Computational Aspects of Nearly Single-Peaked Electorates[*]

Gábor Erdélyi
erdelyi@wiwi.uni-siegen.de
University of Siegen
Siegen, Germany

Martin Lackner
lackner@dbai.tuwien.ac.at
Vienna University of Technology
Vienna, Austria

Andreas Pfandler
pfandler@dbai.tuwien.ac.at
Vienna University of Technology
Vienna, Austria

Abstract

Manipulation, bribery, and control are well-studied ways of changing the outcome of an election. Many voting systems are, in the general case, computationally resistant to some of these manipulative actions. However when restricted to single-peaked electorates, these systems suddenly become easy to manipulate. Recently, Faliszewski, Hemaspaandra, and Hemaspaandra (2011b) studied the complexity of dishonest behavior in nearly single-peaked electorates. These are electorates that are not single-peaked but close to it according to some distance measure.

In this paper we introduce several new distance measures regarding single-peakedness. We prove that determining whether a given profile is nearly single-peaked is NP-complete in many cases. For one case we present a polynomial-time algorithm. Furthermore, we explore the relations between several notions of nearly single-peakedness.

Introduction

Voting is a very useful method for preference aggregation and collective decision-making. It has applications in many settings ranging from politics to artificial intelligence and further topics in computer science (see, e.g., the work of Dwork et al. (2001), Ephrati and Rosenschein (1997), Ghosh et al. (1999)). In the presence of huge data volumes, the computational properties of voting rules are worth studying. In particular, we usually want to determine the winners of an election quickly. On the other hand we want to make various forms of dishonest behavior computationally hard.

The first to study the computational aspects of manipulation in elections were Bartholdi, Tovey, and Trick (1989). In this paper they defined and studied manipulation, i.e., a group of voters casting their votes insincerely in order to reach a desired outcome. Another type of dishonest behavior is control, where an external agent makes structural changes on the election such as adding/deleting/partitioning either candidates or voters (as has been studied, e.g., by Bartholdi,

Tovey, and Trick (1992)) in order to reach a desired outcome. There is also bribery, where an external agent changes some voters' votes in order to change the outcome of the election (see, e.g., the work of Faliszewski, Hemaspaandra, and Hemaspaandra (2009)). For an overview and many natural examples on bribery, control, and manipulation we refer to the work Baumeister et al. (2010), Faliszewski, Hemaspaandra, and Hemaspaandra (2010), Faliszewski and Procaccia (2010), and Brandt, Conitzer, and Endriss (2012).

Traditionally, the complexity of such attacks is studied under the assumption that, in each election, any admissible vote can occur. However, there are many elections where the diversity of the votes is limited in the sense that there are admissible votes nobody would ever cast. One of the best known examples is *single-peakedness*, introduced by Black (1948). It assumes that the votes are polarized along some linear axis. The study of the computational aspects of elections with single-peaked preferences was initiated by Walsh (2007) (see also the work of Faliszewski et al. (2011a), and Brandt et al. (2010)). Many problems which are NP-hard in the general case turn out to be easy for single-peaked societies. A recent line of research initiated by Conitzer (2009) and by Escoffier, Lang, and Öztürk (2008) suggests that many elections are not perfectly single-peaked but are *close* to it with respect to some metric. In the work of Faliszewski, Hemaspaandra, and Hemaspaandra (2011b) various notions of nearly single-peaked elections were introduced and it was shown that the complexity of manipulative actions jumps back to NP-hardness in many cases.

In this paper we present a systematic study of nearly single-peaked electorates. Our main contributions are:

- We introduce three new notions of nearly single-peakedness. In addition, we study four notions that already have been defined or suggested in the literature.

- We explore connections between both existing and new notions by providing inequalities. These allow to compare these notions and better understand their relationship.

- We analyze the computational complexity of computing the distance of arbitrary preference profiles to single-peakedness. In most cases we show NP-completeness. For the *k*-candidate deletion distance, we present a polynomial-time algorithm.

Related Work Our paper fits in the line of research on single-peaked and nearly single-peaked preferences. Ballester and Haeringer (2011) combinatorially characterize single-peaked elections by the means of forbidden subsequences. In the work of Brandt et al. (2010) and Faliszewski et al. (2011a) the complexity of winner problems and of dishonest behavior (e.g., the complexity of manipulation and control) in electorates with single-peaked preferences is investigated. These papers do not consider nearly single-peaked preferences, but mention them as future work.

In the context of nearly single-peaked preferences the most relevant paper is by Faliszewski, Hemaspaandra, and Hemaspaandra (2011b). They introduce several notions of nearly single-peakedness and analyze the complexity of bribery, control, and manipulation in nearly single-peaked elections. In contrast, we are studying the complexity of computing the distance of a preference profile to single-peakedness. The question whether a given profile is single-peaked has been recently investigated by Escoffier, Lang, and Öztürk (2008). The difference to their work is that they have not considered nearly single-peakedness, they only pointed it out as a future research direction.

Two further distance measures have recently been studied. *Single-peaked width* has been been studied by Cornaz, Galand, and Spanjaard (2012). Elkind, Faliszewski, and Slinko (2012) define the *decloning* measure which describes the number of adjacent candidates (adjacent in every vote) that are merged into one candidate in order to obtain single-peakedness. Finally, we remark that single-peaked preferences have been considered in the context of preference elicitation (Conitzer 2009) and in the context of possible and necessary winners under uncertainty regarding the votes (Walsh 2007).

Preliminaries

Let C be a finite set of *candidates*, V be a finite set of *voters*, and let $\succ$ be a *preference relation* (i.e., a tie-free and total order) over C. Without loss of generality let $V = \{1, \ldots, n\}$. We call a candidate c the *peak* (or *top-ranked*) of a preference relation $\succ$ if $c \succ c'$ for all $c' \in C \setminus \{c\}$. Let $\mathscr{P} = (\succ_1, \ldots, \succ_n)$ be a *preference profile* (i.e., a collection of preference relations) over the candidate set C. We say that the preference relation $\succ_i$ is the *vote* of voter i. For simplicity, we will write for each voter $i \in V$ $c_1 c_2 \ldots c_m$ instead of $c_1 \succ_i c_2 \succ_i \ldots \succ_i c_m$. An *election* is defined as a triple $E = (C, V, \mathscr{P})$, where C is the set of candidates, V the set of voters and $\mathscr{P}$ a preference profile over C.

Definition 1. Let an *axis* A be a total order over C denoted by $>$. Furthermore, let $\succ$ be a vote with peak c. The vote $\succ$ is *single-peaked with respect to A* if for any $x, y \in C$, if $x > y > c$ or $c > y > x$ then $c \succ y \succ x$ has to hold.

A preference profile $\mathscr{P}$ is said to be *single-peaked with respect to an axis A* if and only if each vote is single-peaked with respect to A. A preference profile $\mathscr{P}$ is said to be *single-peaked consistent* if there exists an axis A such that $\mathscr{P}$ is single-peaked with respect to A.

By $\mathscr{P}[C']$ we denote the profile $\mathscr{P}$ restricted to the candidates in C'. Analogously if A is an axis over C, we denote by $A[C']$ the axis A restricted to candidates in C'.

Nearly single-peaked preferences

In real-world settings one can expect a certain amount of "noise" in preference data. The single-peakedness property is very fragile and thus susceptible to such noise. The following example illustrates the *fragility of single-peakedness*: Consider the single-peaked election consisting two kinds of votes: $a \succ b \succ c \succ d$ and $d \succ c \succ b \succ a$. Assume that both votes have been cast by a large number of voters. This election is single-peaked only with respect to the axis $a > b > c > d$ and its reverse. Adding a single vote $a \succ b \succ d \succ c$ destroys the single-peakedness property although this vote is almost identical to the first kind of votes.

In this section we formally define different notions of nearly single-peakedness. All these notions define a distance measure to single-peaked profiles. Furthermore, we explore the relation of these distance measures.

▶ *k-Maverick* (**M**) The first formal definition of nearly single-peaked societies was given by Faliszewski, Hemaspaandra, and Hemaspaandra (2011b). Consider a preference profile $\mathscr{P}$ for which most voters are single-peaked with respect to some axis A. All voters that are not single-peaked with respect to A are called *mavericks*. The number of mavericks defines a natural distance measure to single-peakedness. If an axis can be found for a large subset of the voters, this is still a fundamental observation about the structure of the votes.

Definition 2 (Faliszewski, Hemaspaandra, and Hemaspaandra 2011b)**.** Let $E = (C, V, \mathscr{P})$ be an election and k a positive integer. We say that the profile $\mathscr{P}$ is *k-maverick single-peaked consistent* if by removing at most k preference relations (votes) from $\mathscr{P}$ one can obtain a preference profile $\mathscr{P}'$ that is single-peaked consistent.

Let $M(\mathscr{P})$ denote the smallest k such that $\mathscr{P}$ is k-maverick single-peaked consistent. Note that $M(\mathscr{P}) \leq |V| - 1$ always holds.

Example 1. Consider an election with $C = \{a, b, c, d, e\}$ and $V = \{1, 2, \ldots, 202\}$. Let the preference profile $\mathscr{P}$ consist of the votes $a \succ_1 b \succ_1 c \succ_1 e \succ_1 d$ and $e \succ_2 d \succ_2 c \succ_2 a \succ_2 b$ as well as 100 votes of the form $a \succ b \succ c \succ d \succ e$ and 100 votes of the form $e \succ d \succ c \succ b \succ a$. Notice that preference profiles containing $a \succ b \succ c \succ d \succ e$ and $e \succ d \succ c \succ b \succ a$ may only be single-peaked consistent with respect to the axis $a > b > c > d > e$ and its reverse. Since $\succ_1$ and $\succ_2$ are not single-peaked with respect to this axis, $\mathscr{P}$ is not single-peaked. Deleting $\succ_1$ and $\succ_2$ obviously yields single-peaked consistency and thus we have $M(\mathscr{P}) = 2$.

▶ *k-Candidate Deletion* (**CD**) As suggested by Escoffier, Lang, and Öztürk (2008), we introduce outlier candidates. These are candidates that do not have "a correct place" on any axis and consequently have to be deleted in order to obtain a single-peaked consistent profile. Examples could be a candidate that is not well-known (e.g., a new political party) or a candidate that prioritizes other topics than most candidates and thereby is judged by the voters according to different criteria. The votes restricted to the remaining candidates

might still have a clear and significant structure, in particular they might be single-peaked consistent.

Definition 3. Let $E = (C, V, \mathscr{P})$ be an election and k a positive integer. We say that the profile $\mathscr{P}$ is *k-candidate deletion single-peaked consistent* if we can obtain a set $C' \subseteq C$ by removing at most k candidates from C such that the preference profile $\mathscr{P}[C']$ is single-peaked consistent.

Let $CD(\mathscr{P})$ denote the smallest k such that $\mathscr{P}$ is k-candidate deletion single-peaked consistent. Note that $CD(\mathscr{P}) \leq |C| - 2$ always holds.

Example 1 (continued). Consider the preference profile $\mathscr{P}$ as defined above. Observe that for $C' = \{b, c, d\}$, $\mathscr{P}[C']$ is single-peaked consistent. Deleting a single candidate does not yield single-peaked consistency and thus $CD(\mathscr{P}) = 2$.

▶ *k-**Local Candidate Deletion** (LCD)* Personal friendships or hatreds between voters and candidates could move candidates up or down in a vote. These personal relationships cannot be reflected in a global axis. To eliminate the influence of personal relationships to some candidates we define a local version of the previous notion. This notion can also deal with the possibility that the least favorite candidates might be ranked without special consideration or even randomly.

We first have to define partial domains and partial profiles.

Definition 4. Let C be a set of candidates and A an axis over C. A preference relation $\succ$ over a candidate set $C' \subset C$ is called a *partial vote*. It is said to be *single-peaked with respect to A* if it is single-peaked with respect to $A[C']$. A partial preference profile consists of partial votes. It is called *single-peaked consistent* if there exists an axis A such that its partial votes are single-peaked with respect to A.

Definition 5. Let $E = (C, V, \mathscr{P})$ be an election and k a positive integer. We say that the profile $\mathscr{P}$ is *k-local candidate deletion single-peaked consistent* if by removing at most k candidates from each vote in $\mathscr{P}$ we obtain a partial preference profile $\mathscr{P}'$ that is single-peaked consistent.

Let $LCD(\mathscr{P})$ denote the smallest k such that $\mathscr{P}$ is k-local candidate deletion single-peaked consistent. Note that $LCD(\mathscr{P}) \leq |C| - 2$ always holds.

Example 1 (continued). Note that it is sufficient to remove a from vote $\succ_1$ and e from vote $\succ_2$ to obtain single-peaked consistency. Consequently, $LCD(\mathscr{P}) = 1$.

▶ *k-**Additional Axes** (AA)* Another suggestion by Escoffier, Lang, and Öztürk (2008) was to consider the minimum number of axes such that each preference relation of the profile is single-peaked with respect to at least one of these axes. This notion is particularly useful if each candidate represents opinions on several issues (as it is the case in political elections). A voter's ranking of the candidates would then depend on which issue is considered most important by the voter and consequently each issue might give rise to its own corresponding axis.

Definition 6. Let $E = (C, V, \mathscr{P})$ be an election and k a positive integer. We say that the profile $\mathscr{P}$ is *k-additional axes single-peaked consistent* if there is a partition $V_1, \dots, V_{k+1}$ of V such that the corresponding preference profiles $\mathscr{P}_1, \dots, \mathscr{P}_{k+1}$ are single-peaked consistent.

Let $AA(\mathscr{P})$ denote the smallest k such that $\mathscr{P}$ is k-additional axes single-peaked consistent. Note that $AA(\mathscr{P}) < \min\left(|V|, \frac{|C|!}{2}\right)$ always holds. This is because the number of distinct votes is trivially bounded by $|V|$. Furthermore, $AA(\mathscr{P})$ is bounded by $\frac{|C|!}{2}$ since at most $|C|!$ distinct votes exist and each vote and its reverse are single-peaked with respect to the same axes.

Example 1 (continued). We argue that one additional axis is required for single-peaked consistency. Notice that $\succ_1$ and $\succ_2$ are single-peaked consistent with respect to axis $b > a > c > e > d$. The remaining votes are consistent with respect to $a > b > c > d > e$. Thus, one additional axis is required and hence $AA(\mathscr{P}) = 1$.

▶ *k-**Global Swaps** (GS)* There is a second method of dealing with candidates that are "not placed correctly" according to an axis A. Instead of deleting them from either the candidate set C or from a vote, we could try to move them to the correct position. We do this by performing a sequence of swaps of consecutive candidates. We remark that the minimum number of swaps required to change one vote to another is the *Kendall tau distance* (Kendall 1938) of these two votes (permutations). For example, to get from vote $abcd$ to vote $adbc$, we first have to swap candidates c and d, and then we have to swap b and d. Since this changes the votes in a more subtle way, this can be considered a less obtrusive notion than k-(Local) Candidate Deletion.

Definition 7. Let $E = (C, V, \mathscr{P})$ be an election and k a positive integer. We say that the profile $\mathscr{P}$ is *k-global swaps single-peaked consistent* if $\mathscr{P}$ can be made single-peaked by performing at most k swaps in the profile.

Note that these swaps can be performed wherever we want – we can have k swaps in only one vote, or one swap each in k votes. Let $GS(\mathscr{P})$ denote the smallest k such that $\mathscr{P}$ is k-global swaps single-peaked consistent. Note that $GS(\mathscr{P}) \leq \binom{|C|}{2} \cdot |V|$ always holds since rearranging a total order in order to obtain any other total order requires at most $\binom{|C|}{2}$ swaps.

Example 1 (continued). It is possible to make $\mathscr{P}$ single-peaked consistent by swapping d and e in vote $\succ_1$ and swapping a and b in vote $\succ_2$. This gives $GS(\mathscr{P}) = 2$.

▶ *k-**Local Swaps** (LS)* We can also consider a "local budget" for swaps, i.e., we allow up to k swaps per vote. This distance measure has been introduced by Faliszewski, Hemaspaandra, and Hemaspaandra (2011b) as Dodgson$_k$.

Definition 8. Let $E = (C, V, \mathscr{P})$ be an election and k a positive integer. We say that the profile $\mathscr{P}$ is *k-local swaps single-peaked consistent* if $\mathscr{P}$ can be made single-peaked consistent by performing no more than k swaps per vote.

Let $LS(\mathscr{P})$ denote the smallest k such that $\mathscr{P}$ is k-local swaps single-peaked consistent. Note that $LS(\mathscr{P}) \leq \binom{|C|}{2}$ always holds.

Example 1 (continued). Since only one swap is required in $\succ_1$ and $\succ_2$ each, we have $LS(\mathscr{P}) = 1$.

▶ *k-Candidate Partition* (**CP**) Our last nearly single-peaked notion is the candidate analog of *k*-additional axes. In this case we partition the set of candidates into subsets such that all of the restricted profiles are single-peaked consistent. This notion is useful for example in the following situation. Each candidate has an opinion on a controversial Yes/No-issue. Depending on their own preference voters will always rank all Yes-candidates before or after all No-candidates. It might be that when considering only the Yes- or only the No-candidates, the election is single-peaked. Therefore, if we acknowledge the importance of this Yes/No-issue and partition the candidates accordingly, we may obtain two single-peaked elections.

Definition 9. Let $E = (C, V, \mathscr{P})$ be an election and k a positive integer. We say that the profile $\mathscr{P}$ is *k-candidate partition single-peaked consistent* if the set of candidates C can be partitioned into at most k disjoint sets $C_1, \ldots, C_k$ with $C_1 \cup \ldots \cup C_k = C$ such that the profiles $\mathscr{P}[C_1], \ldots, \mathscr{P}[C_k]$ are single-peaked consistent.

Let $CP(\mathscr{P})$ denote the smallest k such that $\mathscr{P}$ is k-candidate partition single-peaked consistent. Note that $CP(\mathscr{P}) \leq \left\lceil \frac{|C|}{2} \right\rceil$ always holds.

Example 1 (continued). We partition the candidates into $C_1 = \{a, e\}$ and $C_2 = \{b, c, d\}$. Notice that $\mathscr{P}[C_1]$ is trivially single-peaked consistent because this holds for all profiles over at most two candidates. Furthermore, $\mathscr{P}[C_2]$ contains only votes of the form $b \succ c \succ d$ or its reverse, which also gives immediately single-peakedness. Thus, $CP(\mathscr{P}) = 2$.

We start with our first result, which shows several inequalities that hold for the distance measures under consideration. We hereby show how these measures relate to each other. Notice that these inequalities do not have an immediate impact on a classical complexity analysis. However, they turn out to be very useful for the complexity analysis of manipulation in nearly single-peaked elections.

Theorem 10. *Let $\mathscr{P}$ be a preference profile. Then the following inequalities hold:*

(1) $LS(\mathscr{P}) \leq GS(\mathscr{P})$.   *(5)* $M(\mathscr{P}) \leq GS(\mathscr{P})$.
(2) $LCD(\mathscr{P}) \leq CD(\mathscr{P})$.   *(6)* $AA(\mathscr{P}) \leq M(\mathscr{P})$.
(3) $CD(\mathscr{P}) \leq GS(\mathscr{P})$.   *(7)* $CP(\mathscr{P}) \leq CD(\mathscr{P}) + 1$.
(4) $LCD(\mathscr{P}) \leq LS(\mathscr{P})$.   *(8)* $CP(\mathscr{P}) \leq LS(\mathscr{P}) + 1$.

This list is complete in the following sense: Inequalities that are not listed here and that do not follow from transitivity do not hold in general. The resulting partial order with respect to $\leq$ is displayed in Figure 1 as a Hasse diagram.

Decision Problems We now introduce the seven problems we will study. We define the following problem for X ∈ {Maverick, Candidate Deletion, Local Candidate Deletion, Additional Axes, Global Swaps, Local Swaps, Candidate Partition}.

X SINGLE-PEAKED CONSISTENCY

| **Given:** | An election $E = (C, V, \mathscr{P})$ and a positive integer k. |
| **Question:** | Is $\mathscr{P}$ k-X single-peaked consistent? |

| GS | … | Global Swaps |
M	…	Maverick
CD	…	Candidate Del.
LS	…	Local Swaps
AA	…	Additional Axes
CP	…	Candidate Partition
LCD	…	Local Cand. Del.

Figure 1: Hasse diagram of the partial order described in Theorem 10.

Nearly Single-Peaked Consistency

In this section we analyze the computational complexity of determining the distance to single-peakedness with respect to one of the measures discussed above. It turns out that in most cases this task is NP-complete with the notable exception of candidate deletion.

Theorem 11. MAVERICK SINGLE-PEAKED CONSISTENCY *is* NP-*complete.*

Theorem 12. LOCAL CANDIDATE DELETION SINGLE-PEAKED CONSISTENCY *is* NP-*complete.*

Theorem 13. ADDITIONAL AXES SINGLE-PEAKED CONSISTENCY *is* NP-*complete. This holds even for $k = 2$, i.e., for checking single-peaked consistency with two additional axes.*

Theorem 14. GLOBAL SWAPS SINGLE-PEAKED CONSISTENCY *is* NP-*complete, even for eight voters.*

Theorem 15. LOCAL SWAPS SINGLE-PEAKED CONSISTENCY *is* NP-*complete.*

The proofs of these theorems had to be omitted due to space constraints. Exemplarily, we give the proof of Theorem 12.

Proof of Theorem 12. We will reduce from the NP-complete problem MINIMUM RADIUS, which was shown to be NP-complete in (Frances and Litman 1997) and is defined as follows:

MINIMUM RADIUS

| **Given:** | A set of strings $S \subseteq \{0,1\}^n$ and a positive integer s. |
| **Question:** | Has S a radius of at most s, i.e., is there a string $\alpha \in \{0,1\}^n$ such that each string in S has a Hamming distance of at most s to α? |

Given a string β, let $\beta(k)$ denote the bit value at the k-th position in β. We are going to construct an LCD SINGLE-PEAKED CONSISTENCY instance. Each string in $S = \{\beta_1, \ldots, \beta_n\}$ will correspond to a voter. Each bit of the strings corresponds to two candidates. In addition, we have $2ms + 2$ extra candidates. Consequently, we have $C = \{c_1^1, c_1^2, c_2^1, c_2^2, \ldots, c_n^1, c_n^2, c_1', \ldots, c_{ms+1}', c_1'', \ldots, c_{ms+1}''\}$.

We define the preference profile with the help of two functions creating total orders: $f_0(a,b) = a \succ b$ and $f_1(a,b) =$

$b \succ a$. The vote $\succ_k$, for each $k \in \{1, \ldots, m\}$, is of the form

$$c'_1 \ \ldots \ c'_{ms+1} \ f_{\beta_k(1)}(c_1^1, c_1^2) \ \ldots \ f_{\beta_k(n)}(c_n^1, c_n^2) \ c''_1 \ \ldots \ c''_{ms+1}.$$

Let $\overline{\succ_i}$ denote vote $\succ_i$ in reverse order. The preference profile $\mathscr{P}$ is now defined as $(\succ_1, \ldots, \succ_n, \overline{\succ_1}, \ldots, \overline{\succ_n})$. It holds that $(C, V, \mathscr{P})$ is s-LCD single-peaked consistent if and only if S has a radius of at most s. Due to lack of space we have to omit the correctness proof of the reduction. □

In contrast to the previous hardness results, we show that CANDIDATE DELETION SINGLE-PEAKED CONSISTENCY can be decided in polynomial time. The algorithm builds upon the $\mathscr{O}(|V| \cdot |C|)$ time algorithm for testing single-peaked consistency by Escoffier, Lang, and Öztürk (2008). For the remainder of this section let $(C, V, \mathscr{P})$ be an election with $|V| = n$ and $C = \{c_1, \ldots, c_m\}$.

Definition 16. $L(\mathscr{P}, C')$ is the set of last ranked candidates in $\mathscr{P}[C']$.

Definition 17. A *partial axis* A is a total order of a subset of the candidates in C. Let $cand(A)$ denote the candidates that are ordered by A. Consequently, any partial axis A is an axis over $cand(A)$.

Definition 18. An *incomplete axis* is a partial axis with a marked position that indicates where further elements may be added. We denote this position by a star symbol, e.g., the incomplete axis $c_1 > c_2 > \star > c_3$ allows additional candidates to be added right of c_2 and left of c_3. The *boundary* of an incomplete axis A, $boundary(A)$, are the two elements left and right of the star, e.g., $boundary(c_1 > c_2 > \star > c_3) = \{c_2, c_3\}$.

The algorithm by Escoffier, Lang, and Öztürk (2008) proceeds iteratively by placing the last ranked candidates that have not yet been placed. Let C' be the set of candidates that have not yet been positioned on the (incomplete) axis A. The algorithm checks what kinds of constraints follow from each vote. If these constraints do not contradict each other, the set of last ranked candidates $L(\mathscr{P}, C')$ is placed. We denote this procedure with $\mathsf{place}(A, X)$ where $X = L(\mathscr{P}, C')$. The procedure $\mathsf{place}(A, X)$ returns either a new incomplete axis (extending A by the candidates in X) or the value INCONSISTENT. The algorithm repeatedly invokes place until all elements have been placed or a contradiction has been found.

Fact 19. The placement of candidates in the place procedure only depends on $boundary(A)$ rather than the full partial axis A.

This fact is the main reason why we can employ dynamic programming in the algorithm for deciding the CANDIDATE DELETION SINGLE-PEAKED CONSISTENCY problem.

The candidate deletion algorithm. The algorithm operates on pairs consisting of an incomplete axis A (as in the single-peaked consistency algorithm) and a set of candidates X that have been placed on A by the previous call of the place procedure. We refer to this pair (A, X) as *state*. The basic data structure is an array $\mathscr{S}$ of states. The total number of states can be exponential in C. However, for the algorithm it suffices to maintain an array of size $\lfloor 1.5 \cdot |C|^2 \rfloor$.

The algorithm utilizes dynamic programming. Given a state (A, X) and a set of candidates X_{new} that are to be placed next, we try to obtain a new incomplete axis A_{new}. If such an incomplete axis A_{new} can be found, it extends A by the candidates in X_{new}. The placement is performed by the place procedure, more precisely we call $\mathsf{place}(A, X_{new})$. Since placing more than two candidates at once is not possible, we always have $|X_{new}| \leq 2$.

In order to allow for a more concise description of the algorithm, we assume that there are two additional candidates c_0 and c'_0. The candidate c_0 is ranked last in every vote and c'_0 is ranked second-to-last, i.e., $L(\mathscr{P}, C) = \{c_0\}$ and $L(\mathscr{P}, C \setminus \{c_0\}) = \{c'_0\}$. All other candidates are ranked above these. Note that these modified votes are single-peaked consistent if and only if the original votes were single-peaked consistent. Indeed, any axis for the original votes is an axis for the modified votes if we add c_0 at the leftmost position and c'_0 at the rightmost position (or vice versa). Conversely, every axis of the modified votes is an axis of the original votes if c_0 and c'_0 are removed. Due to these observations we can assume that the algorithm always starts with the incomplete axis $c_0 > \star > c'_0$. The starting state is consequently $(c_0 > \star > c'_0, \{c'_0\})$.

For a concise description of the algorithm see Algorithm 1. Similar to the single-peaked consistency algorithm we place lower ranked candidates first. However, in contrast to the previously described single-peaked consistency algorithm we may delete candidates. Hence there are several possibilities which candidates are to be placed next. We define a set $next(X)$ containing those candidates that may be placed next. For this let $X = \{x_1, x_2\}$, with $x_1 = x_2$ in case $|X| = 1$. We define

$$next(X) = \{c \in C \mid \forall k \in \{1, \ldots, |V|\} \ (c \succ_k x_1) \vee (c \succ_k x_2)\}.$$

Candidates that are not contained in $next(X)$ have already been processed, i.e., they have already been placed on the axis or they have been deleted. Consequently, the candidates that have been deleted so far are exactly those contained in the set $C \setminus (cand(A) \cup next(X))$.

Recall that placing three or more candidates by the place procedure at once is not possible. Therefore, we consider every set of candidates $X_{new} \subseteq next(X)$ of cardinality 1 or 2. If $|X_{new}| = 2$ an additional condition has to apply. There has to be a vote $\succ$ for which $x_1 \succ x_2$ holds and another vote for which $x_2 \succ x_1$ holds. This condition is equivalent to requiring that $L(\mathscr{P}, X_{new}) = X_{new}$. (If this condition is not satisfied, the lower ranked candidate has to be placed first and the higher ranked candidate in a later iteration step.)

First, we create a copy of $\mathscr{S}$ called $\mathscr{S}_{new}$. We are only modifying $\mathscr{S}_{new}$ while iterating over all elements of $\mathscr{S}$. The algorithm applies $\mathsf{place}(A, X_{new})$ for the incomplete axis A of every state $(A, X) \in \mathscr{S}$ and for every admissible $X_{new} \subseteq next(X)$. If $\mathsf{place}(A, X_{new})$ returns a new incomplete axis A_{new}, we have obtained a new state (A_{new}, X_{new}). We now have to decide whether to store (A_{new}, X_{new}) in $\mathscr{S}_{new}$.

Recall that we keep at most $\lfloor 1.5 \cdot |C|^2 \rfloor$ states in $\mathscr{S}$. This is possible due to the following observations: If two states

Algorithm 1: Polynomial time algorithm for k-CD single-peaked consistency – Theorem 20

```
 1  (A_init, X_init) = (c_0 > ⋆ > c'_0, {c'_0})
 2  𝒮 ← {(A_init, X_init)}
 3  repeat |C| times
 4  │   𝒮_new ← 𝒮
 5  │   foreach state (A, X) ∈ 𝒮 do
 6  │   │   foreach X_new ⊆ next(X) with 1 ≤ |X_new| ≤ 2
 7  │   │   │   A_new ← place(A, X_new)     and L(𝒫, X_new) = X_new do
 8  │   │   │   if A_new ≠ INCONSISTENT then
 9  │   │   │   │   i ← index(A_new, X_new)
10  │   │   │   │   if 𝒮[i] is empty then
11  │   │   │   │   │   𝒮[i] ← (A_new, X_new)
12  │   │   │   │   else
13  │   │   │   │   │   (A_old, X_old) ← 𝒮[i]
14  │   │   │   │   │   if |cand(A_new)| > |cand(A_old)| then
15  │   │   │   │   │   │   𝒮[i] ← (A_new, X_new)
16  │   𝒮 ← 𝒮_new
17  return a state (A, X) ∈ 𝒮 with maximum |cand(A)|
```

have the same X set, they have the same set of candidates that have not yet been placed nor deleted. If two states have the same boundary, they are indistinguishable from the perspective of the place procedure (cf. Fact 19). Therefore if two states have both the same boundary and the same X set, we can discard the state where more candidates had to be deleted so far. This is the same as discarding the state with the smaller incomplete axis. In case that two such states have incomplete axes of the same cardinality, we can make the choice arbitrarily. Since there are $\binom{|C|}{2}$ possible boundaries and only three X sets per boundary, an array of size $\lfloor 1.5 \cdot |C|^2 \rfloor$ suffices.

We use the function *index* to compute the index of a state (A, X) in the array. This position is uniquely determined by the boundary of A and by X. Thus, when deciding whether a state (A, X) is to be stored in $𝒮_{new}$, it only has to be compared with the state stored in $𝒮_{new}$ at index $index(A, X)$. The state with a larger incomplete axis is stored in $𝒮_{new}$. After the iteration over all states in $𝒮$ and all possible X sets is completed, $𝒮_{new}$ becomes $𝒮$.

We repeat the just described procedure $|C|$ times. Any sequence of states leading to a cardinality maximal partial axis has length at most $|C|$ because in each step at least one candidate is placed on the axis. Therefore the algorithm stops after $|C|$ iterations and the array $𝒮$ contains a partial axis of maximum cardinality.

Theorem 20. CANDIDATE DELETION SINGLE-PEAKED CONSISTENCY *can be solved in time* $\mathcal{O}(|V| \cdot |C|^5)$.

Proof. The runtime bound can be seen as follows. The array $𝒮$ has size $\lfloor 1.5 \cdot |C|^2 \rfloor = \mathcal{O}(|C|^2)$. For each of these

Notion	Complexity
k-Maverick	NP-c (Thm. 11)
k-Candidate Deletion	in P (Thm. 20)
k-Local Candidate Deletion	NP-c (Thm. 12)
k-Additional Axes	NP-c (Thm. 13)
k-Global Swaps	NP-c (Thm. 14)
k-Local Swaps	NP-c (Thm. 15)
k-Candidate Partition	open

Table 1: Complexity results for different notions of nearly single-peakedness.

states and for each admissible X set, we employ the place procedure. Since place has a runtime of $\mathcal{O}(|V|)$, we require $\mathcal{O}(|V| \cdot |C|^4)$ time for one iteration step. This is repeated $|C|$ times. We obtain a total runtime of $\mathcal{O}(|V| \cdot |C|^5)$. □

Conclusions and Open Questions

We have investigated the nearly single-peaked consistency problem. We have introduced three new notions of nearly single-peakedness and studied four known notions. We have drawn a complete picture of the relations between all the notions of nearly single-peakedness discussed in this paper. For five notions we have shown that deciding single-peaked consistency is NP-complete and for k-candidate deletion we have presented a polynomial time algorithm. We refer the reader to Table 1 for an overview. An obvious direction for future work is to determine the complexity of CANDIDATE PARTITION SINGLE-PEAKED CONSISTENCY.

NP-completeness, however, does not rule out the possibility of algorithms that perform well in practice. One approach is to search for fixed-parameter algorithms, i.e., an algorithm with runtime $f(k) \cdot poly(n)$ for some computable function f. A fixed-parameter algorithm for MAVERICK SINGLE-PEAKED CONSISTENCY was found by Bredereck (2012). The design of fixed-parameter algorithms for nearly single-peaked consistency deserves further attention. A second approach is the development of approximation algorithms since nearly single-peaked consistency can also be seen as an optimization problem.

Another interesting direction for future work is extending our models to manipulative behavior, such as manipulation, control, and bribery. That is, assuming we have a nearly single-peaked electorate according to one of our notions, how hard is a manipulative action under a certain voting rule computationally? The analysis of manipulation and control in such elections has already been started by Faliszewski, Hemaspaandra, and Hemaspaandra (2011b) for some distance measures. This work has yet to be extended to the distance measures introduced in this paper. Finally, there might be further useful and natural distance measures regarding single-peakedness to be found.

References

Ballester, M. A., and Haeringer, G. 2011. A characterization of the single-peaked domain. *Social Choice and Welfare* 36(2):305–322.

Bartholdi, J.; Tovey, C.; and Trick, M. 1989. The computational difficulty of manipulating an election. *Social Choice and Welfare* 6(3):227–241.

Bartholdi, J.; Tovey, C.; and Trick, M. 1992. How hard is it to control an election? *Mathematical and Computer Modeling* 16(8/9):27–40.

Baumeister, D.; Erdélyi, G.; Hemaspaandra, E.; Hemaspaandra, L.; and Rothe, J. 2010. Computational aspects of approval voting. In Laslier, J., and Sanver, R., eds., *Handbook on Approval Voting*. Springer. chapter 10, 199–251.

Black, D. 1948. On the rationale of group decision making. *Journal of Political Economy* 56(1):23–34.

Brandt, F.; Brill, M.; Hemaspaandra, E.; and Hemaspaandra, L. A. 2010. Bypassing combinatorial protections: Polynomial-time algorithms for single-peaked electorates. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI-10)*, 715–722. AAAI Press.

Brandt, F.; Conitzer, V.; and Endriss, U. 2012. Computational social choice. In Weiss, G., ed., *Multiagent Systems*. MIT Press.

Bredereck, R. 2012. Personal communication.

Conitzer, V. 2009. Eliciting single-peaked preferences using comparison queries. *Journal of Artificial Intelligence Research* 35:161–191.

Cornaz, D.; Galand, L.; and Spanjaard, O. 2012. Bounded single-peaked width and proportional representation. In *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI 2012)*, volume 242 of *Frontiers in Artificial Intelligence and Applications*, 270–275. IOS Press.

Dwork, C.; Kumar, R.; Naor, M.; and Sivakumar, D. 2001. Rank aggregation methods for the web. In *Proceedings of the 10th International World Wide Web Conference*, 613–622. ACM Press.

Elkind, E.; Faliszewski, P.; and Slinko, A. M. 2012. Clone structures in voters' preferences. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC-12)*, 496–513. ACM.

Ephrati, E., and Rosenschein, J. 1997. A heuristic technique for multi-agent planning. *Annals of Mathematics and Artificial Intelligence* 20(1–4):13–67.

Escoffier, B.; Lang, J.; and Öztürk, M. 2008. Single-peaked consistency and its complexity. In *Proceedings of the 18th European Conference on Artificial Intelligence (ECAI 2008)*, volume 178 of *Frontiers in Artificial Intelligence and Applications*, 366–370. IOS Press.

Faliszewski, P., and Procaccia, A. D. 2010. AI's war on manipulation: Are we winning? *AI Magazine* 31(4):53–64.

Faliszewski, P.; Hemaspaandra, E.; Hemaspaandra, L. A.; and Rothe, J. 2011a. The shield that never was: Societies with single-peaked preferences are more open to manipulation and control. *Information and Computation* 209(2):89–107.

Faliszewski, P.; Hemaspaandra, E.; and Hemaspaandra, L. A. 2009. How hard is bribery in elections? *Journal of Artificial Intelligence Research* 35:485–532.

Faliszewski, P.; Hemaspaandra, E.; and Hemaspaandra, L. A. 2010. Using complexity to protect elections. *Communications of the ACM* 53(11):74–82.

Faliszewski, P.; Hemaspaandra, E.; and Hemaspaandra, L. A. 2011b. The complexity of manipulative attacks in nearly single-peaked electorates. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-2011)*, 228–237. Full version available as technical report arXiv:1105.5032v2 [cs.GT], ACM Computing Research Repository (CoRR), July 2012.

Frances, M., and Litman, A. 1997. On covering problems of codes. *Theory of Computing Systems* 30:113–119.

Ghosh, S.; Mundhe, M.; Hernandez, K.; and Sen, S. 1999. Voting for movies: The anatomy of recommender systems. In *Proceedings of the 3rd Annual Conference on Autonomous Agents*, 434–435. ACM Press.

Kendall, M. G. 1938. A new measure of rank correlation. *Biometrika* 30(1/2):pp. 81–93.

Walsh, T. 2007. Uncertainty in preference elicitation and aggregation. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence (AAAI-07)*, 3–8. AAAI Press.