, we ask whether one can identify preference profiles that are close to being single-crossing. We consider three distance measures, which are based on partitioning voters or candidates or performing a small number of swaps in each vote. We prove that it can be efficiently decided if voters can be split into two single-crossing groups. Also, for every fixed $k \geqslant 1$ we can decide in polynomial time if a profile can be made single-crossing by performing at most k candidate swaps per vote. In contrast, for each $k \geqslant 3$ it is NP-complete to decide whether candidates can be partitioned into k sets so that the restriction of the input profile to each set is single-crossing.

1 Introduction

There is a growing body of literature in computational social choice that aims to identify notions of structure in collective preferences, and to provide efficient algorithms for recognising structured preferences; see a recent survey by Elkind, Lackner, and Peters (2017). By discovering structural properties of a preference profile we obtain an explanatory model of voters' behavior. Often such models allow us to solve preference aggregation and elicitation problems more efficiently. This is particularly important in the context of *multiwinner voting*: while multiwinner voting rules are useful in a variety of applications ranging from group recommendation systems and design of marketing campaigns to parliamentary elections, and have attracted intense interest in recent years (see Faliszewski et al. 2017), for many appealing multiwinner rules the winner determination problem is computationally hard (Lu and Boutilier 2011; Aziz et al. 2015; Skowron, Faliszewski, and Lang 2016). Thus, it is important to identify natural subdomains where winning sets can be computed efficiently. The approach based on structured preferences has been very successful in this setting, with a number

distance metric	single-peaked	single-crossing
voter deletion	NP-complete [a,b]	in P [b]
alternative deletion	in P [a]	NP-complete [b]
voter partition	NP-complete [a]	**in P** for $k = 2$
alternative partition	open	**NP-complete**
local swaps	NP-complete [a]	**in XP** wrt k

Table 1: Overview of the complexity of deciding whether a profile is close to being structured. Bold results are new.
[a] proved by Erdélyi, Lackner, and Pfandler (2017).
[b] proved by Bredereck, Chen, and Woeginger (2016).

of polynomial-time algorithms proposed for several rules and various notions of structure (Betzler, Slinko, and Uhlmann 2013; Skowron et al. 2015; Yu, Chan, and Elkind 2013; Clearwater, Puppe, and Slinko 2015; Peters and Elkind 2016; Peters and Lackner 2017).

The two best-known types of structured preferences are *single-peaked preferences* (Black 1948; Arrow 1951) and *single-crossing preferences* (Inada 1964; Mirrlees 1971; Roberts 1977). Informally, the former notion is based on ordering the candidates along a one-dimensional axis, while the latter notion is based on ordering the voters from left to right; in both cases, the voters' preferences are required to be consistent with this ordering. In particular, single-crossing preferences arise when voters' opinions depend only on a single one-dimensional type or parameter. For example, in political elections, this parameter may be the position of the voter on the left-to-right ideological spectrum (if candidates, too, can be positioned on this spectrum, the voters' preferences will be single-peaked, in addition to being single-crossing, but single-crossing preferences may arise even if this is not the case). Similarly, when voters choose among several policies, the policy space may be complicated and not obviously one-dimensional, but the voter's preferences could still be explainable by a single parameter (such as their position on how to deal with a trade-off). The economics literature considers many such scenarios, including taxation, coalition formation, public goods, and education markets (see Saporiti 2009).

Profiles that are single-peaked or single-crossing have many attractive properties: for instance, for both domains it

holds that the majority preferences are transitive, and there is a strategyproof preference aggregation rule. Also, the membership in either of these domains can be efficiently tested, and both domains admit polynomial-time algorithms for problems that are hard for general preferences (see Elkind, Lackner, and Peters 2017).

However, these notions of structure are not robust: even the addition of a single 'misbehaving' voter to a structured profile can make it non-structured. Further, they are restrictive in that only an exponentially small fraction of all preference profiles belongs to one of these domains (Lackner and Lackner 2017). Indeed, most preference profiles elicited in the real world are unlikely to belong to either of these domains: none of the profiles in the PrefLib library (Mattei and Walsh 2013) is single-peaked or single-crossing.

Nevertheless, voters' preferences are often driven by considerations that are essentially single-dimensional (consider, e.g., voting over tax rates), and in such cases we expect the preferences to be 'nearly' structured. For instance, it may be the case that a profile can be made single-peaked or single-crossing by swapping a few pairs of adjacent candidates in each vote, ignoring a small number of 'irrational' voters or 'unconventional' candidates, or splitting the voters in a few groups so that within each group the preferences have the desired structure. Such 'nearly structured' preferences may inherit some of the attractive computational properties of the original domain. Nearly structured preferences have been explored in a number of recent papers (Faliszewski, Hemaspaandra, and Hemaspaandra 2014; Cornaz, Galand, and Spanjaard 2012; 2013; Bredereck, Chen, and Woeginger 2016; Erdélyi, Lackner, and Pfandler 2017), which introduced several distance measures capturing how 'close' a given profile is to being structured, studied the complexity of computing such distances, and proposed algorithms for social choice problems that are hard for general preferences for instances from 'nearly structured' domains.

In particular, these papers argue that, for nearly structured domains to be algorithmically useful, they need to admit efficient recognition algorithms. To address this challenge, Erdélyi, Lackner, and Pfandler (2017) put forward an extensive list of distance measures, and show that, for almost all distances on their list, it is NP-hard to compute how far a given profile is from being single-peaked. In contrast, for the single-crossing domain the picture is far from complete: Bredereck, Chen, and Woeginger (2016) provide a hardness result for candidate deletion and a polynomial-time algorithm for voter deletion, and Cornaz, Galand, and Spanjaard (2012; 2013) provide positive algorithmic results for a measure they call the *single-peaked/single-crossing width*, but for many distance measures introduced by Erdélyi, Lackner, and Pfandler (2017) the complexity of deciding if a given profile is close to being single-crossing remains open.

Our contribution. We study profiles that are almost single-crossing in three different senses:

• *Voter partition*. We ask whether the input profile can be partitioned into k single-crossing subprofiles, i.e., whether the electorate can be subdivided into few sub-communities each of which is well-structured. We solve this problem

for $k = 2$ by reducing it to 2SAT.

• *Local swaps*. We ask whether the input profile can be made single-crossing by making at most k swaps of adjacent alternatives in the preferences of each voter, i.e., whether it can be the case that the input profile fails to be single-crossing simply because voters made small 'errors' when reporting their preferences. We show that this problem is in XP with respect to k, i.e., for each fixed k we can answer this question in polynomial time.

• *Alternative partition*. We ask whether the set of alternatives can be partitioned into k subsets $A = A_1 \cup \cdots \cup A_k$ so that the restriction of the input profile to each subset is single-crossing. We show that this problem is NP-hard for every fixed $k \geqslant 3$, by providing a reduction from the k-colouring problem.

Our results are summarised in Table 1.

2 Preliminaries

Let A be a finite set of m *alternatives*, or *candidates*, and let $N = \{1, \ldots, n\}$ be a set of n voters; each voter i is associated with a strict total order v_i over A, which we call v_i's *preference order*. If alternative a is ranked above alternative b in a strict total order v, we write $a >_v b$. The collection $P = (v_1, \ldots, v_n)$ is called a *preference profile*. Given a subset A' of A, we write $P[A']$ to denote a profile $(u_1, \ldots, u_n)$ of strict linear orders over A' such that for each $i \in [n]$ and all $a, b \in A'$ we have $a >_{u_i} b$ if and only if $a >_{v_i} b$; we refer to $P[A']$ as the *restriction of P to A'*. We write $N_{ab} = \{i \in N : a >_{v_i} b\}$ to denote the set of voters who prefer a to b.

Definition. *Let L be a linear order over the set N of voters. We say that P is* single-crossing with respect to L *if for all pairs $a, b \in A$ of alternatives, the sets N_{ab} and N_{ba} are intervals of L.*

Thus, as we move along L from left to right, we observe that the voters' preferences over a and b 'cross' at most once. We say that a profile P is *single-crossing* if there exists some linear order L over N such that P is single-crossing with respect to L. We also say that $P = (v_1, \ldots, v_n)$ is *single-crossing in the given order* if P is single-crossing with respect to the order L given by $1 < 2 < \cdots < n$. Figure 1 shows an example of a single-crossing profile.

v_1	v_2	v_3	v_4	v_5
a	b	b	d	d
b	a	d	b	c
c	d	a	c	b
d	c	c	a	a

Figure 1: A preference profile that is single-crossing with respect to the voter ordering $1 < 2 < 3 < 4 < 5$. The 'trajectories' of any two alternatives cross at most once.

We now introduce an alternative way of looking at single-crossing preferences, and present a method to decide whether

a given profile is single-crossing. Given two linear orders u and v, we define the *conflict set* $\Delta(u, v)$ to be the set of pairs of alternatives on which u and v disagree:

$$\Delta(u, v) = \{\{a, b\} \subseteq A : a >_u b \text{ and } b >_v a\}.$$

The *Kendall-tau distance* between u and v is then defined as $d(u, v) = |\Delta(u, v)|$.

Consider a profile $P = (v_1, \ldots, v_n)$ that is single-crossing in the given order. Then, as i grows from 1 to n, the number of conflicts between v_1 and v_i increases with each crossing. Thus, we have

$$\varnothing = \Delta(v_1, v_1) \subseteq \Delta(v_1, v_2) \subseteq \cdots \subseteq \Delta(v_1, v_n). \quad (1)$$

Clearly, the converse is also true: If (1) holds, then P is single-crossing in the given order. For example, in Figure 1, we have

$$\Delta(v_1, v_2) = \{ab, cd\},$$
$$\Delta(v_1, v_3) = \{ab, cd, ad\},$$
$$\Delta(v_1, v_4) = \{ab, cd, ad, bd, ac\},$$
$$\Delta(v_1, v_5) = \{ab, cd, ad, bd, ac, bc\},$$

and these sets form a chain with respect to $\subseteq$.

Given a profile P and a voter i, we say that P is *single-crossing with first voter i* if P is single-crossing with respect to an ordering L in which i is the first (leftmost) voter. Then (1) implies the following characterisation.

Proposition 1. *A profile $P = (v_1, \ldots, v_n)$ is single-crossing with first voter i if and only if for all pairs of voters j, k either $\Delta(v_i, v_j) \subseteq \Delta(v_i, v_k)$ or $\Delta(v_i, v_k) \subseteq \Delta(v_i, v_j)$.*

The condition formulated in Proposition 1 is easy to check. Thus, fixing the first voter substantially simplifies the task of finding an ordering L such that P is single-crossing with respect to L. In particular, using this characterisation, we can decide in polynomial time whether a given profile P is single-crossing.

Proposition 2. *We can decide whether a given profile is single-crossing in time $O(n^3m^2)$.*

Proof. Guess the first voter i, compute the sets $\Delta(v_i, v_j)$, and check if they form a chain with respect to $\subseteq$. $\square$

This basic algorithm can be optimised in various ways, leading to faster runtimes. For instance, one can work with the Kendall-tau distances $d(v_i, v_j)$ rather than the conflict sets (see, e.g., Doignon and Falmagne 1994; Elkind, Faliszewski, and Slinko 2012). Our algorithms for detecting nearly single-crossing profiles, however, build on this simple template based on conflict sets.

Another corollary to Proposition 1 is that, in order for a profile P to be single-crossing with first voter i, the distances $d(v_i, v_j)$, $v_j \in P$, must be pairwise distinct.

3 Voter Partition

Arguably, if a preference profile is single-crossing, then there is a form of *consensus* among the voters: while individuals may have very different preferences, there is a common understanding of what the underlying issue space is, and how

different positions in this issue space translate into preference rankings (List et al. 2012 make a similar argument in the context of single-peaked preferences). However, it may also happen that, while there is no common agreement over the issue space, voters can be split into a few disjoint groups so that each group shares the same perspective on issues and, consequently, each group has single-crossing preferences.

This is the underlying idea of the *voter partition* metric, which will be considered in this section. Formally, we ask if we can partition N into k subsets, $N = N^1 \cup N^2 \cup \cdots \cup N^k$, so that for each $j \in [k]$ the subprofile $P^j = (v_i)_{i \in N^j}$ is single-crossing. Certainly, the answer is 'yes' if k is large enough: e.g., every two-voter profile is single-crossing, so for $k \geqslant |N|/2$ the answer is positive. But given our original motivation, we are particularly interested in solving this problem for small values of k. An example of a profile where the answer is 'yes' for $k = 2$ is shown in Figure 2.

We note that voter partition is related to *voter deletion*, where we look for a maximum-cardinality subset N' of N such that the profile $P' = (v_i)_{i \in N'}$ is single-crossing; in contrast, in the voter partition problem we want both P' and $P \setminus P'$ to be single-crossing.

	v_1	v_2	v_3	v_4	v_5	v_6
	a	a	d	b	d	a
	b	c	c	c	a	d
	c	d	b	d	b	c
	d	b	a	a	c	b

Figure 2: A preference profile that can be partitioned into two single crossing profiles, $\{v_1, v_2, v_3\}$ and $\{v_4, v_5, v_6\}$.

The main result of this section is a polynomial-time algorithm for the voter partition problem with $k = 2$. Our algorithm proceeds by creating and solving several instances of 2SAT, and uses the characterisation of single-crossing preferences with a fixed first voter from Proposition 1.

Theorem 3. *Given a profile P with n voters and m alternatives, we can decide in time $O(n^4m^2)$ whether the voter set N can be partitioned as $N = N^1 \cup N^2$ so that both profiles $P^1 = (v_i)_{i \in N^1}$ and $P^2 = (v_i)_{i \in N^2}$ are single-crossing.*

Proof. As a first step, we guess two voters $i, \ell \in N$, $i \neq \ell$, and let $s^1 = v_i$, $s^2 = v_\ell$. We then check whether there is a partition $P = P^1 \cup P^2$ such that for $j = 1, 2$ it holds that $s^j \in N^j$ and P^j is single-crossing with first vote s^j. We perform this check by a reduction to 2SAT, the problem of deciding if a formula of propositional logic in 2-CNF is satisfiable. 2SAT is known to be solvable in linear time (see, e.g., Papadimitriou 1993).

To construct this formula, we introduce one variable z_v for each vote $v \in P$. Intuitively, an assignment α to these variables corresponds to a partition of P into subprofiles

$$P^1 = \{v : \alpha(z_v) = \text{true}\} \text{ and } P^2 = \{v : \alpha(z_v) = \text{false}\}. \quad (2)$$

We now introduce clauses that force each of the profiles P^j, $j = 1, 2$, to be single-crossing with first vote s^j. Given two

sets A and B, we write $A \parallel B$ if A and B are incomparable under $\subseteq$, i.e., $A \not\subseteq B$ and $B \not\subseteq A$. For each $j = 1, 2$ and every pair of votes $x, y \in P^j$ the sets $\Delta(s^j, x)$ and $\Delta(s^j, y)$ must be comparable under $\subseteq$, i.e., $\Delta(s^j, x) \not\parallel \Delta(s^j, y)$. That is, if $\Delta(s^j, x) \parallel \Delta(s^j, y)$ for some $j = 1, 2$, then votes x and y cannot both belong to P^j. Hence our condition is captured by the following 2-CNF formula:

$$\varphi^{s^1, s^2} \equiv \bigwedge_{\substack{x, y \in P \\ \Delta(s^1, x) \parallel \Delta(s^1, y)}} (\neg z_x \vee \neg z_y) \wedge \bigwedge_{\substack{x, y \in N \\ \Delta(s^2, x) \parallel \Delta(s^2, y)}} (z_x \vee z_y). \quad (3)$$

If formula (3) is not satisfiable, there is no partition of P into two single-crossing profiles with s^1 and s^2 as first votes. Conversely, a satisfying assignment α can be converted into a partition via (2).

Formula (3) has n variables and $O(n^2)$ clauses, and can be constructed in $O(n^2 m^2)$ time. Since 2SAT can be solved in linear time, we can decide whether (3) can be satisfied in time $O(n^2)$. As we consider up to $\binom{n}{2}$ different formulas, the overall runtime of this algorithm is $O(n^4 m^2)$. $\qquad \square$

The proof of Theorem 3 does not extend to $k = 3$, since the reduction to SAT would yield clauses with three literals. We leave the complexity of voter partition with $k \geqslant 3$ open. This problem seems structurally similar to the problem of 3-colouring a graph formed by a union of three incomparability graphs. However, existing results on incomparability graphs (see, e.g., Bosek, Krawczyk, and Matecki 2013) do not seem to be directly applicable to our problem.

We note that for the case of single-peaked preferences, voter partition is known to be NP-hard for each $k \geqslant 3$, but the case $k = 2$ is open (Erdélyi, Lackner, and Pfandler 2017).

4 Local Swaps

We now consider the *local swaps* distance, where we have a budget of k swaps per vote: for each voter, we may (successively) perform k swaps of two candidates that are adjacent in the vote, with the aim of making the profile single-crossing. A swap of adjacent candidates is, in some sense, a minimal change of a preference order, so this distance corresponds to fine-grained perturbations of the input profile.

Figure 3 shows an example of a profile that can be made single-crossing by performing a single swap per voter. Note that the original version of the profile on the left looks rather chaotic, whereas the perturbed version on the right is visibly structured. If we are in a setting where we expect the true profile to be single-crossing, but we observe the profile in the left part of Figure 3, we may well come to the conclusion that there have been errors in the process of eliciting the voters' preferences, and that the perturbed profile is closer to the truth.

As in the previous section, our aim is to decide whether we can make a profile single-crossing by using at most k swaps per vote. Note that even the case $k = 1$ is not trivial: a naive brute-force search requires considering $(m - 1)^n$ possibilities. However, we show that for $k = 1$ this problem can be solved in polynomial time. In fact, we obtain a stronger result: our problem can be solved in polynomial time for every fixed k, though the exponent in the runtime may depend on k. This
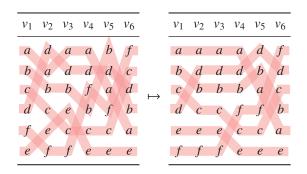

Figure 3: A profile that becomes single-crossing after we swap a single pair of adjacent candidates in each vote.

places our problem in the parameterised complexity class XP with respect to k.

Our algorithm proceeds by splitting the profile into several pieces. Then for each piece we enumerate all ways in which that piece can be made single-crossing by using at most k swaps per voter. Finally, we use dynamic programming to find a way to combine the results for individual pieces.

Theorem 4. *Given a profile P with n voters and m alternatives, we can decide in time $O((nm)^{8k^3 + 2k^2} \text{poly}(n, m))$ whether P can be made single-crossing by making at most k swaps in the preferences of each voter.*

Proof. Note that if P contains multiple copies of some preference order, we can remove all but one copy without changing the answer to our question. Thus, in what follows we can treat P as a *set* of (distinct) preference orders. Recall that the Kendall-tau distance $d(v_i, v_j)$ between two preference orders is the number of pairs on which they disagree. Given a set $P = \{v_1, \ldots, v_q\}$ of preference orders, we say that a set $P' = \{u_1, \ldots, u_r\}$ is a *k-variant* of P if for each $i \in [q]$ there is a $j \in [r]$ such that $d(v_i, u_j) \leqslant k$ and for each $j \in [r]$ there is an $i \in [q]$ such that $d(v_i, u_j) \leqslant k$; note that we may have $q \neq r$. Thus, a profile is obtainable from P by performing k local swaps if and only if it is a k-variant. We can now formalise our question as follows: is there a k-variant P' of P that is single-crossing?

Throughout our analysis of the problem, we will fix the number k of allowed swaps per vote. Thus, 'polynomial size' and 'polynomial time' refer to values that are polynomial for constant k, that is, to values bounded by $(nm)^{f(k)}$ for some function f.

Let P be the input profile, interpreted as a set of linear orders. The first step in our algorithm is to guess which voter will appear first in the order of voters witnessing that the target profile P' is single-crossing. We implement this guessing by iterating through all preference orders $v \in P$, and iterating through all possibilities of applying at most k swaps to v, yielding s. There are $O(nm^k)$ possibilities for this. Having guessed the starting voter, we have to check whether P is within k local swaps from being single-crossing with first voter s.

We now partition P into blocks based on the distance to s.

For each $r = 1, \ldots, \binom{m}{2}$, let

$$P_r := \{v_i \in P : d(s, v_i) = r\}$$

be the set of preference orders at distance exactly r from s. At first we will handle each P_r separately, by enumerating all promising ways of applying k swaps to the orders in P_r. To this end, let $B_k(P_r)$ denote the set of all k-variants P'_r of P_r such that $P'_r \cup \{s\}$ is single-crossing with first vote s.

Lemma 5. *The set $B_k(P_r)$ is of polynomial size and can be enumerated in polynomial time.*

Proof. Consider a set $P'_r \in B_k(P_r)$. We have $d(s, v) \in [r - k, r + k]$ for each $v \in P'_r$. For P'_r to be single-crossing with first vote s, the votes in P'_r must be at pairwise different distances from s. Thus, for each $d \in [r - k, r + k]$, there is at most one preference order v in P'_r with $d(s, v) = d$. Hence $|P'_r| \leqslant 2k + 1$.

It follows that P'_r is a k-variant of some subset $P^*_r \subseteq P_r$ of size at most $2k + 1$. Thus, to enumerate $B_k(P_r)$, we can enumerate all $O(|P_r|^{2k+1})$ subsets $P^*_r \subseteq P_r$ of size at most $2k + 1$, then, for each P^*_r, enumerate all $O((m^k)^{2k+1})$ k-variants of P^*_r, and finally cross out all sets that are not k-variants of P_r or are not single-crossing with first voter s. □

The procedure described in Lemma 5 is rather inefficient even for relatively small k. For the case $k = 1$, we can show that $B_k(P_r)$ can be enumerated in time *linear* in m. We omit the proof due to space constraints.

Having enumerated all potential k-variants for each P_r separately, we now need a way to decide whether there is a way to combine them into a single-crossing profile. Note that in this profile votes from different sets P'_r may be interleaved, which makes our problem more difficult. Nevertheless, we will now argue that it can be reduced to the problem of finding a path of a certain length in the directed acyclic graph that we will now construct.

Let $P_{r_1}, \ldots, P_{r_q}$ be a list of all non-empty sets P_r, where $r_1 < \cdots < r_q$. For each $i = 1, \ldots, q - 4k + 1$, let

$$C_{r_i} := \{(P'_{r_i}, \ldots, P'_{r_{i+4k-1}}) \in B_k(P_{r_i}) \times \cdots \times B_k(P_{r_{i+4k-1}}) :$$
$$\{s\} \cup P'_{r_i} \cup \cdots \cup P'_{r_{i+4k-1}} \text{ is single-crossing with first vote } s\}$$

Thus, C_{r_i} is the set of all combinations of k-variants from $4k$ successive non-empty sets P_r that together are single-crossing with first vote s. We now construct a directed graph D with vertex set $V = C_{r_1} \cup \cdots \cup C_{r_{q-4k+1}}$ and add an arc

$$(P'_{r_i}, \ldots, P'_{r_{i+4k-1}}) \to (P''_{r_{i+1}}, \ldots, P''_{r_{i+4k}})$$

whenever these two vectors are *compatible*, that is, whenever $P'_{r_{i+t}} = P''_{r_{i+t}}$ for all $t \in [4k - 1]$. Note that there are only arcs between vertices from two successive sets C_{r_i}. To finish the proof, we need the following lemma.

Lemma 6. *Profile P admits a k-variant that is single-crossing with first vote s if and only if D contains a path of length $q - 4k + 1$.*

Proof. Suppose P' is such a k-variant. Then, for each r_i, there is some $P'_{r_i} \in B_k(P_{r_i})$ such that $P'_{r_i} \subseteq P'$ by the correctness of the algorithm in Lemma 5. Then D contains the path

$$(P'_{r_1}, \ldots, P'_{r_{4k}}) \to \cdots \to (P'_{r_{q-4k+1}}, \ldots, P'_{r_q}).$$

Conversely, suppose D contains a path of length $q - 4k + 1$. This induces a choice of $P'_{r_i} \in B_k(P_{r_i})$ for each $i \in [q]$. Let $P' = \{s\} \cup P'_{r_1} \cup \cdots \cup P'_{r_q}$. We claim that P' is single-crossing with first vote s. We check this using Proposition 1.

Suppose, aiming for a contradiction, that there exist preference orders $x \in P_{r_i}$ and $y \in P_{r_j}$ with $i \leqslant j$ such that $\Delta(s, x)$ and $\Delta(s, y)$ are incomparable under $\subseteq$. Choose x and y so as to minimize $j - i$. Clearly $j - i \geqslant 4k$ by construction of C_{r_i}. Take any preference order $z \in P'_{r_{i+2k}}$. Then by minimality of $j - i$, the sets $\Delta(s, x)$ and $\Delta(s, z)$ must be comparable, and the sets $\Delta(s, z)$ and $\Delta(s, y)$ must be comparable (under $\subseteq$). Since we are only allowed k swaps per vote, we have

$$d(s, x) \leqslant r_i + k = (r_i + 2k) - k \leqslant d(s, z), \text{ and}$$
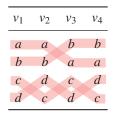$$d(s, z) \leqslant r_i + 3k \leqslant r_j - k \leqslant d(s, y).$$

Hence, we must have $\Delta(s, x) \subseteq \Delta(s, z)$ and $\Delta(s, z) \subseteq \Delta(s, y)$. Thus $\Delta(s, x) \subseteq \Delta(s, y)$, which is a contradiction. □

Now, the digraph D is acyclic, since it admits a topological ordering; also, for every fixed k it has polynomially many vertices. Thus, we can decide in polynomial time whether D contains a directed path with $q - 4k + 1$ vertices. By Lemma 6, such paths correspond to k-variants of P that are single-crossing with first vote s. Thus, the proof is complete. □

5 Alternative Partition

The *alternative partition* metric is similar in spirit to the voter partition metric: given a profile P over a set of alternatives A, we ask whether we can partition A into k subsets $A = A_1 \cup \cdots \cup A_k$ so that for each $j \in [k]$ the restriction of P to A_j is single-crossing. Importantly, each restricted profile $P[A_j]$ may be single-crossing with respect to a different voter ordering. We will now argue that this problem is NP-complete.

In our hardness reduction, it will be useful to have some small profiles that are *not* single-crossing. For example, consider the profile on the right. Suppose this profile was single-crossing with respect to some order $<_L$. By considering the pair $\{a, b\}$, we conclude that $\{1, 2\} <_L \{3, 4\}$ or $\{3, 4\} <_L \{1, 2\}$. Similarly, by considering the pair $\{c, d\}$, we conclude that either $\{1, 3\} <_L \{2, 4\}$ or $\{2, 4\} <_L \{1, 3\}$. The conditions imposed by $\{a, b\}$ are incompatible with those imposed by $\{c, d\}$, so this profile is not single-crossing. We can use this profile as a gadget in our reduction, to ensure that alternatives a, b, c, d cannot all be contained in the same part of our partition.

Profiles of this type are known as *δ-configurations*, which together with 'γ-configurations' are used by Bredereck, Chen, and Woeginger (2013) to characterise the domain of single-crossing profiles in terms of forbidden subprofiles. Given a

profile P, we will say that a subset $A^\dagger \subseteq A$ of alternatives *induces a δ-configuration* if there are four voters i_1, i_2, i_3, i_4 and we can write $A^\dagger = \{a, b, c, d\}$ (where a, b, c, d are not necessarily distinct) so that

$$a \succ_{i_1} b \text{ and } \quad c \succ_{i_1} d, \quad a \succ_{i_2} b \text{ and } \quad d \succ_{i_2} c,$$
$$b \succ_{i_3} a \text{ and } \quad c \succ_{i_3} d, \quad b \succ_{i_4} a \text{ and } \quad d \succ_{i_4} c.$$

As argued above, when $A^\dagger$ induces a δ-configuration, and $A' \subseteq A$ is such that $P[A']$ is single-crossing, it holds that $A^\dagger \not\subseteq A'$, i.e., the alternatives in $A^\dagger$ cannot all be in the same part of the partition.

Our reduction is from the k-colouring problem. An instance of this problem is an undirected graph $G = (V, E)$ and a parameter k; it is a 'yes'-instance if it is possible to partition V into k independent sets, and a 'no'-instance otherwise. The k-colouring problem is NP-complete for each $k \geqslant 3$ (Garey and Johnson 1979). The high-level idea of our reduction is that the vertices of G become alternatives, and we use voters and further dummy alternatives to build δ-configurations that force alternatives corresponding to adjacent vertices to lie in different partition parts. The implementation details are rather involved. Our reduction is inspired by the proof that it is NP-hard to decide if a given profile can be made single-crossing by removing k alternatives (Bredereck, Chen, and Woeginger 2016).

Theorem 7. *For every $k \geqslant 3$, it is* NP-*complete to decide whether given a profile P we can partition the alternative set A into k sets $A = A_1 \cup \cdots \cup A_k$ so that for each $j \in [k]$ the profile $P[A_j]$ is single-crossing.*

Proof. This problem is in NP, since one can check in polynomial time whether k sets form a partition and whether a given profile is single-crossing. To show NP-hardness, we reduce from the k-colouring problem.

Given an instance G of the k-colouring problem, where G has vertex set $F = \{f_1, \ldots, f_n\}$ and edge set $E = \{e_1, \ldots, e_m\}$, we will construct a profile P over a set of alternatives A in such a way that A can be partitioned into k sets $A = A_1 \cup \cdots \cup A_k$ with each profile $P[A_j]$ being single-crossing if and only if G is k-colourable.

For every vertex $f_i \in F$ we introduce a candidate q_i. Let
$$Q := \{q_i : i \in [n]\}.$$
Define the canonical order of Q to be
$$\langle Q \rangle := q_1 \succ \cdots \succ q_n.$$
For every edge $e_i \in E$ we introduce $4k$ candidates $a_i^1, \ldots, a_i^{2k}, x_i^1, \ldots, x_i^{2k}$, and define
$$B_i := \{a_i^1, \ldots, a_i^{2k}\}, \quad X_i := \{x_i^1, \ldots, x_i^{2k}\},$$
with the canonical orders of B_i and X_i being defined as
$$\langle B_i \rangle := a_i^1 \succ \cdots \succ a_i^{2k}, \quad \langle X_i \rangle := x_i^1 \succ \cdots \succ x_i^{2k}.$$
Furthermore, for $j \in [2k]$ we let $\ell = 2k - j$ and define the following linear orders over the sets B_i and X_i:

$$\widehat{\langle B_i \rangle}_j := a_i^1 \succ a_i^2 \succ \cdots \succ a_i^\ell \succ a_i^{2k} \succ a_i^{2k-1} \succ \cdots \succ a_i^{\ell+1},$$
$$\widehat{\langle X_i \rangle}_j := x_i^1 \succ x_i^2 \succ \cdots \succ x_i^\ell \succ x_i^{2k} \succ x_i^{2k-1} \succ \cdots \succ x_i^{\ell+1},$$
$$\widetilde{\langle B_i \rangle}_j := a_i^j \succ a_i^{j-1} \succ \cdots \succ a_i^1 \succ a_i^{j+1} \succ a_i^{j+2} \succ \cdots \succ a_i^{2k},$$
$$\widetilde{\langle X_i \rangle}_j := x_i^j \succ x_i^{j-1} \succ \cdots \succ x_i^1 \succ x_i^{j+1} \succ x_i^{j+2} \succ \cdots \succ x_i^{2k}.$$

Define the candidate set A to be
$$A := Q \cup B_1 \cup X_1 \cup \cdots \cup B_m \cup X_m.$$
Now for every $j \in [2k]$ we construct two votes v_j, v_j' that rank the candidates as follows:

$$v_j: \quad \langle Q \rangle \succ \widehat{\langle B_1 \rangle}_j \succ \widehat{\langle X_1 \rangle}_j \succ \cdots \succ \widehat{\langle B_m \rangle}_j \succ \widehat{\langle X_m \rangle}_j,$$
$$v_j': \quad \langle Q \rangle \succ \widetilde{\langle B_1 \rangle}_j \succ \widetilde{\langle X_1 \rangle}_j \succ \cdots \succ \widetilde{\langle B_m \rangle}_j \succ \widetilde{\langle X_m \rangle}_j.$$

For each $i \in [m]$ it holds that any three distinct candidates a_i^r, a_i^s, a_i^t in B_i such that $r < s < t$, together with votes $v_1, v_{2k-r}, v_s', v_{2k}$, induce a δ-configuration, since

$$v_1: \quad a_i^r > a_i^s > a_i^t,$$
$$v_{2k-r}: \quad a_i^r > a_i^t > a_i^s,$$
$$v_s': \quad a_i^s > a_i^r > a_i^t,$$
$$v_{2k}: \quad a_i^t > a_i^s > a_i^r.$$

Similarly, any three distinct candidates x_i^r, x_i^s, x_i^t in X_i with $r < s < t$, together with votes $v_1, v_{2k-r}, v_s', v_{2k}$, induce a δ-configuration.

Now, for every edge $e_i = \{f_j, f_k\}$ with $j < k$ we construct four votes $w_{4i-3}, w_{4i-2}, w_{4i-1}, w_{4i}$ that rank the candidates as follows:

$$\langle B_1 \rangle \succ \langle X_1 \rangle \succ \cdots \succ \langle B_{i-1} \rangle \succ \langle X_{i-1} \rangle \succ$$
$$q_1 \succ \cdots \succ q_{j-1} \succ \{q_j\} \cup B_i \succ q_{j+1} \succ \cdots \succ$$
$$q_{k-1} \succ \{q_k\} \cup X_i \succ q_{k+1} \succ \cdots \succ q_n \succ$$
$$\langle B_{i+1} \rangle \succ \langle X_{i+1} \rangle \succ \cdots \succ \langle B_m \rangle \succ \langle X_m \rangle.$$

We still need to define the relative ranking of q_j and B_i, as well as of q_k and X_i:

$$w_{4i-3} \text{ ranks } q_j \succ \langle B_i \rangle \text{ and } q_k \succ \langle X_i \rangle,$$
$$w_{4i-2} \text{ ranks } q_j \succ \langle B_i \rangle \text{ and } \langle X_i \rangle \succ q_k,$$
$$w_{4i-1} \text{ ranks } \langle B_i \rangle \succ q_j \text{ and } q_k \succ \langle X_i \rangle,$$
$$w_{4i} \text{ ranks } \langle B_i \rangle \succ q_j \text{ and } \langle X_i \rangle \succ q_k.$$

Note that for all candidates $s \in B_i$ and $t \in X_i$ it holds that votes $w_{4i-3}, \ldots, w_{4i}$ and candidates q_j, q_k, s, t induce a δ-configuration.

Define the profile P to be
$$\{v_i, v_i' : i \in [2k]\} \cup \{w_{4i-3}, w_{4i-2}, w_{4i-1}, w_{4i} : i \in [m]\}.$$

The profile P can be constructed in polynomial time. We will now show that G is k-colourable if and only if there exists a partition $A_1 \cup \cdots \cup A_k$ of A such that $P[A_i]$ is single-crossing for each $i \in [k]$.

$\impliedby$ Suppose $A_1, \ldots, A_k$ is a partition of A such that $P[A_1], \ldots, P[A_k]$ are all single-crossing. For each $i \in [k]$, define
$$F_i := \{f_j \in F : q_j \in A_i\}.$$
We show that $F_1, \ldots, F_k$ is a valid k-colouring for G.

Clearly, $F_1, \ldots, F_k$ partition F; hence it suffices to show that each F_i is an independent set. Suppose, aiming for a contradiction, that one of the sets F_i, say F_1, is not an independent set. Then there exists an edge $e_\ell = (f_i, f_j) \in E$

such that $f_i, f_j \in F_1$. By definition of F_1 both q_i and q_j are in A_1.

Recall that any three distinct candidates $a_\ell^r, a_\ell^s, a_\ell^t$ in B_ℓ, where $r < s < t$, together with votes $v_1, v_{2k-r}, v'_s, v_{2k}$ induce a δ-configuration. Since $P[A_1], \ldots, P[A_k]$ are all single-crossing, it follows that $|B_\ell \cap A_t| \leqslant 2$ for all $t \in [k]$. Because $|B_\ell| = 2k$, it follows that $|B_\ell \cap A_t| = 2$ for all $t \in [k]$; similarly, $|X_\ell \cap A_t| = 2$ for all $t \in [k]$. Now let $a \in B_\ell \cap A_1$ and $x \in X_\ell \cap A_1$. Then votes $w_{4\ell-3}, \ldots, w_{4\ell}$ and candidates a, x, q_i, q_j induce a δ-configuration, which contradicts the assumption that $P[A_1]$ is single-crossing.

Thus F_1, and similarly all other sets F_i, are independent sets. Hence $F_1, \ldots, F_k$ is a valid colouring of G.

$\Longrightarrow$ Now suppose that G is k-colourable. Let $F_1, \ldots, F_k$ be a k-colouring of G. We construct a partition of $A = A_1 \cup \cdots \cup A_k$, where for each $i \in [k]$ we take

$$A_i := \{q_j \in A : f_j \in F_i\}$$
$$\cup \{a_j^{2i-1}, a_j^{2i}, x_j^{2i-1}, x_j^{2i} \in A : j \in [m]\}.$$

The sets $A_1, \ldots, A_k$ form a partition of A, and thus we only need to show that $P[A_i]$ is single-crossing for each $i \in [k]$. To do so, we define an order of votes in each profile $P[A_i]$, $i \in [k]$. We set

$$L_i := v_{2(k-i+1)} > \cdots > v_{2k} > v'_{2i} > \cdots > v'_{2k} >$$
$$v_1 > \cdots > v_{2k-2i+1} > v'_1 > \cdots > v'_{2i-1} >$$
$$\{w_1, w_2, w_3, w_4\} > \cdots > \{w_{4m-3}, w_{4m-2}, w_{4m-1}, w_{4m}\}.$$

It remains to define L_i on the sets $\{w_{4j-3}, w_{4j-2}, w_{4j-1}, w_{4j}\}$ for $j \in [m]$. Suppose that $e_j = (f_k, f_\ell)$. Then if $f_\ell \notin F_i$, then

$$L_i : \quad w_{4j-3} > w_{4j-2} > w_{4j-1} > w_{4j},$$

and if $f_\ell \in F_i$, then

$$L_i : \quad w_{4j-3} > w_{4j-1} > w_{4j-2} > w_{4j}.$$

Note that for all $j \in [m]$ in all votes up to v'_{2k} in the linear order L_i candidate a_j^{2i} is ranked above a_j^{2i-1} and candidate x_j^{2i} is ranked above x_j^{2i-1}, whereas in all votes from v_1 onwards candidate a_j^{2i-1} is ranked above a_j^{2i} and candidate x_j^{2i-1} is ranked above x_j^{2i} for all $j \in [m]$.

It is straightforward to check that $P[A_i]$ is single-crossing with respect to L_i by arguing that for all distinct candidates $s, t \in A_i$ the set of votes where s is ranked above t is connected in L_i, and so is the set of votes where t is ranked above s. We omit this argument due to space constraints; details can be found in the full version of the paper. $\square$

6 Conclusions and Future Work

This work contributes to our understanding of nearly structured preferences. For two natural notions of distance we have obtained efficient algorithms for identifying preferences that are very nearly single-crossing. We consider our algorithmic result for the local swaps distance to be particularly appealing, because it concerns an operation that only performs small perturbations to voters' reported preferences. Indeed, we believe that when preferences are a few local swaps away from

being structured, we can sometimes be justified in taking the structured version of the profile to be the starting point of the analysis. It is instructive to contrast the local swap distance and the voter deletion distance: while for single-crossing preferences the latter admits an efficient algorithm even when k is part of the input (Bredereck, Chen, and Woeginger 2016), the voter deletion distance forces us to completely ignore some voters' preferences; in this sense, the local swaps distance is more egalitarian.

The running time of our algorithm for local swaps, whose exponent depends on k, will only be acceptable for small k (and moderate values of n and m). However, this is the practically relevant case: we are presumably not interested in profiles that can be made structured by very large perturbations, since then the structure does not explain much anyway. It would be interesting to obtain other positive results of this type, i.e., to design algorithms that can efficiently detect if a given preference profile is *very* close to being structured.

In contrast to our tractability results for voter partition and local swaps, we obtain a hardness result for alternative partition. In a sense, this result is not surprising, as for the related measures of voter and alternative *deletion*, a similar phenomenon is known: the smallest number of voters that need to be deleted to make a profile single-crossing can be computed in polynomial time, whereas for alternative deletion this problem is NP-hard (Bredereck, Chen, and Woeginger 2016). Intuitively, this is because the definition of single-crossing preferences focuses on ordering the voters, and therefore computational problems that deal with rearranging the voters are easier than those that deal with rearranging the alternatives. Indeed, for single-peaked preferences, which are defined in terms of an ordering of the alternatives, the alternative deletion problem is easy, but the voter deletion problem is NP-hard (Erdélyi, Lackner, and Pfandler 2017).

Our work leaves a number of interesting open questions. In particular, for the local swaps distance, it would be desirable to improve the running time of our algorithm, and to investigate the existence of an FPT algorithm. For the voter partition algorithm, the complexity of partitioning into $k > 2$ sets remains open. It is also natural to ask if there are analogues of our results for the single-peaked domain. We note that existing NP-completeness results for this domain (Erdélyi, Lackner, and Pfandler 2017), do not rule out tractability of voter partition for $k = 2$, or an XP result for local swaps.

Another research direction is to investigate the complexity of deciding whether a given profile is close to being structured, for more general notions of structure, such as being single-peaked or single-crossing on graphs other than the line (such as, e.g., trees or cycles). To the best of our knowledge, this topic has not yet been explored.

Finally, to make our positive results more practically applicable, it would be desirable to extend tractability results for winner determination of various preference aggregation rules (particularly multiwinner voting rules) from the perfectly structured profiles to those that are almost structured. Such results have been obtained for single-peaked or single-crossing width (Cornaz, Galand, and Spanjaard 2012; 2013; Skowron et al. 2015), and, very recently, for a few other distance measures (Misra, Sonar, and Vaidyanathan 2017).

Results of this type also appear in the literature on manipulation and control (see, e.g., Faliszewski, Hemaspaandra, and Hemaspaandra 2014; Yang and Guo 2015; 2017). Extending this line of research to a wider range of domains and distance measures is a promising direction for future work.

References

Arrow, K. J. 1951. *Social Choice and Individual Values*. John Wiley and Sons.

Aziz, H.; Gaspers, S.; Gudmundsson, J.; Mackenzie, S.; Mattei, N.; and Walsh, T. 2015. Computational aspects of multi-winner approval voting. In *AAMAS'15*, 107–115.

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. *Journal of Political Economy* 56(1):23–34.

Bosek, B.; Krawczyk, T.; and Matecki, G. 2013. First-fit coloring of incomparability graphs. *SIAM Journal on Discrete Mathematics* 27(1):126–140.

Brederick, R.; Chen, J.; and Woeginger, G. J. 2013. A characterization of the single-crossing domain. *Social Choice and Welfare* 41(4):989–998.

Brederick, R.; Chen, J.; and Woeginger, G. J. 2016. Are there any nicely structured preference profiles nearby? *Mathematical Social Sciences* 79:61–73.

Clearwater, A.; Puppe, C.; and Slinko, A. 2015. Generalizing the single-crossing property on lines and trees to intermediate preferences on median graphs. In *IJCAI'15*, 32–38.

Cornaz, D.; Galand, L.; and Spanjaard, O. 2012. Bounded single-peaked width and proportional representation. In *ECAI'12*, 270–275.

Cornaz, D.; Galand, L.; and Spanjaard, O. 2013. Kemeny elections with bounded single-peaked or single-crossing width. In *IJCAI'13*, 76–82.

Doignon, J.-P., and Falmagne, J.-C. 1994. A polynomial time algorithm for unidimensional unfolding representations. *Journal of Algorithms* 16(2):218–233.

Elkind, E.; Faliszewski, P.; and Slinko, A. 2012. Clone structures in voters' preferences. In *ACM EC'12*, 496–513.

Elkind, E.; Lackner, M.; and Peters, D. 2017. Structured preferences. In Endriss, U., ed., *Trends in Computational Social Choice*. AI Access. chapter 10, 187–207.

Erdélyi, G.; Lackner, M.; and Pfandler, A. 2017. Computational aspects of nearly single-peaked electorates. *Journal of Artificial Intelligence Research* 58:297–337.

Faliszewski, P.; Skowron, P.; Slinko, A.; and Talmon, N. 2017. Multiwinner voting: A new challenge for social choice theory. In Endriss, U., ed., *Trends in Computational Social Choice*. AI Access. chapter 2, 27–47.

Faliszewski, P.; Hemaspaandra, E.; and Hemaspaandra, L. 2014. The complexity of manipulative attacks in nearly single-peaked electorates. *Artificial Intelligence* 207:69–99.

Garey, M. R., and Johnson, D. S. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman.

Inada, K. 1964. A note on the simple majority decision rule. *Econometrica* 32(4):525–531.

Lackner, M.-L., and Lackner, M. 2017. On the likelihood of single-peaked preferences. *Social Choice and Welfare* 48(4):717–745.

List, C.; Luskin, R. C.; Fishkin, J. S.; and McLean, I. 2012. Deliberation, single-peakedness, and the possibility of meaningful democracy: Evidence from deliberative polls. *Journal of Politics* 75(1):80–95.

Lu, T., and Boutilier, C. 2011. Budgeted social choice: From consensus to personalized decision making. In *IJCAI'11*, 280–286.

Mattei, N., and Walsh, T. 2013. Preflib: A library for preferences. In *ADT'13*, 259–270. http://www.preflib.org.

Mirrlees, J. A. 1971. An exploration in the theory of optimum income taxation. *Review of Economic Studies* 38(2):175–208.

Misra, N.; Sonar, C.; and Vaidyanathan, P. R. 2017. On the complexity of Chamberlin–Courant on almost structured profiles. In *ADT'17*, 124–138.

Papadimitriou, C. H. 1993. *Computational complexity*. Pearson.

Peters, D., and Elkind, E. 2016. Preferences single-peaked on nice trees. In *AAAI '16*, 594–600.

Peters, D., and Lackner, M. 2017. Preferences single-peaked on a circle. In *AAAI'17*, 649–655.

Roberts, K. W. 1977. Voting over income tax schedules. *Journal of Public Economics* 8(3):329–340.

Saporiti, A. 2009. Strategy-proofness and single-crossing. *Theoretical Economics* 4(2):127–163.

Skowron, P.; Yu, L.; Faliszewski, P.; and Elkind, E. 2015. The complexity of fully proportional representation for single-crossing electorates. *Theoretical Computer Science* 569:43–57.

Skowron, P.; Faliszewski, P.; and Lang, J. 2016. Finding a collective set of items: From proportional multirepresentation to group recommendation. *Artificial Intelligence* 241:191–216.

Yang, Y., and Guo, J. 2015. How hard is control in multi-peaked elections: A parameterized study. In *AAMAS'15*, 1729–1730.

Yang, Y., and Guo, J. 2017. The control complexity of r-approval: from the single-peaked case to the general case. *Journal of Computer and System Sciences* 89:432–449.

Yu, L.; Chan, H.; and Elkind, E. 2013. Multiwinner elections under preferences that are single-peaked on a tree. In *IJCAI'13*, 425–431.