

Incomplete Preferences in Single-Peaked Electorates

Zack Fitzsimmons

*College of the Holy Cross
Department of Mathematics and Computer Science
One College Street, Worcester, MA 01610, USA*

ZFITZSIM@HOLYCROSS.EDU

Martin Lackner

*TU Wien
Databases and Artificial Intelligence Group
Favoritenstraße 9-11, 1040 Vienna, Austria*

LACKNER@DBAI.TUWIEN.AC.AT

Abstract

Incomplete preferences are likely to arise in real-world preference aggregation scenarios. This paper deals with determining whether an incomplete preference profile is single-peaked. This is valuable information since many intractable voting problems become tractable given single-peaked preferences. We prove that the problem of recognizing single-peakedness is NP-complete for incomplete profiles consisting of partial orders. Despite this intractability result, we find several polynomial-time algorithms for reasonably restricted settings. In particular, we give polynomial-time recognition algorithms for weak orders, which can be viewed as preferences with indifference.

1. Introduction

Both human and automated decision making often have to rely on incomplete information. The same issue arises in collective decision making—voting—in multi-agent systems. Konczak and Lang (2005) distinguish two main sources of incompleteness: The first one is *intrinsic* incompleteness where a voter is unable or unwilling to give complete information. The second one is *epistemic* incompleteness where voters are able to provide complete information but at the time of decision making the required information is not fully available. Clearly, also a combination of these two scenarios may arise.

Incomplete qualitative preferences are often modeled as partial orders (see, e.g., the survey by Boutilier & Rosenschein, 2016). For example, if voters provide some pairwise comparisons between alternatives, then partial orders arise naturally as the transitive closure of these comparisons. In this sense, partial orders can be viewed as the most general model for (qualitative) incomplete preferences while maintaining transitivity as a rationality assumption. A further standard assumption in this setting (Boutilier & Rosenschein, 2016) is that voters have a true, *complete* preference order, modeled as a total order. This total order is not available (due to intrinsic or epistemic incompleteness), but the known partial order is compatible with the underlying total order, i.e., it can be extended to this total order.

Real-world preferences—whether they are incomplete or not—tend to possess an underlying structure, as they are formed by an implicit or explicit rationale. In this paper, we consider a particularly advantageous structural property: single-peaked preferences (Black, 1948). The single-peaked domain is one of the most commonly studied preference restrictions in social choice theory (Gaertner, 2002). Single-peaked preferences arise if preferences are formed based on an underlying ordering of alternatives, a so-called axis. It is assumed that voters have a most-

preferred alternative, their peak, and alternatives farther away on this axis are increasingly less preferred.

Single-peaked preference profiles enjoy many positive properties (see, e.g., surveys by Elkind, Lackner, & Peters, 2016, 2017), both for algorithmic purposes and from the axiomatic point of view from classical social choice. For example, the problem of determining if a candidate is a winner for Dodgson, Kemeny, and Young elections is computationally hard to compute (Bartholdi, Tovey, & Trick, 1989; Hemaspaandra, Hemaspaandra, & Rothe, 1997; Hemaspaandra, Spakowski, & Vogel, 2005; Rothe, Spakowski, & Vogel, 2003), but efficient, polynomial-time algorithms exist when preferences are single-peaked (Brandt, Brill, Hemaspaandra, & Hemaspaandra, 2015). Another example of a desirable property is that single-peaked profiles are guaranteed to possess a transitive majority relation and thus a (weak) Condorcet is guaranteed to exist (Black, 1948, 1958). Furthermore, for single-peaked profiles there exist nondictatorial strategyproof voting rules (Moulin, 1984, 1991) (e.g., selecting a Condorcet winner); such rules do not exist for arbitrary preference profiles by the Gibbard–Satterthwaite Theorem (Gibbard, 1973; Satterthwaite, 1975).

If only incomplete preference information is available, it is a challenge to determine whether voters' preferences can be explained by a single-peaked axis. This is the main goal of this paper: to investigate the algorithmic problem of recognizing whether an incomplete preference profile may be single-peaked, i.e., whether the voters' underlying complete preference orders are single-peaked. The motivation to obtain this information is twofold: First, to know that incomplete preferences may be explainable by a single-peaked axis provides fundamental understanding about the nature of the given preference data set, since the preferences of all voters can be explained by optimal points (peaks) on such an axis. Second, the single-peaked combinatorial structure is highly advantageous for algorithmic purposes, also in the case of incomplete preferences. This claim has been corroborated by recent work, which we discuss below in Section 1.2. In contrast to these two positive aspects, the axiomatic benefits of the single-peaked domain do *not* fully extend to incomplete preferences. In particular, a transitive majority relation is no longer guaranteed (as we discuss later in Section 6.1). It thus deserves special attention that the algorithmic usefulness of incomplete single-peaked preferences is not an immediate consequence of their social choice-theoretic properties.

We now give an overview of the main results in this paper.

1.1 Results

This paper deals with the question of how to determine whether an incomplete preference profile (e.g., a preference profile comprised of partial orders) is single-peaked. This requires a definition of single-peakedness that is applicable to more general types of orders than total orders, i.e., partial orders. To this end, we propose the notion of *possibly single-peakedness*, which is satisfied if—given an incomplete preference profile—there exist a compatible complete profile that is single-peaked. To be more precise, we ask whether a given profile of partial orders can be extended to total orders such that the resulting complete profile is single-peaked (according to the usual definition). Based on this definition, we analyze the corresponding algorithmic recognition problem. In the following, let n denote the size of the preference profile (number of voters) and let m denote the number of candidates. Our main computational results are as follows:

- We prove that determining whether an incomplete preference profile is single-peaked is NP-complete. This is in contrast to the case of complete preferences for which single-peakedness can be determined in linear time (Escoffier, Lang, & Öztürk, 2008). Furthermore, we strengthen this result by showing that NP-completeness still holds if one voter completely specifies his or her preferences, and also for the case of local weak orders.

Apart from these hardness results, we present polynomial-time algorithms for recognizing single-peakedness in relevant subcases:

- First, we consider the problem of determining whether a profile of partial orders is single-peaked for a given axis. We prove this problem to be polynomial-time solvable even for incomplete profiles consisting of partial orders.
- Our most general polynomial-time algorithm is based on the consecutive ones problem. By providing a reduction from single-peaked recognition to the consecutive ones problem, we obtain an $\mathcal{O}(m^2 \cdot n)$ algorithm for profiles of weak orders, i.e., rankings with ties. We have implemented this algorithm; see the discussion section (Section 7) for further details.
- We also provide a direct $\mathcal{O}(m \cdot n)$ algorithm for single-peaked consistency for weak orders, with the requirement that the preference profile contains at least one implicitly specified total order. This algorithm is an improvement over the single-peaked recognition algorithm by Escoffier et al. (2008) since it is applicable to a broader class of preference profiles (weak orders instead of total orders) while maintaining its runtime.
- Furthermore, we present a 2-SAT-based algorithm, which also requires an implicitly specified total order, but is applicable to local weak orders (a generalization of weak orders). This more general algorithm requires $\mathcal{O}(m^3 \cdot n)$ time.
- For the case of top orders (weak orders with all incomparable candidates ranked last), we provide another direct, combinatorial algorithm with a runtime of $\mathcal{O}(m^2 \cdot n)$. This algorithm does not rely on solving consecutive ones instances and thus is conceptually simpler, in particular for implementations.

These algorithms differ in their worst-case runtime estimates, which range from $\mathcal{O}(m \cdot n)$ to $\mathcal{O}(m^3 \cdot n)$. Note that these runtime estimates differ only with respect to the number of candidates. In particular in scenarios with a large number of candidates (e.g., when aggregating web search results, as studied by Betzler, Brederbeck, & Niedermeier, 2014), the actual runtime of these algorithms can be expected to vary.

Finally, we consider the case that voters are unable to provide total orders because their true preferences include some form of indifference. This gives another perspective on weak orders: viewing them not as incomplete but as rankings with ties. We compare our definition of possibly single-peaked preferences to other applicable single-peaked concepts, highlighting benefits and drawbacks of our more general approach.

1.2 Related Work

As mentioned before, the single-peaked property is highly advantageous for algorithmic purposes. This has been demonstrated for many social choice problems based on total (linear)

order preferences, e.g., in the work of Betzler, Slinko, and Uhlmann (2013), Brandt et al. (2015), and Elkind and Ismaili (2015). Recent work has shown that similar benefits also hold for possibly single-peaked preferences, in particular for profiles of weak orders. Peters (2018) shows that several rules that are NP-hard otherwise can be computed in polynomial time for possibly single-peaked profiles of weak orders. This holds for multi-winner rules such as Chamberlin-Courant, Proportional Approval Voting, and the very general class of OWA-based rules. Furthermore, Young’s rule is also polynomial-time computable for possibly single-peaked profiles of weak orders; this follows from work of Peters and Lackner (2017), who prove this result for the more general domain of preferences that are single-peaked on a circle. If we restrict weak orders to two equivalence classes, i.e., we only distinguish between good (approved) and bad (disapproved) candidates, we obtain the class of dichotomous preferences. The notion of possibly single-peakedness directly translates to this setting. Liu and Guo (2016) have shown that Minimax Approval Voting, although NP-hard in general (LeGrand, Markakis, & Mehta, 2007), can be solved in polynomial time for possibly single-peaked dichotomous preferences (see the work of Elkind & Lackner, 2015 containing a detailed discussion of single-peakedness and dichotomous preferences). Overall, the notion of possibly single-peaked preferences has been shown to be beneficial for many algorithmic purposes.

Another interesting effect of single-peaked preferences is that the complexity of manipulative actions (such as manipulation and control) often decreases for single-peaked preferences (Faliszewski, Hemaspaandra, Hemaspaandra, & Rothe, 2011; Hemaspaandra, Hemaspaandra, & Rothe, 2016). When we instead consider more general types of preference orders than total orders we get different behavior depending on the model of single-peakedness used. In general, the model of possibly single-peakedness does not yield such a reduction in complexity compared to the general case for different manipulative actions (Fitzsimmons & Hemaspaandra, 2015, 2016; Menon & Larson, 2017). This is in contrast to the algorithmic benefits mentioned in the last paragraph. However, here a decrease in complexity would not necessarily be desirable. In particular, if one seeks to reduce insincere behavior by choosing voting rules that are computationally hard to manipulate, then a reduction in computational complexity removes this kind of protection (Faliszewski et al., 2011). We also mention that Fitzsimmons and Hemaspaandra (2016) consider the complexity of manipulative actions of the models of single-peakedness discussion in Section 6 as well as another related restriction, single-peaked preferences with outside options (Cantala, 2004).

Recognition algorithms for restricted preference domains are also indispensable for characterization results. The single-peaked domain (Ballester & Haeringer, 2011), the single-crossing domain (Bredereck, Chen, & Woeginger, 2013), and the group-separable domain (Ballester & Haeringer, 2011) have all been characterized by a set of *forbidden subprofiles*. These characterizations are obtained by analyzing under which conditions recognition algorithms certify that a profile is *not* contained in a preference domain. The Unguided Algorithm, as presented in this paper, was used for characterizing preferences that are single-peaked on a circle (Peters & Lackner, 2017). This was possible by relating the problem of recognizing preferences that are single-peaked on a circle to the problem of recognizing possibly single-peaked preferences. Such characterizations do not always exist, e.g., the Euclidean domain cannot be characterized by a finite set of forbidden subprofiles (Chen, Pruhs, & Woeginger, 2017).

Our definition of possibly single-peaked preferences resembles that of a possible winner (Konczak & Lang, 2005): A candidate is a possible winner if—given an incomplete preference profile—there exists a compatible complete profile in which this candidate is a winner. The possible

winner problem is a central, well-studied problem in computational social choice (Walsh, 2007; Betzler, Hemmann, & Niedermeier, 2009; Betzler & Dorn, 2010; Pini, Rossi, Venable, & Walsh, 2011; Xia & Conitzer, 2011; Baumeister & Rothe, 2012; Dery, Kalech, Rokach, & Shapira, 2014), see also the survey by Boutilier and Rosenschein (2016).

The single-peaked domain is far from the only well-studied preference domain; many other domains have been considered from a computational viewpoint. We refer to the survey of Elkind et al. (2017) for an overview. Here, we briefly mention papers that build upon or are directly related to our work. Another major preference restriction is the single-crossing domain, characterized by an ordering of voters. Elkind, Faliszewski, Lackner, and Obraztsova (2015) analyzed the recognition of incomplete single-crossing preferences. This work studies possibly single-crossing profiles (a notion analogously defined to our definition of possibly single-peaked profiles), but found that approaches like consecutive ones and 2-SAT reductions—which we use here—yield weaker results in the single-crossing domain. In particular, fixing the underlying order appears to have a much weaker effect on the recognition problem. Another notion is top-monotonicity (Barberà & Moreno, 2011), for which questions similar to ours have been studied. The results for top-monotonicity resemble those for single-peaked preferences, although new techniques are required: The recognition problem for partial orders is NP-hard (Aziz, 2014) and solvable in polynomial time for weak orders (Magiera & Faliszewski, 2019). Furthermore, Peters (2017) highlights the inherent complexity of detecting more-dimensional Euclidean preferences and discusses possible definitions for weak orders.

More broadly, recommender systems routinely have to deal with incomplete preferences and therefore this issue has received significant attention in the field. We refer to the following works for an overview (Schafer, Frankowski, Herlocker, & Sen, 2007; De Gemmis, Iaquinta, Lops, Musto, Narducci, & Semeraro, 2010; Masthoff, 2015).

1.3 Outline of the Paper

In Section 2, we provide basic definitions from social choice theory. This is followed by Section 3, in which we introduce our notion of possibly single-peaked preferences and study their basic properties. In Section 4, we prove the computational hardness of recognizing possibly single-peaked preferences. All algorithmic results are then presented in Section 5; this section contains the most important contributions of this paper. In Section 6, we discuss indifference, weak orders, and other models of single-peakedness. A final discussion of our results, a note on an implementation of the consecutive ones approach, and suggestions for future research can be found in Section 7. Some proof details (correctness proofs of some algorithms) are delegated to an appendix, Sections A.1, and A.2.

2. Preliminaries

In this paper, preferences are represented by different types of orders (see Figure 1 for examples). The most general type are partial orders: a (*strict*) *partial order* P on a set X is a binary relation on X that is transitive (xPy and yPz implies xPz for all $x, y, z \in X$) and asymmetric (if xPy then not yPx for all $x, y \in X$). We say that y is *ranked above* x if xPy holds. If for two elements $x, y \in X$ neither xPy nor yPx holds, we say that these two elements are *incomparable*; we write $x \sim y$. An element $x \in X$ is *minimal* if there is no element $y \in X$ with yPx ; maximal elements are defined analogously.

A (*strict*) *weak order* is a partial order where the incomparability relation is transitive. Weak orders have the same expressiveness as total preorders¹ (also known as nonstrict weak orders or preference orders). Since a weak order can be considered a ranking with ties, we refer to \sim also as “indifference” in this context (see Section 6, where we discuss several other models of single-peakedness for preferences with indifference). Weak orders are also referred to as bucket orders (e.g. in Fagin, Kumar, Mahdian, Sivakumar, & Vee, 2006); elements that tie are in the same “bucket”.

A *top order* is a weak order where incomparability appears only among minimal elements, i.e., if $a \sim b$ then both a and b are minimal elements in a top order. The *ranked* (i.e., nonminimal) elements of a top order T are those that are not incomparable to any other candidate. We would like to remark that top orders are also known as *top lists* (Dwork, Kumar, Naor, & Sivakumar, 2001; Fagin, Kumar, & Sivakumar, 2003) and as top-truncated votes (Baumeister, Faliszewski, Lang, & Rothe, 2012).

A partial order with no incomparable elements is called *total order*, or, equivalently, a total order is a complete (xPy or yPx for all $x, y \in X$) partial order. Any partial order P can be *extended* to some total order T such that aPb implies aTb ; T is then a (not necessarily unique) *extension* of P .

Finally, we define a *local weak order* P on a set X to be a partial order on X with the following property: there exist sets X_1, X_2 with $X_1 \cup X_2 = X$ such that the elements in X_1 are incomparable to all other elements in X and the profile P restricted to X_2 is a weak order. Intuitively, a local weak order is a weak order together with some isolated elements for which absolutely no information is available.

Throughout the paper, total orders are denoted by $\langle c_1 > c_2 > \dots > c_k \rangle$; the brackets allow us to unambiguously denote total orders consisting of one or even zero elements, i.e, we use $\langle \rangle$ to denote the empty order relation. For top orders, we write $\langle c_1 > c_2 > \dots > c_k > \bullet \rangle$ to denote a top order where c_1, \dots, c_k are ranked as stated and all other elements (usually all remaining candidates in C) are ranked last, i.e., are minimal elements. We sometimes use set operators (\cup, \cap, \setminus) on top orders with the intended meaning that we apply these operators to the corresponding sets of ranked candidates. We describe weak orders with a notation similar to top orders and use \sim to signify incomparability, e.g., $\langle c_1 > c_2 \sim c_3 > c_4 \rangle$.

We would now like to address the usefulness of these types of orders for expressing preferences. Total orders allow the voter to fully specify a strict ranking of options. Given a large set of options, this might be unfeasible. Partial orders, on the other hand, allow the voter to specify the relative strict order of any pair of options subject to transitivity, which can be viewed as a rationality constraint. Thus they can be seen as a very general formalism for representing incomplete preferences. They are compatible with total orders in the sense that partial orders can always be extended to total orders. Weak orders are less general than partial orders but arise in many natural scenarios. For example every real-valued utility function implies a weak order (candidates with the same utility tie, i.e., are incomparable in our sense). Local weak orders correspond to partial real-valued utility functions and thus arise in scenarios where voters do not have knowledge about all candidates. If the elicitation of preferences is costly, one might ask only for the most important (top ranked) options of each voter; in such a case we obtain top

1. Total preorders are binary relations that are transitive and complete (xPy or yPx for all $x, y \in X$). For our algorithmic purposes, it is not relevant whether to use strict weak-orders or total preorders. Our main reason to use strict weak orders as the standard definition is that total preorders are not necessarily partial orders (they are not asymmetric) and hence do not fit in our hierarchy of orders.




Figure 1: The order zoo: examples of different types of orders that are used to describe preferences.

orders. Top orders also are the natural type of order for specifying preferences in some scoring rules. We will further comment on scoring rules and top orders at the end of Section 5.2.

Throughout this paper we use C to denote the set of candidates or options. Votes (or preference orders) are considered to be either partial, local weak, weak, top, or total orders. For a vote V_i , we use $x \succ_i y$ to denote that x is ranked above y . If there is only one vote under consideration, usually denoted by V , we omit the index and write $x \succ y$. If V_i is a weak or top order, we write $x \sim_i y$ if voter i is indifferent between x and y , $x \succsim_i y$ if voter i is either indifferent or ranks x above y .

A tuple (V_1, \dots, V_n) of votes is called a (*preference*) *profile* of $\{\text{partial orders, local weak orders, weak orders, top orders, total orders}\}$, depending on the type of orders. Given a vote V and a set of candidates $C' \subseteq C$, we define $V[C']$ to be the order V restricted to elements in C' . Analogously, given a preference profile $\mathcal{P} = (V_1, \dots, V_n)$, we define $\mathcal{P}[C']$ to be $(V_1[C'], \dots, V_n[C'])$. We denote the number of candidates with m and the number of votes with n .

3. Single-Peaked Profiles

We start by giving a definition for single-peaked profiles of total orders and then extend this definition to partial orders. A central concept is that of an *axis*, which is a total order on C . Let A be an axis. Throughout this paper we write $x \triangleright y$ instead of xAy . (Note that we use \succ for votes and \triangleright for axes.) We formally define single-peakedness for total orders below using triples of candidates, so-called *valleys*. Valleys will be important for many of the results in our paper.

Definition 1. *Let V be a partial order on C . A vote V contains a v-valley with respect to an axis A if there exist $c_1, c_2, c_3 \in C$ such that $c_1 \triangleright c_2 \triangleright c_3$, $c_1 \succ c_2$ and $c_3 \succ c_2$.*

The definition of v-valleys suffices to define single-peakedness for profiles of total orders.

Definition 2. *A profile \mathcal{P} of total orders is single-peaked with respect to A if no vote $V \in \mathcal{P}$ contains a v-valley with respect to A (and thus every vote has only a single “peak”). A profile of total orders is single-peaked consistent (or simply, single-peaked) if there exists some axis A such that \mathcal{P} is single-peaked with respect to A .*

Note that this is similar to how single-peakedness is defined in (Brandt et al., 2015; Elkind et al., 2017), and easily seen to be equivalent. We now want to extend this definition to profiles of partial orders. A natural way is to consider extensions of partial orders to total orders.




Figure 2: Examples of v-valleys and u-valleys

Definition 3. Let $\mathcal{P} = (V_1, \dots, V_n)$ be a profile of partial orders. The profile \mathcal{P} is possibly single-peaked with respect to an axis A if for every $k \in \{1, \dots, n\}$, V_k can be extended to a total order V'_k such that the profile of total orders $\mathcal{P}' = (V'_1, \dots, V'_n)$ is single-peaked with respect to A . A profile of partial orders is possibly single-peaked consistent if it is possibly single-peaked with respect to some axis.

While it is also conceivable to require that every extension is single-peaked (“necessarily” instead of “possibly”), this would yield an extremely restrictive definition, basically requiring that the profile of partial orders already contains “almost” total orders. We consider this definition at the end of Section 6. Apart from that, we focus our attention on the more general definition of possibly single-peaked profiles.

We now define an equivalent definition for possibly single-peakedness based on valleys. For this, we first need the following definition of u-valleys.

Definition 4. Let V be a partial order on C . The vote V contains a u-valley with respect to A if there exist distinct elements $a, b, c, d \in C$ with $a \triangleright b \triangleright d$ and $a \succ b$ as well as $a \triangleright c \triangleright d$ and $d \succ c$.

In Figure 2 a graphical comparison of v- and u-valleys is shown. In the following, we see that these two types of valleys characterize possibly single-peaked incomplete profiles.

Lemma 1. Let $\mathcal{P} = (V_1, \dots, V_n)$ be a profile of partial orders and A an axis. The following two statements are equivalent.

- (i) The profile \mathcal{P} is possibly single-peaked with respect to A .
- (ii) Every vote $V \in \mathcal{P}$ contains neither a u-valley nor a v-valley with respect to A .

Proof. To see that (i) implies (ii), note that if some V_k contained a u-valley in the sense of Definition 4 then every extension V'_k would contain a v-valley. More concretely, if $a, b, c, d \in C$ form a u-valley in V_k then any extension of V_k either contains a v-valley with respect to A on the candidates a, b, c or on b, c, d . Furthermore, if V_k contained a v-valley then so would every extension.

For the other direction, we show that a vote V not containing a valley can be extended to a total order that is single-peaked with respect to A . We recursively define its extension V' starting with its last ranked candidate. Let $V'(1)$ denote the last ranked candidate of V' , $V'(2)$ the second-to-last, etc. For the definition we require two functions:

- $left_A(X)$ is the smallest (leftmost) candidate in X with respect to the axis A .
- $right_A(X)$ is the largest (rightmost) candidate in X with respect to the axis A .

We now define for every $i \in \{1, \dots, m\}$,

$$X_i = C \setminus \{V'(1), \dots, V'(i-1)\} \text{ and}$$

$$V'(i) = \begin{cases} \text{right}_A(X_i) & \text{if } \text{right}_A(X_i) \text{ is a minimal element in } V[X_i], \\ \text{left}_A(X_i) & \text{otherwise.} \end{cases}$$

This definition immediately yields that V' is single-peaked with respect to A : By always choosing one of the two outermost candidates on A (that have not yet been chosen) for the next higher ranked candidate, valleys cannot arise.

It remains to show that V' is indeed an extension of V , i.e., we have to show that for every pair of candidates $a, b \in C$, $a \succ b$ implies $a \succ' b$. Towards a contradiction assume that $a \succ b$ and $b \succ' a$. Let $i \in \{1, \dots, m\}$ such that $V'(i) = a$. We have to consider two cases: $a = \text{left}_A(X_i)$ and $a = \text{right}_A(X_i)$.

Let $a = \text{left}_A(X_i)$ and $d = \text{right}_A(X_i)$. Since $V'(i) = a$, we know that there has to exist a $c \in X_i$ with $d \succ c$. Observe that $a \triangleright c \triangleright d$ has to hold. Furthermore, either $a \triangleright b \triangleright d$ or $b = d$ holds. If $a \triangleright b \triangleright d$ holds then a, b, c, d form a u-valley. If $b = d$ then $a \triangleright c \triangleright b$, $a \succ b$ and $b \succ c$ holds: a v-valley. Both cases contradict our assumption that V does not contain valleys with respect to A .

Now, let $a = \text{right}_A(X_i)$. This immediately yields a contradiction to the definition of $V'(i)$, since $b \in X_i$, $a \succ b$ and hence a is not a minimal element in $V[X_i]$. \square

This lemma immediately yields a polynomial-time algorithm for checking whether an incomplete profile is possibly single-peaked with respect to a given axis:

Proposition 1. *Verifying whether a profile of partial orders is single-peaked with respect to a given axis can be done in $\mathcal{O}(m^4 \cdot n)$ time.*

Proof. For every quadruple of candidates and every vote, an algorithm has to check whether a u- or v-valley arises (by Lemma 1). \square

Let $\mathcal{T} \in \{\text{partial order, local weak order, weak order, top order, total order}\}$ be a type of order. The remainder of the paper is dedicated to a more challenging computational problem:

\mathcal{T} SINGLE-PEAKED CONSISTENCY	
Instance:	A profile \mathcal{P} of type \mathcal{T} and a set of candidates C .
Question:	Is \mathcal{P} possibly single-peaked consistent?

Note that in contrast to the problem in Proposition 1, the input of this problem does not include an axis. The TOTAL ORDER SINGLE-PEAKED CONSISTENCY problem is known to be solvable in polynomial time as witnessed by several algorithms. Historically, the first algorithm to solve this problem was due to Bartholdi and Trick (1986) and is based on the consecutive ones problem, which we will encounter in Section 5.1. This approach yields a runtime of $\mathcal{O}(m^2 \cdot n)$, which was improved by direct, combinatorial algorithms first to $\mathcal{O}(m \cdot n + m^2)$ by Doignon and Falmagne (1994) and finally to $\mathcal{O}(m \cdot n)$ by Escoffier et al. (2008). In the next section, we show that a polynomial-time result is unlikely to exist for the case of partial orders and even local weak orders.

4. Hardness Results

Theorem 1. LOCAL WEAK ORDER SINGLE-PEAKED CONSISTENCY is NP-complete.

Proof. We reduce from the NP-complete BETWEENNESS problem (Opatrny, 1979). A BETWEENNESS instance consists of a finite set S and a set T containing (ordered) triples of distinct elements of S . The decision problem asks whether there is a total order L such that for every triple $(a, b, c) \in T$ we have either $aLbLc$ or $cLbLa$. Intuitively, a triple $(a, b, c) \in T$ corresponds to the constraint that b has to lie “in between” a and c on the total order L .

We construct a profile of local weak orders \mathcal{P} with a candidate set $C = S$, i.e., we identify elements in S with candidates in C . The preference profile \mathcal{P} consists of two votes for each triple (a, b, c) : the partial order $\{a \succ c, b \succ c\}$ and the partial order $\{b \succ a, c \succ a\}$. These two votes form a valley on any axis with c between a and b and on any axis with a between b and c . Thus b has to be between a and c on any single-peaked axis. We are now going to show that \mathcal{P} is possibly single-peaked consistent if and only if the BETWEENNESS instance is a yes-instance.

“ \Rightarrow ” Assume that there exists an extension of \mathcal{P} , \mathcal{P}^{ext} , and an axis, A , such that \mathcal{P}^{ext} is single-peaked with respect to A . By Lemma 1 we know that this implies that no v-valleys exist. Since for every triple $(a, b, c) \in T$ both the vote $\{a \succ c, b \succ c\}$ and $\{b \succ a, c \succ a\}$ are contained in \mathcal{P} , we have that neither $a \triangleright c \triangleright b$, $b \triangleright c \triangleright a$, $b \triangleright a \triangleright c$, nor $c \triangleright a \triangleright b$ can hold. Consequently it has to hold that either $a \triangleright b \triangleright c$ or $c \triangleright b \triangleright a$ holds and thus b is “in between” a and c .

“ \Leftarrow ” Assume that there exists a set T such that all constraints in T are satisfied. It is easy to verify that \mathcal{P} is possibly single-peaked with respect to T . □

Corollary 1. PARTIAL ORDER SINGLE-PEAKED CONSISTENCY is NP-complete.

The proof of Theorem 1 uses preference profiles where the votes contain very little information: only two pairs of candidates are comparable in each vote. We know that determining single-peaked consistency is possible in polynomial time if every vote is a total order, i.e., all votes contain complete information. What happens if one voter provides complete information? Having a single completely specified vote has been found to be helpful in a related context: It allows for the efficient elicitation of single-peaked preferences using only few comparison queries (Conitzer, 2009) and thus the communication complexity of preference elicitation is reduced. However, in our case such a voter does not provide enough additional information for a decrease in (computational) complexity.

Theorem 2. The PARTIAL ORDER SINGLE-PEAKED CONSISTENCY problem is NP-complete, even if the given preference profile contains a total order.

Proof. We reduce from SET SPLITTING: Let X be a finite set. Given a collection Z of subsets of X , is there a partition of X into two subsets X_1 and X_2 such that no subset of Z is contained entirely in either X_1 or X_2 ? This problem is NP-complete even if all sets in Z have cardinality three (Garey & Johnson, 1979).

Let $X = \{c_1, \dots, c_m\}$. For the construction, we identify the elements of X with candidates and add an additional candidate x . For each set $\{c_i, c_j, c_k\} \in Z$ with $i < j < k$ we introduce one vote: $\{c_i \succ c_j, x \succ c_k\}$. In addition, we add the vote $x \succ c_m \succ \dots \succ c_1$. We claim that the resulting preference profile \mathcal{P} is possibly single-peaked consistent if and only if (X, Z) is a SET SPLITTING yes-instance.

Assume that \mathcal{P} is possibly single-peaked with respect to an axis A . We define X_1 to be the candidates on A left of x and X_2 those that are right of x . We will show that there is no

subset of Z entirely contained in X_1 or X_2 . Towards a contradiction assume that $\{c_i, c_j, c_k\} \in Z$ with $i < j < k$ are contained in X_1 . Then it has to hold that, on A , c_i, c_j, c_k are all left of x . Furthermore, from the vote $\langle x \succ c_m \succ \dots \succ c_1 \rangle$ then it follows that the relative order on A of x, c_i, c_j, c_k has to be $c_i \triangleright c_j \triangleright c_k \triangleright x$. However, then the vote $\{c_i \succ c_j, x \succ c_k\}$ contains a u-valley with respect to this order. Assuming that $\{c_i, c_j, c_k\} \in Z$ are contained in X_2 leads to the same contradiction. Thus, X_1 and X_2 indeed certify that (X, Z) is a yes-instance.

For the other direction, assume that (X, Z) is a yes-instance and that X_1 and X_2 are the corresponding partition. Let an axis A be defined as the elements in X_1 with indices in increasing order followed by x followed by the elements in X_2 with indices in decreasing order. We claim that A is an axis for \mathcal{P} . Clearly, the vote $\langle x \succ c_m \succ \dots \succ c_1 \rangle$ is single-peaked with respect to A . Let us consider a vote $\{c_i \succ c_j, x \succ c_k\}$ with $i < j < k$. Since X_1 and X_2 are a valid partition, at least one of c_i, c_j, c_k has to be left of x and another one right. This rules out that a u-valley is formed and thus all votes are possibly single-peaked with respect to A . \square

As we will see in the following section, these two hardness results establish the tractability frontier.

5. Algorithms

In this section, we present several polynomial-time algorithms for recognizing possibly single-peaked profiles.

5.1 The Consecutive Ones Approach

Our first algorithm in this section solves the WEAK ORDER SINGLE-PEAKED CONSISTENCY problem in polynomial time. It uses a reduction to the problem of detecting the *consecutive ones property* in a binary matrix, i.e., a matrix consisting of zeros and ones. Such a matrix has the consecutive ones property if its columns can be permuted in such a way that in all rows the ones appear consecutively. The corresponding decision problem is the following:

CONSECUTIVE ONES	
Instance:	A binary matrix M .
Question:	Does M possess the consecutive ones property, i.e., does there exist a permutation of the columns of M such that in each row all of the ones are consecutive?

The consecutive ones property was originally defined by Fulkerson and Gross (1965) and shown solvable in polynomial time. More specifically, they showed that given an $s \times t$ matrix this problem can be solved in $\mathcal{O}(s \cdot t^2)$. Booth and Lueker (1976) improved on this result by finding an $\mathcal{O}(s \cdot t)$ algorithm² through the development and use of the novel PQ-tree data structure. With this data structure it is not only possible to find one valid permutation of columns (if it exists), but to compactly represent all possible column permutations that witness the consecutive ones property. Subsequent work has improved these results in various ways (Meidanis, Porto, & Telles, 1998; Habib, McConnell, Paul, & Viennot, 2000; Hsu, 2002; McConnell, 2004; Raffinot, 2011).

2. More precisely, their algorithm has a runtime of $\mathcal{O}(s+t+f)$, where f is the number of ones in M . In our case, the matrices obtained from the reduction have $\Theta(s \cdot t)$ one entries and hence this algorithm has a runtime of $\mathcal{O}(s \cdot t)$.

Bartholdi and Trick (1986) were the first to relate the problem of recognizing single-peaked preferences to the consecutive ones problem. We slightly modify their approach for profiles of total orders to be applicable to profiles of weak orders, and by that solve the WEAK ORDER SINGLE-PEAKED CONSISTENCY problem. The reduction works as follows:

Construction 1. Let $\mathcal{P} = (V_1, \dots, V_n)$ be a profile of weak orders over candidate set C with $|C| = m$. For each V_i we construct an $m \times m$ binary matrix X_i . We assume that the rows and columns of X_i are indexed by C . For $a, b \in C$, the entry of X_i is defined as:

$$X_i(a, b) = \begin{cases} 0 & \text{if } a \succ_i b \\ 1 & \text{if } b \succ_i a \text{ or } a \sim_i b \end{cases}.$$

Finally, the matrices X_1, \dots, X_n are row-wise concatenated to obtain the $mn \times m$ matrix $X_{\mathcal{P}}$.

Example 1. Consider the preference profile $\mathcal{P} = (V_1, V_2)$ with $V_1 = \langle a \sim c \succ b \succ e \sim d \succ f \rangle$ and $V_2 = \langle a \succ b \succ c \succ e \sim d \succ f \rangle$. We construct X_1 and X_2 :

$$X_1 = \begin{matrix} & a & b & c & d & e & f \\ \begin{matrix} a \\ b \\ c \\ d \\ e \\ f \end{matrix} & \begin{bmatrix} 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 \end{bmatrix} \end{matrix} \qquad X_2 = \begin{matrix} & a & b & c & d & e & f \\ \begin{matrix} a \\ b \\ c \\ d \\ e \\ f \end{matrix} & \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 \end{bmatrix} \end{matrix}$$

We then row-wise concatenate X_1 and X_2 to construct $X_{\mathcal{P}}$.

$$X_{\mathcal{P}} = \begin{matrix} & a & b & c & d & e & f \\ \begin{matrix} a \\ b \\ c \\ d \\ e \\ f \\ a \\ b \\ c \\ d \\ e \\ f \\ a \\ b \\ c \\ d \\ e \\ f \end{matrix} & \begin{bmatrix} 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 \end{bmatrix} \end{matrix} \qquad X'_{\mathcal{P}} = \begin{matrix} & b & a & c & d & e & f \\ \begin{matrix} b \\ a \\ c \\ d \\ e \\ f \\ b \\ a \\ c \\ d \\ e \\ f \\ b \\ a \\ c \\ d \\ e \\ f \end{matrix} & \begin{bmatrix} 0 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 \end{bmatrix} \end{matrix}$$

Observe that $X_{\mathcal{P}}$ has the consecutive ones property, as witnessed by $X'_{\mathcal{P}}$. The corresponding permutation of columns (bacdef) directly corresponds to an ordering of candidates, and indeed \mathcal{P} is possibly single-peaked with respect to $b \triangleright a \triangleright c \triangleright d \triangleright e \triangleright f$.

To prove the correctness of Construction 1, we first define a valley-based characterization of possibly single-peaked profiles of weak orders, analogous to Lemma 1. We see that for profiles of weak orders, u-valleys are not relevant.

Lemma 2. *Let $\mathcal{P} = (V_1, \dots, V_n)$ be a profile of weak orders and A an axis. The following two statements are equivalent.*

- (i) *The profile \mathcal{P} is possibly single-peaked with respect to A .*
- (ii) *Every vote $V \in \mathcal{P}$ does not contain a v-valley with respect to A .*

Proof. That statement (i) implies (ii) is a special case of Lemma 1. For the other direction, every u-valley on $a \triangleright b \triangleright c \triangleright d$ (see Figure 2) implies a v-valley either on $a \triangleright b \triangleright d$ or on $a \triangleright c \triangleright d$. \square

Using Construction 1 and Lemma 2, we can now show that:

Theorem 3. *The WEAK ORDER SINGLE-PEAKED CONSISTENCY problem can be solved in $\mathcal{O}(m^2 \cdot n)$ time.*

Proof. Let \mathcal{P} be a preference profile of weak orders. We will show that V is possibly single-peaked if and only if the matrix $X_{\mathcal{P}}$, as obtained by Construction 1, has the consecutive ones property. This proof closely follows the argument made by Bartholdi and Trick (1986) for total orders, but requires Lemma 2 to relate possible single-peaked weak order to the consecutive ones property.

If \mathcal{P} is possibly single-peaked with respect to an axis A then by Lemma 2 we know that no preference order $V \in \mathcal{P}$ contains a v-valley with respect to A . When the columns of the matrix $X_{\mathcal{P}}$ are permuted with respect to the axis A , no row will contain the sequence $\dots 1 \dots 0 \dots 1 \dots$ since this corresponds to a preference order that strictly decreases and then strictly increases along the axis A (a v-valley). Therefore X has the consecutive ones property.

For the other direction suppose that V is not possibly single-peaked. Then by Lemma 2 we know that for every possible axis there exists a preference order $V \in \mathcal{P}$ such that V contains a v-valley with respect to that axis. So every permutation of the columns of $X_{\mathcal{P}}$ will correspond to an axis where some preference order has a v-valley. As stated in the proof of the other direction, a v-valley corresponds to a row containing the sequence $\dots 1 \dots 0 \dots 1 \dots$, so clearly X does not have the consecutive ones property.

Since it takes $\mathcal{O}(m^2 \cdot n)$ time to construct $X_{\mathcal{P}}$ and $\mathcal{O}(m^2 \cdot n)$ time to solve the corresponding CONSECUTIVE ONES problem, we can solve the WEAK ORDER SINGLE-PEAKED CONSISTENCY problem in $\mathcal{O}(m^2 \cdot n)$ time. \square

As mentioned earlier, the approach using PQ-trees by Booth and Lueker (1976) is able to compute all possible permutations that witness the consecutive ones property. Thus, Theorem 3 actually allows us to compute all possibly single-peaked axes. As these permutations are stored in a compact way, the runtime of $\mathcal{O}(m^2 \cdot n)$ time holds even for a potentially exponential number of axes.

To sum up, the consecutive ones approach grants us a simple and quick solution for the WEAK ORDER SINGLE-PEAKED CONSISTENCY problem. The downside is that by relying on the full consecutive ones machinery, any practical implementation has to deal with the nontrivial PQ-tree data structure or related concepts (Meidanis et al., 1998; Habib et al., 2000; McConnell, 2004). In the following two sections we present direct, combinatorial algorithms.

5.2 The Guided Algorithm

We now present a second polynomial-time algorithm for profiles of weak orders. In contrast to consecutive ones approach of Section 5.1, this algorithm requires an additional condition: The input profile must contain at least one total order, the *guiding vote*, to guide the placement of candidates on the axis. Due to this extra information, we gain two major benefits: First, the algorithm is conceptually simpler than the more complex algorithms for solving the consecutive ones problem. This has clear benefits for implementing the algorithm, and is also necessary for the use in characterization proofs (see, e.g., the work of Peters and Lackner (2017) as discussed in Section 1.2). Second, this algorithm is faster. We achieve a runtime of $\mathcal{O}(m \cdot n)$ compared to the runtime of $\mathcal{O}(m^2 \cdot n)$ by the consecutive ones approach.

Theorem 4. *If the profile contains a total order, the WEAK ORDER SINGLE-PEAKED CONSISTENCY problem can be solved in $\mathcal{O}(m \cdot n)$ time.*

Theorem 4 is based on Algorithm 1, which we refer to as the *Guided Algorithm*. Without loss of generality, we assume that the guiding vote is $\langle c_m \succ c_{m-1} \succ \dots \succ c_1 \rangle$, i.e., we number the candidates based on the guiding vote. The requirement that this guiding vote is part of the input can be weakened considerably. It suffices that a guiding vote is implicitly contained in the profile. The following procedure finds such an implicit guiding vote, if one exists: Look for some vote with a unique last ranked candidate. This candidate is ranked last in the guiding vote. Remove this candidate from the profile and repeat this step to obtain the second-to-last element in the guiding vote, etc. If at some point no vote has a unique last-ranked candidate, no implicit guiding vote exist. As in each step there might be more than one last-ranked candidate to choose from, such a guiding vote is not unique. However, it is straightforward to prove that adding any implicit guiding vote obtained in this way does not change whether the profile is possibly single-peaked. This procedure can be implemented in $\mathcal{O}(m \cdot n)$ time.

Example 2. *Let us consider the profile $\mathcal{P} = (V_1, V_2, V_3)$ with $V_1 = \langle a \succ b \succ c \sim d \succ e \rangle$, $V_2 = \langle a \sim b \sim c \succ d \succ e \rangle$, and $V_3 = \langle e \sim d \succ b \succ c \succ a \rangle$. This profile clearly does not contain a total order. However, a guiding vote can be found implicitly: Candidate e is (uniquely) ranked last in V_1 , so we remove this candidate. In this reduced profile, V_2 ranks d last. We again remove this candidate and see that V_1 now equals $\langle a \succ b \succ c \rangle$. Thus, we can use $\langle a \succ b \succ c \succ d \succ e \rangle$ as our guiding order.*

The algorithm has a simple structure: The lowest ranked candidate in the guiding vote, c_1 , is placed on the rightmost position of the axis (this choice is arbitrary.) Starting with the second lowest ranked candidate in the guiding vote, c_2 , the candidates are successively placed on the axis—either at the leftmost or rightmost still-available position. The lists A_L and A_R correspond to the left-hand and right-hand side of the axis under construction. For each candidate, we test whether it can be placed on the right-hand side or left-hand side without creating a v-valley; u-valleys can be ignored by Lemma 2. If only one of these options is viable, the candidate is placed accordingly. If both left and right are possible, we place the candidate arbitrarily right. If neither is possible, the preference profile is not single-peaked.

Testing whether a vote V_k imposes restrictions on the placement of a candidate is achieved by four conditions. These conditions distinguish four categories of candidates: candidates in A_R , candidates in A_L , candidates that have not yet been placed ($C_{>i} = \{c_{i+1}, \dots, c_m\}$), and the candidate that is currently under consideration (c_i). We are only checking for valleys that




Figure 3: Graphical representation of the conditions testing whether c_i can be placed on the right-hand side (R1, R2) or on the left-hand side (L1, L2)

include c_i . This gives rise to the following four conditions: (R1) and (R2) test whether placing c_i on the right-hand side leads to valleys, (L1) and (L2) do the same for the left-hand side. Figure 3 displays a graphical representation. Note that we do not verify whether a v-valley arises with $a_\ell \succ c_i$ and $a_r \succ c_i$, where $a_\ell \in A_L$ and $a_r \in A_R$; such a valley would have already been detected at an earlier stage of the algorithm (cf. correctness proof).

Since we are working with weak orders, we do not have to consider every candidate triple possibly fulfilling these conditions but have to check only maximal or minimal candidates. More specifically, checking whether there is a candidate $c \in A_L$ and $c' \in C_{>i}$ with $c \succ c'$ is equivalent to whether any maximal element in A_L is preferred to some minimal element in $C_{>i}$ (subject to \succ). For $k \in [n]$, let $\min_k(X)$ denote a function that picks an element in X that is minimal with respect to \succ_k . The function $\max_k(X)$ is defined analogously. Now, we can formally define the four conditions:

$$c_i \succ_k \min_k(C_{>i}) \quad \text{and} \quad \max_k(A_L) \succ_k \min_k(C_{>i}) \quad (\text{R1})$$

$$\max_k(C_{>i}) \succ_k c_i \quad \text{and} \quad \max_k(A_R) \succ_k c_i \quad (\text{R2})$$

$$c_i \succ_k \min_k(C_{>i}) \quad \text{and} \quad \max_k(A_R) \succ_k \min_k(C_{>i}) \quad (\text{L1})$$

$$\max_k(C_{>i}) \succ_k c_i \quad \text{and} \quad \max_k(A_L) \succ_k c_i \quad (\text{L2})$$

Using these four definitions, we can give a succinct description of the Guided Algorithm, (Algorithm 1). For a correctness proof of Algorithm 1, we refer the reader to the appendix, Section A.1.

Example 2 (continued). We apply the Guided Algorithm to $\mathcal{P} = (V_1, V_2, V_3)$ with $V_1 = \langle a \succ b \succ c \sim d \succ e \rangle$, $V_2 = \langle a \sim b \sim c \succ d \succ e \rangle$, and $V_3 = \langle e \sim d \succ b \succ c \succ a \rangle$. We use the implicitly found guiding order $\langle a \succ b \succ c \succ d \succ e \rangle$ and thus start with $A_L = \langle \rangle$ and $A_R = \langle e \rangle$. We now want to place d on the axis. None of the four conditions is satisfied for voters V_1 and V_2 , but V_3 satisfies condition (L1). Thus, d has to be placed right, i.e., $A_R \leftarrow \langle d \triangleright e \rangle$. We continue with the third-lowest candidate on the guiding order: c . Here, we encounter a problem: when considering V_3 , we see that c satisfies both (R2) and (L1). Thus, the algorithm cannot place c and returns `not_single_peaked`. Using other implicitly given guiding orders, e.g., $\langle d \succ b \succ c \succ e \succ a \rangle$, would have yielded the same outcome.

Theorem 4 claims that the Guided Algorithm requires $\mathcal{O}(m \cdot n)$ time. Note that this is only possible if the conditions (R1, R2, L1, L2) can be checked in constant time. Thus, the minima

Input: A set of candidates C , a preference profile of weak orders $\mathcal{P} = (V_1, \dots, V_n)$ including a guiding vote $c_m \succ c_{m-1} \succ \dots \succ c_1$.

Output: An axis A or `not_single_peaked`.

```

1  $A_L \leftarrow \langle \rangle, A_R \leftarrow \langle c_1 \rangle$ 
2 for  $i \leftarrow 2 \dots m$  do
3    $right \leftarrow true; left \leftarrow true$ 
4   for  $k \leftarrow 2 \dots n$  do
5     if condition (R1) or (R2) holds then
6        $right \leftarrow false$ 
7     if condition (L1) or (L2) holds then
8        $left \leftarrow false$ 
9   if  $right = true$  then
10     $A_R \leftarrow \langle c_i \triangleright A_R \rangle$ 
11  else
12    if  $left = true$  then
13       $A_L \leftarrow \langle A_L \triangleright c_i \rangle$ 
14    else
15      return not_single_peaked
16 return  $A_L \triangleright A_R$ 

```

Algorithm 1: The Guided Algorithm

and maxima have to be computable in constant time. For $\max_k(A_L)$ and $\max_k(A_R)$ this is easily possible by storing and updating these two values: if c_i is placed left, we update $\max_k(A_L)$ in case c_i is the new maximum (with respect to \succ_k); if c_i is placed right, we proceed analogously $\max_k(A_R)$. For computing a minimal value of $C_{>i}$, observe that the set $C_{>i}$ becomes smaller with increasing i . Thus, a minimal value of $C_{>i}$ might disappear at some point and a new (larger) value has to be chosen. The new minimum is the smallest element (with respect to \succ_k) in $C_{>i}$ that is at least as large as the old minimum. If we maintain pointers to the minimum elements, the amortized cost of this update procedure is $\mathcal{O}(1)$. A maximal value of $C_{>i}$ can be found analogously.

5.3 The Unguided Algorithm

Our next algorithm is applicable to top orders. As it is not dependent on a guiding vote, we refer to it as the Unguided Algorithm. The Unguided Algorithm has a runtime of $\mathcal{O}(m^2 \cdot n)$, the same as achieved by the consecutive ones approach. We see that from the perspective of worst-case complexity the Unguided Algorithm is inferior to the consecutive ones approach: while having the same worst-case runtime, it is applicable only to a smaller domain (top orders vs. weak orders). However, a point made in favor of the Guided Algorithm also applies here: The main strength of the Unguided Algorithm is its relative simplicity. It can be seen as a solution method for certain consecutive ones instances and—due to this specialization—does not require the full power of consecutive ones machinery. It is thus easier to implement and can be used in characterization proofs.³

3. Preliminary work with Dominik Peters indicates that the Unguided Algorithm can be used to characterize the single-peaked-or-caved domain: A total order A is single-peaked-or-caved with respect to A if either T

Input: A set of candidates C and a connected preference profile of top orders $\mathcal{P} = (V_1, \dots, V_n)$.

Output: An axis A or *not_single_peaked*.

```

1  foreach  $c_{start} \in C$  do
2       $A \leftarrow \langle c_{start} \rangle$ 
3      for  $i \leftarrow 1 \dots m$  do
4          foreach vote  $V \in \mathcal{P}$  that has  $a_i$  as its top-ranked candidate do
5              if  $A \oplus V = \text{incompatible}$  then
6                  | Continue with next  $c_{start} \in C$  in line 1.
7              else  $A \leftarrow A \oplus V$ 
8          if  $|A| = i$  and  $i < m$  then
9               $V \leftarrow \text{IntersectingVote}(A)$ 
10             if  $a_i \notin V$  then
11                 | Continue with next  $c_{start} \in C$  in line 1.
12             Let  $x$  be a new candidate, distinct from those in  $C$ .
13              $C' \leftarrow \{c \in V \mid c \succ a_i\} \cup \{a_i, x\}$ 
14             for  $k \leftarrow 1 \dots n$  do
15                 |  $V'_k \leftarrow \text{RepTop}(V_k, C \setminus (A \cup C'), x)$ 
16              $\mathcal{P}' \leftarrow (V'_1, \dots, V'_n)[C']$ 
17              $A' \leftarrow$  axis returned by Guided Alg. with input  $(\mathcal{P}', V[C'], a_i, x)$ 
18             if  $A' = \text{not\_single\_peaked}$  then
19                 | Continue with next  $c_{start} \in C$  in line 1.
20             else  $A \leftarrow A \triangleright A'[C' \setminus \{x\}]$ 
21         return  $A$ 
22 return not_single_peaked
    
```

Algorithm 2: The Unguided Algorithm

In the description of the algorithm we assume that the input preference profile is *connected*. Let us consider a simple graph with vertices corresponding to candidates. We connect two vertices with an edge whenever the corresponding two candidates are both ranked in some top order. A profile of top orders is called connected if this graph has only one connected component. This assumption does not limit the applicability: if two or more connected components exist in this graph, we can use the algorithm for each component (i.e., its respective candidates and votes) and concatenate the resulting axes in arbitrary order.

The Unguided Algorithm (Algorithm 2) works as follows: First, we choose a candidate c_{start} which is going to be the leftmost candidate on the axis A . Since we have no guiding vote, each candidate might be placed at the leftmost position. Hence we loop over all candidates (line 1). The corresponding axis under construction is $A = \langle c_{start} \rangle$. We now aim to complete this axis by adding candidates to the right in such a way that all votes are single-peaked with respect to this axis. To this end we employ the loop in line 3. In this loop (variable i) we infer from the already placed candidate a_i (the i -th candidate on A from left) the candidate a_{i+1} (or even more

or its reverse is single-peaked with respect to A . This domain is clearly more general than the single-peaked domain, but less general than the single-peaked-on-a-circle domain (Peters & Lackner, 2017).

candidates farther to the right), or infer that A cannot be extended to a single-peaked axis and thus try another start candidate.

Lines 4 to 7 are based on the following observation: Let us assume that at a certain point $A = \langle c_1 \triangleright c_2 \triangleright c_3 \rangle$ and $V = \langle c_3 \succ c_2 \succ c_4 \succ c_5 \succ \bullet \rangle \in \mathcal{P}$. Since c_3 , the peak of V , is already contained in A , there is only one compatible extension of A : $\langle c_1 \triangleright c_2 \triangleright c_3 \triangleright c_4 \triangleright c_5 \rangle$. We formalize this extension operation with the \oplus operator:

Definition 5. *Let A be an incomplete axis and V a top order. Furthermore, let $V[C \setminus A] = \langle c'_1 \succ c'_2 \succ \dots \succ c'_j \succ \bullet \rangle$. We define $A \oplus V = \langle A \triangleright c'_1 \triangleright c'_2 \triangleright \dots \triangleright c'_j \rangle$ if V is possibly single-peaked with respect to this axis and $A \oplus V = \mathbf{incompatible}$ otherwise.*

The loop in line 4 considers all votes V that have candidate a_i as their top-ranked candidate (i.e., their maximal element). If $A \oplus V = \mathbf{incompatible}$ then A cannot be extended to a complete, single-peaked axis and we consider the next $c_{\text{start}} \in C$ in line 1. Otherwise, we obtain a new incomplete axis $A \leftarrow A \oplus V$.

It might be the case that the $(i + 1)$ st candidate on A has not yet been determined after these steps. The lines 8 to 20 deal with this case. As the input profile is connected, there has to be at least one vote that ranks both a candidate on A and a candidate that has not been placed yet. The procedure `IntersectingVote` in line 9 returns such a vote V with $A \cap V \neq \emptyset$ and $V \setminus A \neq \emptyset$. This procedure can be efficiently precomputed in such a way that it requires only $\mathcal{O}(m)$ time to provide an answer; details can be found in the proof of Theorem 5.

Let V be a vote returned by `IntersectingVote`. It holds that V 's top-ranked candidate is not placed on A yet: If its top-ranked candidate were contained in A , then V would have been already considered in the first part of the algorithm (lines 4 to 7). If V does not contain a_i (and thus a_i is ranked last in V), A cannot be extended to a single-peaked axis (lines 10 and 11). Now, we employ the Guided Algorithm of Section 5.2 to find a further extension of A . The main idea is to use V as a guiding vote and find an axis for the candidates in $\{c \in V \mid c \succ a_i\}$. In principle, this axis can be found independently of the existing axis A . However, the leftmost and rightmost candidates have to be chosen with regard to “external” considerations: The leftmost candidate has to be a_i , otherwise A and the newly obtain partial axis A' could not be combined. For the rightmost candidate, we have to consider votes with candidates placed on the axis in future steps. The following example illustrates the issue.

Example 3. *Let $A = \langle a \rangle$, $V_1 = \langle b \succ c \succ a \succ \bullet \rangle$ and $V_2 = \langle c \succ d \succ \bullet \rangle$. The vote V_1 intersects A and hence $C' = \{a, b, c\}$. We employ the Guided Algorithm and might obtain $A' = \langle a \triangleright c \triangleright b \rangle$.⁴ Now observe that $A \oplus A' = A'$ can no longer be extended in a way that it is single-peaked for V_2 . This would have been possible if c had been chosen as the rightmost candidate in A' .*

As we can see from this example, we sometimes have to “force” the rightmost candidate in A' . We do this by adding an additional candidate x to every vote (line 14 to 16). Hence, our candidate set under consideration is at this point $C' \leftarrow \{c \in V \mid c \succ a_i\} \cup \{a_i, x\}$. The new candidate x is now placed in each vote at the position of the highest-ranked candidate among those not contained in $A \cup C'$. This is done by the `RepTop` function: `RepTop(V, D, x)` finds the one candidate in vote V that is the highest-ranked among the candidates in D and replaces it with candidate x . By forcing this element x to be the rightmost candidate, we ensure that A' is chosen with consideration to all votes with ranked candidates not in C' .

4. Whether we obtain this axis or $\langle a \triangleright b \triangleright c \rangle$ depends on whether the algorithm prefers placing candidates to the left or to the right if both choices are possible.

Example 3 (continued). We apply $\text{RepTop}(V, D, x)$ to the votes V_1 and V_2 with candidate sets $C' = \{a, b, c, x\}$ and $D = \{d\}$. We obtain the votes $V_1'[C'] = \langle b \succ c \succ a \succ x \rangle$ and $V_2'[C'] = \langle c \succ x \succ \bullet \rangle$. Now, the Guided Algorithm can only return the axis $\langle a \triangleright b \triangleright c \triangleright x \rangle$.

We obtain a profile \mathcal{P}' on the candidate set C' . We now run the Guided Algorithm with input $(\mathcal{P}', V[C'], A', a_i, x)$, i.e., we employ the Guided Algorithm for the profile \mathcal{P}' and guiding vote $V[C']$. Furthermore, we require that the leftmost candidate on the axis is a_i and the rightmost is x . The Guided Algorithm either returns `not_single_peaked` or an axis A' . If it returns `not_single_peaked`, the next $c_{\text{start}} \in C$ is considered. Otherwise, we continue with the extended axis $A \leftarrow A \oplus A'[C' \setminus \{x\}]$, i.e., candidate x is not placed on A .

Theorem 5. *The TOP ORDER SINGLE-PEAKED CONSISTENCY problem can be solved in $\mathcal{O}(m^2 \cdot n)$ time.*

We refer the reader to the appendix, Section A.2, for the correctness proof of Algorithm 2 and its runtime calculations.

Example 4. Let us illustrate the Unguided Algorithm with a full example. Let $\mathcal{P} = (V_1, V_2, V_3, V_4)$ with $V_1 = \langle b \succ c \succ a \succ \bullet \rangle$ (as in Example 3), $V_2 = \langle c \succ d \succ \bullet \rangle$ (also as in Example 3), $V_3 = \langle f \succ g \succ h \succ e \succ a \succ \bullet \rangle$, and $V_4 = \langle h \succ g \succ f \succ \bullet \rangle$. We assume that the algorithm starts with $c_{\text{start}} = h$, as this choice leads to a successful run. We have exactly one vote with h as its top-ranked candidate: V_4 . Thus, $A \leftarrow A \oplus V_4 = \langle h \rangle \oplus \langle h \succ g \succ f \succ \bullet \rangle = \langle h \triangleright g \triangleright f \rangle$. We now look for votes with f as their top-ranked choice; there is again only one: V_3 . As before, $A \leftarrow A \oplus V_3 = \langle h \triangleright g \triangleright f \rangle \oplus \langle f \succ g \succ h \succ e \succ a \succ \bullet \rangle = \langle h \triangleright g \triangleright f \triangleright e \triangleright a \rangle$. This is by Definition 5, since V_3 is possibly single-peaked with respect to $\langle h \triangleright g \triangleright f \triangleright e \triangleright a \rangle$. We now have reached the situation described in Example 3. Thus, using the Guided Algorithm, we complete the axis to $A = \langle h \triangleright g \triangleright f \triangleright e \triangleright a \triangleright b \triangleright c \rangle$ and see that $\mathcal{P} = (V_1, V_2, V_3, V_4)$ is possibly single-peaked with respect to this so found axis.

5.4 A 2-SAT-Based Algorithm

Theorem 2, Theorem 3, and Theorem 4 leave open the case of profiles of local weak orders which contain at least one total order. Here, we show that this case is also polynomial-time solvable.

Theorem 6. *If the given profile contains a total order, the LOCAL WEAK ORDER SINGLE-PEAKED CONSISTENCY problem can be solved in $\mathcal{O}(n \cdot m^3)$ time.*

We encode a LOCAL WEAK ORDER SINGLE-PEAKED CONSISTENCY instance in a 2-SAT instance. The 2-SAT problem asks whether a Boolean formula in conjunctive normal form (a conjunction of disjunctions of literals) where each clause has size two (e.g., $(a \vee b) \wedge (\neg a \vee c)$) is satisfiable. Solving 2-SAT requires only linear time (Aspvall, Plass, & Tarjan, 1979). The Boolean variables in our instance correspond to pairs of candidates, i.e., for each $a, b \in C$ we have a variable ab . The intended meaning of these variables is that $ab = \text{true}$ if and only if a is left of b on the axis. Now, for each vote V and triple $a, b, c \in C$, if $a \succ b$ and $c \succ b$ (a, b, c may form a v-valley), we add the clauses

$$(ba \vee cb) \quad \text{and} \quad (1)$$

$$(ab \vee bc) \quad (2)$$

to the 2-SAT instance. These clauses correspond the requirement that b must not be placed between a and c . Finally, we add for each pair of variables a, b the clauses

$$(ab \vee ba) \wedge (\neg ab \vee \neg ba), \tag{3}$$

corresponding to an exclusive or operator. Solving the 2-SAT instance either yields the information that the instance is not satisfiable or a true/false assignment to the variables. In the first case, the profile is not single-peaked (as shown in Lemma 3). In the second case, we obtain a relation $A = \{(a, b) : ab = \text{true}\}$ which is a total order and our desired axis (as shown in Lemma 4). Since the instance contains at most $\mathcal{O}(n \cdot m^3)$ clauses, we obtain the stated runtime.

Lemma 3. *If \mathcal{P} is single-peaked with respect to an axis A , then the corresponding 2-SAT instance is satisfiable.*

Proof. We construct a valid truth assignment as follows: If $a \triangleright b$ on A , then $ab = \text{true}$ and $ba = \text{false}$. Clauses of the form (1) and (2) are satisfied because no v-valleys occur on A ; Clauses of the form (3) are satisfied because v-valleys cannot occur (\mathcal{P} is single-peaked with respect to an axis A). □

Lemma 4. *The relation $A = \{(a, b) : ab = \text{true}\}$, as returned by the 2-SAT algorithm, is a total order and \mathcal{P} is single-peaked with respect to A .*

Proof. First, we want to show that A is a total order. Asymmetry and totality follow immediately from (3). Towards a contradiction assume that A is not transitive, i.e., there exist three candidates x, y, z such that $\{(x, y), (y, z), (z, x)\} \subseteq A$. Thus, $xy = yz = zx = \text{true}$ and $yx = zy = xz = \text{false}$. Let V be a total order contained in \mathcal{P} (there exists at least one). We distinguish three cases:

- The last ranked candidate of x, y, z in V is y : By (1), it has to hold that $(yx \vee zy)$, which is not the case.
- The last ranked candidate of x, y, z in V is x : By (2), it has to hold that $(yx \vee xz)$, which is not the case.
- The last ranked candidate of a, b, c in V is c : By (1), it has to hold that $(zy \vee xz)$, which is not the case.

Thus, A is transitive. It remains to show that \mathcal{P} is single-peaked with respect to A . Assume that there is a valley $a \succ b, c \succ b$ in some vote and it holds that $\{(a, b), (b, c)\} \subseteq A$. Due to this valley, our 2-SAT instance contains the clause $(ba \vee cb)$. Thus, $(b, a) \in A$ or $(c, b) \in A$ and $ba = \text{true}$ or $cb = \text{true}$. Hence, $ab = \text{false}$ or $bc = \text{false}$, which contradicts our assumption that $\{(a, b), (b, c)\} \subseteq A$. □

6. Other Single-Peaked Concepts for Weak Orders

We now turn to three other models of single-peaked preferences for weak orders—Black single-peaked and single-plateaued preferences, and at the end of the section, necessarily single-peaked preferences.

For weak orders, it is more common to view $a \sim b$ as indifference instead of missing information (i.e., incompleteness). From the algorithmic point of view, which we had taken so




Figure 4: Overview of Single-Peakedness Models for Weak Orders

far, this distinction is irrelevant. From a conceptional point of view, it clearly makes a difference, in particular for the underlying assumption of voters’ true preferences. It may even be that both incompleteness and indifference are simultaneously be present in a voter’s preference information; we however disregard this possibility.

Let us start with the definition of single-peaked preferences as introduced by Black, which applies to weak orders as well.

Definition 6. *A preference profile \mathcal{P} of weak orders is Black single-peaked with respect to an axis A if for every $V \in \mathcal{P}$, A can be split at the most-preferred candidate (peak) of V into two segments X and Y (one of which can be empty) such that V has strictly increasing preferences along X and strictly decreasing preferences along Y .*

In other words, for a weak preference order to be Black single-peaked it must have a single most-preferred candidate and can only contain indifference between at most two candidates at each position in the order. Otherwise the segments X and Y referred to in Definition 6 would not be *strictly* increasing/decreasing.

A slightly weaker restriction than Black single-peakedness for weak orders is the single-plateaued restriction (Black, 1958, Chapter 5), which extends Black single-peakedness to allow voters to state multiple most-preferred candidates (i.e., a single plateau instead of a single peak).⁵ The model of possibly single-peaked preferences considered in this paper generalizes both single-plateaued and Black single-peaked preferences. Examples and visualizations of these three restrictions can be found in Figure 4.

In this section, we have three main goals: In Section 6.1 we differentiate single-peakedness models with respect to the axiomatic properties they guarantee. In Sections 6.2 and 6.3 we show that the consecutive ones approach can also be used to solve the consistency problem for single-plateaued and for Black single-peaked preferences. Finally, in Section 6.4, we briefly consider the concept of *necessarily* single-peaked preferences.

6.1 Transitive Majority Relations and Condorcet Winners

Preference profiles that satisfy either the Black single-peaked or single-plateaued restriction have several desirable properties. One of them is that both Black single-peaked and single-plateaued

⁵ We note that single-plateaued preferences are occasionally referred to as single-peaked preferences, e.g., in (Fishburn, 1973, Chapter 9).

preferences guarantee transitive majority relations (Black, 1948, 1958). To be precise, we say that candidate a is preferred to b by a (strict) majority if $|V \in \mathcal{P} : a \succ b| > |V \in \mathcal{P} : b \succ a|$; in this case we write $a >_m b$, where $>_m$ denotes the majority relation. A transitive majority relation further guarantees the existence of one or more weak Condorcet winners. A candidate a is a weak Condorcet winner if there is no candidate b with $b >_m a$. Upon closer inspection, the Black single-peaked and the single-plateaued restriction show different characteristics; we refer the reader to Barberà (2007), who further discusses how the amounts of indifference permitted in these restrictions impact their properties.

The gain in generality for possibly single-peaked preferences comes at a price: This condition is no longer restrictive enough to guarantee a transitive majority relation. This can be seen by considering the following simple profile reproduced from Fishburn (1973, Table 9.1). Let $\mathcal{P} = (V_1, V_2, V_3, V_4, V_5)$ with

$$\begin{aligned} V_1 &: \langle b \succ a \succ c \rangle \\ V_2, V_3 &: \langle c \succ b \succ a \rangle \\ V_4, V_5 &: \langle a \succ b \sim c \rangle \end{aligned}$$

The corresponding majority relation satisfies $a >_m c >_m b >_m a$, and thus is clearly not transitive. Furthermore, no Condorcet winner exists. Note that \mathcal{P} is possibly single-peaked with respect to the axis $a \triangleright b \triangleright c$ and that this profile is comprised of top orders. Thus, there exist preference profiles of top orders that are possibly single-peaked and do not have a transitive majority relation nor a Condorcet winner.

6.2 Single-Plateaued Profiles and Their Recognition

Single-plateaued preferences are a much more restrictive model than possibly single-peaked preferences since they are essentially Black single-peaked except that each preference order can have multiple most-preferred candidates (Black, 1958, Chapter 5).

We first provide a characterization of single-plateaued preferences, similar to the v-valley based characterization of possibly single-peaked preferences (Lemma 2). Since a preference order must be strictly increasing and then strictly decreasing with respect to an axis (except for its most-preferred candidates), we can again use the v-valley substructure. Furthermore, we will need an additional substructure to prevent two candidates that are ranked indifferent in a voter's preference order from appearing on the same side of that voter's peak (plateau).

Definition 7. *A preference order V contains a plateau with respect to an axis A if there exist candidates $a, b \in C$ such that a and b are adjacent in A and V states $a \sim b$. A preference order V contains a nonpeak plateau with respect to an axis A if there exist candidates $a, b, c, \in C$ such that $a \triangleright b \triangleright c$ on A and V states either $a \succ b \sim c$ or $c \succ b \sim a$.*

We can now state the following analogue of Lemma 2.

Lemma 5. *Let $\mathcal{P} = (V_1, \dots, V_n)$ be a profile of weak orders. The following two statements are equivalent.*

- (i) *The profile \mathcal{P} is single-plateaued with respect to A .*
- (ii) *Every vote $V \in \mathcal{P}$ contains neither a v-valley nor a nonpeak plateau with respect to A .*

Proof. Let \mathcal{P} be a preference profile of weak orders, and let A be an axis.

If \mathcal{P} is single-plateaued with respect to A then for every preference order $V \in \mathcal{P}$, A can be split into segments X , Y , and Z such that V is strictly increasing along X , remaining constant along Y , and strictly decreasing along Z . Since v is only ever strictly decreasing along Z and Z is the rightmost segment of A , V cannot contain a v-valley with respect to A . For a nonpeak plateau to exist with respect to A there must exist candidates $a, b, c \in C$ such that $a \succ b \succ c$ in A and V states either $a \succ b \sim c$ or $c \succ b \sim a$.

We first consider the case where V states $a \succ b \sim c$. Since V strictly prefers a to b and a to c , and both b and c are to the right of a on the axis, we know that both b and c must be in segment Z . However, V is strictly decreasing along Z , so V cannot have a nonpeak plateau of this form.

We now consider the case where V states $c \succ b \sim a$. Since V strictly prefers c to a and c to b , and both a and b are to the left of c on the axis, we know that both a and b must be in segment X . However, V is strictly increasing along X , so V cannot have a nonpeak plateau of this form.

For the other direction, assume that no preference order $V \in \mathcal{P}$ contains a v-valley with respect to A and no preference order $V \in \mathcal{P}$ contains a nonpeak plateau with respect to A . Since no preference order $V \in \mathcal{P}$ contains a v-valley with respect to A , we know from Lemma 2 that V is possibly single-peaked with respect to A . Since we also know that no preference order $V \in \mathcal{P}$ contains a nonpeak plateau with respect to A it is easy to see that V is single-plateaued with respect to A . \square

We now extend the consecutive ones approach from Section 5.1 to single-plateaued profiles. Since the nonpeak plateau substructure is needed in addition to the v-valley substructure, we modify Construction 1 so that if a preference order contains a nonpeak plateau with respect to an axis A , then when the columns of its corresponding preference matrix are permuted according to A the matrix will contain a row with the sequence $\dots 1 \dots 0 \dots 1 \dots$. Further notice that if a preference order ranks three candidates indifferent to each other below its peak (plateau) that it will have a nonpeak plateau with respect to *every* possible axis. To handle this case in the extension to Construction 1 we need to ensure that its corresponding preference matrix will contain a row with the sequence $\dots 1 \dots 0 \dots 1 \dots$ for every permutation of its columns.

Construction 2. Let $\mathcal{P} = (V_1, \dots, V_n)$ be a profile of weak orders over candidate set C with $|C| = m$. For each V_i we initially construct an $m \times m$ binary matrix X_i as described in Construction 1.

The following extensions to Construction 1 ensure that if V_i has a nonpeak plateau with respect to an axis A then when the columns of X_i are permuted according to A it will not have the consecutive ones property.

If there exist three candidates $a, b, c \in C$ such that V_i states $a \sim b \sim c$ and they are not V_i 's most-preferred candidates, then output a matrix that does not have the consecutive ones property.

Otherwise, for each pair of candidates $a, b \in C$ such that (i) V_i ranks $a \sim b$ and (ii) a and b are not maximal in V_i , append three additional rows to the matrix X_i : to the column corresponding to a append $[0 \ 1 \ 1]'$, to the column corresponding to b append $[1 \ 1 \ 0]'$, to each column corresponding to a candidate strictly preferred to a and b append $[1 \ 1 \ 1]'$, and to each column corresponding to a remaining candidate append $[0 \ 0 \ 0]'$. After constructing a matrix X_i for each V_i , the matrices X_1, \dots, X_n are row-wise concatenated to yield a matrix X .

Construction 1 ensures that no preference order contains a v-valley and the extensions made in Construction 2 ensure that no preference order contains a nonpeak plateau. So the following theorem can be shown by an argument similar to the proof of Theorem 3. Now the presence of v-valleys *or* nonpeak plateaus, not just v-valleys, is equivalent to a row containing the sequence $\dots 1 \dots 0 \dots 1 \dots$.

Theorem 7. *The WEAK ORDER SINGLE-PLATEAUED CONSISTENCY problem can be solved in $\mathcal{O}(m^2 \cdot n)$ time.*

6.3 Black Single-Peaked Consistency

Recall that a preference order is Black single-peaked with respect to an axis A if it is strictly increasing to a single most-preferred candidate (peak) and then strictly decreasing with respect to A , this is consequently a special case of single-plateaued preferences. So we again use the v-valley substructure to obtain a characterization result, but as before we need an additional substructure. Even if no preference order has a v-valley with respect to A it may not be Black single-peaked because it is indifferent between two candidates on the same side of its peak or has more than one most-preferred candidate. We can handle the first condition just mentioned with the nonpeak plateau substructure used in the previous section, but the second condition requires us to forbid *any* kind of plateau.

Lemma 6. *Let $\mathcal{P} = (V_1, \dots, V_n)$ be a profile of weak orders. The following two statements are equivalent.*

- (i) *The profile \mathcal{P} is Black single-peaked with respect to A .*
- (ii) *Every vote $V \in \mathcal{P}$ contains neither a v-valley nor a plateau with respect to A .*

We omit the proof, as the argument is very similar to the proof of Lemma 5.

Now, we extend Construction 2 so that if a preference order contains a plateau with respect to an axis A , then when the columns of its preference matrix are permuted according to A the matrix will contain the sequence $\dots 1 \dots 0 \dots 1 \dots$. Since Construction 2 already ensures this for the case of nonpeak plateaus, our extended construction below only needs to add a condition for plateaus that contain the most-preferred candidates in a given preference order.

Construction 3. *Follow Construction 2 except add the following condition while constructing a preference matrix X_i for each preference order $V_i \in \mathcal{P}$. If there exist two candidates $a, b \in C$ such that V_i states $a \sim b$ and they are V_i 's most-preferred candidates, then output a matrix that does not have the consecutive ones property.*

When a preference order has a unique most-preferred candidate and is single-plateaued, it is clearly also Black single-peaked. Construction 3 ensures that no preference order contains more than one most-preferred candidate the same way that Construction 2 ensures that no preference order contains three or more candidates that are all ranked indifferent to each other and that are not the most preferred candidates, since this always results in a nonpeak plateau. So the following theorem can be shown by essentially the same argument as for Theorem 7, but using Lemma 6 instead of Lemma 5.

Theorem 8. *The WEAK ORDER BLACK SINGLE-PEAKED CONSISTENCY problem can be solved in $\mathcal{O}(m^2 \cdot n)$ time.*

6.4 Necessarily Single-Peaked Preferences

The central definition of this paper is that of possibly single-peaked preferences. As mentioned in Section 3, one could also consider *necessarily single-peaked* preferences, i.e., incomplete preferences for which every extension is single-peaked. We conclude this section by considering this much more restrictive variant. For weak orders, we can give a clear answer how these two concepts compare: necessarily single-peaked preferences can be characterized as a subclass of single-plateaued preferences.

Proposition 2. *A profile of weak orders is necessarily single-peaked with respect to an axis A if and only if it is single-plateaued with respect to A in such a way that plateaus have a size of at most two.*

Proof. Let \mathcal{P} , a profile of weak orders, be single-plateaued with respect to an axis A such that all plateaus have a size ≤ 2 . It is easy to see that for each $V \in \mathcal{P}$ that every extension of V to a total order cannot contain a v-valley with respect to A .

For the other direction, we consider two cases: (i) \mathcal{P} is not single-plateaued with respect to A and (ii) \mathcal{P} is single-plateaued with respect to axis A but a $V \in \mathcal{P}$ contains a plateau of size larger than two. In case (i), we use Lemma 5. There are two ways in which the single-plateaued property can be violated: a vote V contains a v-valley or a nonpeak plateau with respect to A . If V contains a v-valley, then none of its extensions is single-peaked with respect to A . If V contains a nonpeak plateau, then it is of the form $p \succ a \sim b$, where p is a maximal element of V , with $a \triangleright b \triangleright p$ on A (or reversed). Let V' be any extension of V for which $p \succ' a \succ' b$ holds. Then V' is not single-peaked with respect to A , and hence \mathcal{P} is not necessarily single-peaked with respect to A . Now, consider case (ii). Let $V \in \mathcal{P}$ with $a \sim b \sim c$, all of which are maximal elements in V , i.e., we have a plateau of size at least three. Without loss of generality assume that $a \triangleright b \triangleright c$ on A . Let V' be any extension of V that satisfies $a \succ' b$ and $c \succ' b$. Then V' contains a v-valley and thus \mathcal{P} is not necessarily single-peaked with respect to A . \square

Interestingly, this “plateau of size at most two” definition is exactly the way Arrow (1951) defined single-peakedness, introducing a slight deviation from Black’s original definition (as mentioned by Dummett & Farquharson, 1961).

7. Discussion and Future Work

We have analyzed the \mathcal{T} SINGLE-PEAKED CONSISTENCY problem for $\mathcal{T} \in \{\text{PARTIAL ORDER, LOCAL WEAK ORDER, WEAK ORDER, TOP ORDER, TOTAL ORDER}\}$. An overview of the results is displayed in Table 1. Despite the NP-completeness of PARTIAL ORDER SINGLE-PEAKED CONSISTENCY, we have found four fast algorithms for plausible application scenarios. The Guided Algorithm and the 2-SAT based algorithm each require a guiding vote. Such an order is likely to exist—at least implicitly—for large preference profiles. In the case that no guiding order exists, the consecutive ones approach and the Unguided Algorithm are applicable. Also, we have found that PARTIAL ORDER SINGLE-PEAKED CONSISTENCY is solvable in polynomial time if the axis is already part of the input. This covers a large spectrum of possible scenarios of real-world preferences where our algorithms could be applied.

\mathcal{T}	general	with guiding vote
PARTIAL ORDER	NP-c (Cor 1)	NP-c (Thm 2)
LOCAL WEAK O.	NP-c (Thm 1)	$\mathcal{O}(m^3 \cdot n)$ (Thm 6)
WEAK ORDER	$\mathcal{O}(m^2 \cdot n)$ (Thm 3)	$\mathcal{O}(m \cdot n)$ (Thm 4)
TOP ORDER	$\mathcal{O}(m^2 \cdot n)$ (Thm 5)	$\mathcal{O}(m \cdot n)$ (Thm 4)
TOTAL ORDER	$\mathcal{O}(m \cdot n)$ (Escoffier et al. (2008), also Thm 4)	n/a

Table 1: Overview of the complexity results and algorithms for \mathcal{T} SINGLE-PEAKED CONSISTENCY

7.1 Implementation

We implemented our consecutive ones-based algorithm for the recognition of possibly single-peaked preferences for weak orders (and the extensions to the construction for single-plateaued and Black single-peaked preferences). The code is written in Python 3.7 and is available at <https://github.com/zmf6921/incompletesp>. We use a SAT solver for the consecutive ones instances. (This is a much simpler and more reliable approach than implementing the available linear-time algorithms and still yields very fast runtimes.) We ran the algorithm on all currently available profiles of weak orders available on PrefLib.org (Mattei & Walsh, 2013); these are in total 379 toc-instances⁶. We found that none of these instances are possibly single-peaked. This highlights a disadvantage of the single-peaked domain: its poor robustness. Even a minor change in a single voter’s preferences can violate the single-peaked condition. And although the possibly single-peaked domain is a more general restriction, this drawback remains. Given this fragility, the result is not surprising but leads us to ask the question of how *close* these profiles are to being possibly single-peaked. We discuss this research direction further below. On a more positive note, even large instances with $> 1,000$ candidates and instances with $> 25,000$ distinct votes could be solved in seconds (the maximum running time was 20.2 seconds, the average was 0.8 seconds, the median 0.02 seconds, processed on a 4GHz Intel i7 with 16 GB of RAM).

7.2 Top Orders and Scoring Rules

We would like to mention one particular application of our algorithms concerning a specific class of voting rules, so-called scoring rules. Scoring rules are specified by a family of scoring vectors. For an election over m candidates, a scoring rule uses the corresponding m -candidate scoring vector $(\alpha_1, \dots, \alpha_m)$ to determine the winner(s). A vote $\langle c_1 \succ \dots \succ c_m \rangle$ gives α_1 points to c_1 , α_2 points to c_2 , etc. The candidate(s) with the highest score win. Often, scoring vectors of the type $(\alpha_1, \dots, \alpha_k, 0, \dots, 0)$ with $\alpha_1 \geq \dots \alpha_k > 0$ are considered. For example, the voting in the Eurovision Song Contest uses the scoring vector $(12, 9, 8, 7, 6, 5, 4, 3, 2, 1, 0, \dots, 0)$ (Ginsburgh & Noury, 2008). For such scoring vectors, top orders (with k ranked candidates) constitute *full* information for determining the winner. It is therefore debatable whether the input of such voting rules may be considered to be given as a profile of total orders, as total orders contain problem-irrelevant information. Brandt et al. (2015) study the constructive coalitional weighted manipulation problem for scoring rules in single-peaked elections. The authors consider the axis to be part of the input (for good reasons as explained in their paper). The computation of such

6. At PrefLib, profiles of weak orders are referred to as “tied-order complete (toc)”, c.f. <http://www.preflib.org/data/format.php#toc>.

an axis with existing algorithms is only possible if preferences are specified by total orders and thus contain problem-irrelevant information. If only relevant information is given, i.e., the input consists of top orders, algorithms such as those presented in this paper are required.

7.3 Future Research

Let us discuss some directions for future research. When discussing our implementation and the outcome of our PrefLib analysis, we mentioned the lack of robustness of (possibly) single-peaked profiles. This issue of robustness has been addressed by considering *nearly* single-peaked profiles, i.e., profiles that are close to being single-peaked according to some measure (Elkind, Faliszewski, & Slinko, 2012; Cornaz, Galand, & Spanjaard, 2012, 2013; Faliszewski, Hemaspaandra, & Hemaspaandra, 2014; Bredereck, Chen, & Woeginger, 2016; Erdélyi, Lackner, & Pfandler, 2017). This line of work is limited to profiles of total orders; an extension to possibly single-peaked profiles could further generalize the applicability of single-peakedness. The first step in this direction has been taken by Menon and Larson (2016), who consider profiles of top orders in which most voters have single-peaked preferences. Furthermore, work on matrices that almost satisfy the consecutive ones property (see Section 5.1) is highly relevant for this purpose; we refer the reader to a survey by Dom (2009) for a literature overview. It would be of great interest to see how close to being possibly single-peaked the profiles on PrefLib are.

Finally, work by Durand (2003), Lackner and Lackner (2017), and Chen and Finnendahl (2018) analyzed the number of single-peaked profiles of total orders. To extend this work to profiles of partial orders could shed light on how much generality can be gained by moving from total orders to partial orders. Even more interesting would be a comparison of single-peaked concepts for weak orders (see Section 6). Clearly, the possibly single-peaked definition is more general than, e.g., the single-plateaued definition—but a quantitative answer is missing.

Acknowledgements

This work is based on a conference paper by Lackner (2014) and a follow-up conference paper by Fitzsimmons (2015).

Martin Lackner was supported by the Austrian Science Foundation FWF, grant P31890. Zack Fitzsimmons was supported in part by the National Science Foundation under grant no. CCF-1101452 and by a National Science Foundation Graduate Research Fellowship under grant no. DGE-1102937. The research was done in part while Zack Fitzsimmons was on research leave at Rensselaer Polytechnic Institute. The collaboration between the two authors of this paper has been supported by COST Action IC1205 on Computational Social Choice.

We thank Edith Elkind, Piotr Faliszewski, Edith Hemaspaandra, David Narváez, Stanisław Radziszowski, and the anonymous referees for their helpful comments and suggestions. We thank Dominik Peters for helpful discussions and the permission to use the *TikZ* script that produced Figure 4.

Appendix A. Proof details

A.1 Correctness of the Guided Algorithm

We are going to prove that the Guided Algorithm (Algorithm 1) is correct. In the following, we write $a \succsim b$ to denote that either $a \succ b$ or $a \sim b$. We consider the Guided Algorithm at any given point during its runtime. In particular, we consider the sets A_L and A_R as constructed by the algorithm. Let V_g be the guiding vote in \mathcal{P} .

Lemma 7. *Let $k \in [n]$, $a_\ell = \max_k(A_L)$, $a_r = \max_k(A_R)$, and $c_j \in \{c_i, \dots, c_m\}$. It either holds that $c_j \succsim_k a_r$ or that $c_j \succsim_k a_\ell$.*

Proof. Without loss of generality we assume that a_ℓ was placed before a_r . Towards a contradiction assume that both $a_r \succ_k c_j$ and $a_\ell \succ_k c_j$. Let us consider the algorithm at the point when a_r was placed ($c_{i'} = a_r$ for some $i' < i$). We will show that (R1) is true and thus a_r could not have been placed on the right-hand side. Recall condition (R1):

$$c_i \succ_k \min_k(C_{>i}) \quad \text{and} \quad \max_k(A_L) \succ_k \min_k(C_{>i})$$

Since $c_{i'} = a_r \succ_k c_j \succ_k \min_k(C_{>i'})$ and $\max_k(A_L) = a_\ell \succ_k c_j \succ_k \min_k(C_{>i'})$, (R1) is true and $c_{i'}$ could not have been placed on the right side of the axis. \square

Lemma 8. *Let $k \in [n]$ and $c_j = \max_k(c_i, \dots, c_m)$. Furthermore, let $a, a' \in A_R$ such that candidate a has been placed on A_R before a' . Then it either holds that $a' \succsim_k c_j$ or it holds that $a' \succsim_k a$.*

Proof. We consider the algorithm at the point where a' was placed on the right-hand side, i.e., in A_R . At this point, condition (R2) had to be false; the fact that $a' \succsim_k c_j$ or $a' \succsim_k a$ holds is a direct consequence. \square

Proposition 3. *The Guided Algorithm (Algorithm 1) is correct, i.e., it outputs an axis if and only if the given preference profile is single-peaked, and, furthermore, \mathcal{P} is single-peaked with respect to any axis that is returned by the algorithm.*

Proof. We first show that if an axis A is found, the profile \mathcal{P} is single-peaked with respect to A . Note that the guiding vote V_g is single-peaked with respect to A , as the Guided Algorithm constructs A based on V_g . Towards a contradiction assume that there is a vote $V \in \mathcal{P}$ that is not single-peaked with respect to A . This means that there are three candidates a, b, c with order $a \triangleright b \triangleright c$ on A , $a \succ b$ and $c \succ b$. We have to distinguish six cases of how a, b, c are ordered by the guiding vote \succ_g :

- $a \prec_g b \prec_g c$ (a is placed first, then b , then c ; other candidates in arbitrary order): Let us consider the algorithm at the point when b is being placed, i.e., $b = c_i$, and when the conditions for vote V are being checked. It holds that either $a \in A_L$ or $a \in A_R$. If $a \in A_L$, observe that condition (L2) is satisfied since $a \succ b$ and $c \succ b$. Consequently, b has to be placed on the right side (*left = false*). Then it holds that $a \triangleright c \triangleright b$ on the axis generated by the algorithm which contradicts our assumption that $a \triangleright b \triangleright c$ holds. In the case that $a \in A_R$ condition (R2) is satisfied. This leads to a contradiction by the same argument.
- $c \prec_g b \prec_g a$: This case is analogous.

- $a \prec_g c \prec_g b$: Now we consider the point where c is being placed, i.e., $c = c_i$, and when the conditions for vote V are being checked. It holds that either $a \in A_L$ or $a \in A_R$. If $a \in A_L$, observe that (R1) is satisfied and hence c has to be placed on the left side (*right = false*). Then it holds that $a \triangleright c \triangleright b$ on the axis generated by the algorithm which contradicts our assumption that $a \triangleright b \triangleright c$ holds. In the case that $a \in A_R$, condition (L1) is satisfied and we obtain a contradiction by the same argument.
- $c \prec_g a \prec_g b$: This case is analogous to the previous one.
- $b \prec_g c \prec_g a$ or $b \prec_g a \prec_g c$: Since b is placed on A before a and c , the resulting axis cannot be $a \triangleright b \triangleright c$, a contradiction.

For the other direction, let us show that if the algorithm returns `not_single_peaked`, then the profile \mathcal{P} is not single-peaked. First, let us observe under what conditions the algorithm returns `not_single_peaked`. There are four cases: Either conditions (R1) and (L1), (R1) and (L2), (R2) and (L1) or conditions (R2) and (L2) hold. These pairs of conditions may either hold for the same vote or for two distinct votes; we denote these two votes V and V' although it might be that these two are the same.

- While placing c_i , condition (L1) holds for some vote V and condition (R1) holds for some vote V' :

We have the following five candidates in these conditions: $a_r = \max_V(A_R)$, c_i , $c_j = \min_V(C_{>i})$ in condition (L1) and $a_\ell = \max_{V'}(A_L)$, c_i , $c'_j = \min_{V'}(C_{>i})$ in condition (R1). In Figure 5, the known information about the votes V and V' is shown. Since condition (R1) and (L1) are symmetrical, we can assume without loss of generality that a_ℓ is placed before a_r . Thus, we can assume that the guiding vote is $\{c_j, c'_j\} \succ_g c_i \succ_g a_r \succ_g a_\ell$, as shown in the figure (the order of c_j and c'_j is not relevant and thus can be arbitrary).

There are four types of axes possible that are compatible with the guiding vote V_g . (The order of candidates in sets is arbitrary.)

- $\langle a_\ell \triangleright c_i \triangleright \{c_j, c'_j\} \triangleright a_r \rangle$: Vote V is not single-peaked with respect to any axis of this type (or their reverse).
- $\langle a_\ell \triangleright \{c_j, c'_j\} \triangleright c_i \triangleright a_r \rangle$: Vote V' is not single-peaked with respect to any axis of this type (or their reverse).
- $\langle a_\ell \triangleright a_r \triangleright \{c_j, c'_j\} \triangleright c_i \rangle$: Vote V is not single-peaked with respect to any axis of this type (or their reverse).
- $\langle a_\ell \triangleright a_r \triangleright c_i \triangleright \{c_j, c'_j\} \rangle$: Consider vote V' and Lemma 7. Since $a_\ell \succ' c'_j$ it has to hold that $c'_j \succ' a_r$. Consequently, $a_\ell \succ' a_r$ and $c_i \succ' a_r$. Thus, the candidates a_ℓ, a_r, c_i form a v-valley for vote V' .

Since these are all possible axes, we can conclude that the profile is not single-peaked.

- While placing c_i , condition (L2) holds for some vote V and condition (R2) holds for some vote V' :

This case is similar to the previous one. In particular, we use the same candidate variables. In Figure 6, the known information about the votes V and V' is shown.




Figure 5: Conditions (L1) and (R1)

Figure 6: Conditions (L2) and (R2)

There are four types of axes possible that are compatible with V_g . (The order of candidates in sets is arbitrary.)

- $\langle a_\ell \triangleright c_i \triangleright \{c_j, c'_j\} \triangleright a_r \rangle$: Vote V is not single-peaked with respect to any axis of this type (or their reverse).
 - $\langle a_\ell \triangleright \{c_j, c'_j\} \triangleright c_i \triangleright a_r \rangle$: Vote V' is not single-peaked with respect to any axis of this type (or their reverse).
 - $\langle a_\ell \triangleright a_r \triangleright \{c_j, c'_j\} \triangleright c_i \rangle$: Consider vote V and Lemma 7. Since $a_\ell \succ c_i$ it has to hold that $c_i \succsim a_r$. Thus, the candidates a_ℓ, a_r, c_j form a v-valley for vote V .
 - $\langle a_\ell \triangleright a_r \triangleright c_i \triangleright \{c_j, c'_j\} \rangle$: Vote V' is not single-peaked with respect to any axis of this type (or their reverse).
- While placing c_i , condition (L1) holds for vote V and condition (R2) holds for vote V' : We have the following five candidates in these conditions: $a_r = \max_V(A_R)$, $c_i = \min_V(C_{>i})$ in condition (L1) and $a'_r = \max_{V'}(A_R)$, $c_i, c'_j = \min_{V'}(C_{>i})$ in condition (R2). In Figure 7, the known information about the votes V and V' is shown. In the following arguments it is irrelevant which of c_j and c'_j is placed first. However, for a_r and a'_r this is relevant. We will consider both possibilities (cases 1 and 2 in Figure 7). There are four types of axes possible that are compatible with either of these two guiding votes. (The order of candidates in sets is arbitrary.)

- $\langle \{a_r, a'_r\} \triangleright \{c_j, c'_j\} \triangleright c_i \rangle$: Vote V is not single-peaked with respect to any axis of this type (or their reverse).
- $\langle \{a_r, a'_r\} \triangleright c_i \triangleright \{c_j, c'_j\} \rangle$: Vote V' is not single-peaked with respect to any axis of this type (or their reverse).
- $\langle a_r \triangleright \{c_j, c'_j\} \triangleright c_i \triangleright a'_r \rangle$: Both V and V' are not single-peaked with respect to any axis of this type (or their reverse).
- $\langle a_r \triangleright c_i \triangleright \{c_j, c'_j\} \triangleright a'_r \rangle$: Here we have to distinguish whether a_r or a'_r is placed first (case 1 and 2 in Figure 7).

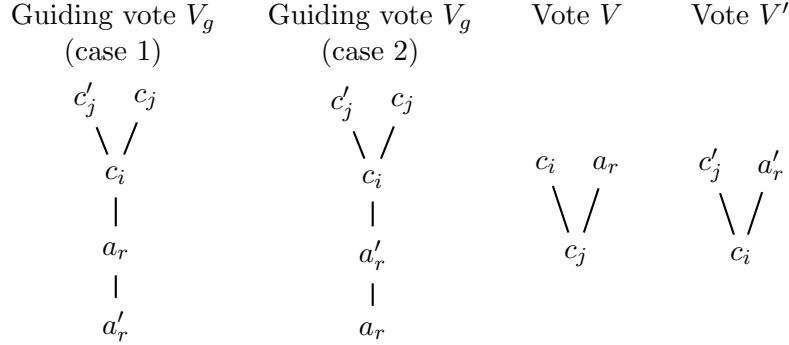


Figure 7: Condition (L1) and Condition (R2)

Case 1 ($a'_r \prec_g a_r$): Lemma 8 applied to V' implies that either (i) $a_r \succ' c_i$ or (ii) $a_r \succ' a'_r$. In both cases we encounter a valley with respect to vote V' for the candidates $\{a_r, a'_r, c_i\}$.

Case 2 ($a_r \prec_g a'_r$): This case is similar to the previous. Lemma 8 applied to V implies that either (i) $a_r \succ c_j$ or (ii) $a'_r \succ a_r$. In both cases we encounter a valley with respect to vote V for the candidates $\{a'_r, a_r, c_j\}$.

We see that in both cases either V or V' is not single-peaked with respect to any axis of this type (or their reverse).

- While placing c_i , condition (L2) holds for some vote and condition (R1) holds for some vote:

This can be shown analogously to the previous case since (R1) and (L1) are symmetrical as well as (R2) and (L2), see Figure 3.

We have shown that if the algorithm returns `not_single_peaked` then the profile \mathcal{P} is indeed not single-peaked. \square

A.2 Runtime and Correctness for the Unguided Algorithm

Theorem 5. *The TOP ORDER SINGLE-PEAKED CONSISTENCY problem can be solved in $\mathcal{O}(m^2 \cdot n)$ time.*

Proof. Let us first discuss the runtime of Algorithm 2. Some functions have to be precomputed to achieve the $\mathcal{O}(m^2 \cdot n)$ runtime, in particular the function `IntersectingVote`. The function `IntersectingVote(A)` returns a vote V with $A \cap V \neq \emptyset$ and $V \setminus A \neq \emptyset$. We show that it suffices to compute a list of $2m$ votes to answer `IntersectingVote` function calls in constant time. Let us first make the following observation: Let $c \in C$. Consider the set of votes for which the sets $\{c' \in C \mid c' \succ c\}$ are maximal (with respect to \subseteq). If we consider a single-peaked axis, then candidates in such a set have to form a contiguous subsequence either directly left or directly right of c . Since these sets are maximal, only two of them can exist (assuming single-peakedness). Consequently, we compute these maximal sets for each candidate. If three or more exist for one candidate, we can terminate the algorithm already at this point. Also, if two maximal sets have a non-empty intersection, the algorithm terminates. (The candidates in the intersection would

have to lie left and right of c). Otherwise we store the (at most) two corresponding votes for each candidate.

Let $A = \langle a_1 \triangleright \dots \triangleright a_{i-1} \triangleright a_i \rangle$, i.e., a_i is the rightmost candidate in the incomplete axis A . The function call `IntersectingVote(A)` can now be answered by considering the one or two maximal votes for a_i . The function simply returns the vote where a_{i-1} is not ranked higher than a_i . (It might be that both votes do not rank a_{i-1} higher than a_i . In this case A cannot be extended to a single-peaked axis, but this is going to be detected by algorithm. Any of the two axes can be returned.)

It remains to observe that finding the (at most two) maximal votes for a candidate c requires $\mathcal{O}(m \cdot n)$ time. This has to be done for every candidate and consequently this preprocessing requires $\mathcal{O}(m^2 \cdot n)$ time.

We can now analyze the runtime of the algorithm. The main loop (line 1) iterates over all m candidates. The loop in line 4 iterates over every vote at most once. Consequently, the \oplus operator is applied at most n times. Since $A \oplus V$ can be computed in $\mathcal{O}(m)$ time, the lines 4 to 7 have a total runtime of $\mathcal{O}(m^2 \cdot n)$.

It remains to determine the runtime of the lines 8 to 20. Due to the preprocessing of the `IntersectingVote` procedure we can obtain V in constant time. The profile \mathcal{P}' can be generated in $\mathcal{O}(|C'| \cdot n)$ time. Applying the Guided Algorithm requires $\mathcal{O}(|C'| \cdot n)$ time as well (Theorem 4). Observe that after applying the Guided Algorithm the candidates in C' are placed on the axis. Consequently, the Guided Algorithm is always applied to a disjoint set of candidates (except maybe a_i). Hence for a fixed $c_{\text{start}} \in C$, the total runtime of the Guided Algorithm is $\mathcal{O}(m \cdot n)$. Taking the loop in line 1 into account, we obtain a total runtime of $\mathcal{O}(m^2 \cdot n)$.

Let us now show that the Unguided Algorithm (Algorithm 2) is correct, i.e., it outputs an axis if and only if the given preference profile is single-peaked and, furthermore, \mathcal{P} is single-peaked with respect to any axis that is returned by the algorithm. If the algorithm outputs an axis, it is certainly single-peaked since this is tested for every vote in line 5. By the same argument, one can conclude that if the profile is not single-peaked, the algorithm returns `not_single_peaked`.

It remains to prove that the algorithm always returns a valid axis in case of a single-peaked profile. Let us consider the algorithm at the time when c_{start} is the leftmost candidate of a valid axis. We show that the algorithm will find a complete axis with c_{start} as leftmost candidate. First, observe that for every i in line 3 either there is a vote V with a_i as its top-ranked candidate or the condition in line 8 is true. In the first case, it is easy to verify that the \oplus operator adds candidates to the axis in the only possible way. In the second case, an intersecting vote is found as guaranteed by the connectedness condition. Then the Guided Algorithm is applied and the axis is extended by the candidates in C' . It remains to verify that the resulting axis A is single-peaked for all votes with a non-empty intersection with C' . This is guaranteed by the x element, which ensures that candidates outside of $A \cup C'$ are taken into account. \square

References

- Arrow, K. (1951). *Social Choice and Individual Values*. John Wiley and Sons.
- Aspvall, B., Plass, M., & Tarjan, R. (1979). A linear-time algorithm for testing the truth of certain quantified Boolean formulas. *Information Processing Letters*, 8(3), 121–123.
- Aziz, H. (2014). Testing top monotonicity. Tech. rep. arXiv:1403.7625 [cs.GT], arXiv.org.

- Ballester, M. A., & Haeringer, G. (2011). A characterization of the single-peaked domain. *Social Choice and Welfare*, 36(2), 305–322.
- Barberà, S. (2007). Indifferences and domain restrictions. *Analyse & Kritik*, 29(2), 146–162.
- Barberà, S., & Moreno, B. (2011). Top monotonicity: A common root for single peakedness, single crossing and the median voter result. *Games and Economic Behavior*, 73(2), 345–359.
- Bartholdi, III, J., Tovey, C., & Trick, M. (1989). Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare*, 6(2), 157–165.
- Bartholdi, III, J., & Trick, M. (1986). Stable matching with preferences derived from a psychological model. *Operations Research Letters*, 5(4), 165–169.
- Baumeister, D., Faliszewski, P., Lang, J., & Rothe, J. (2012). Campaigns for lazy voters: Truncated ballots. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, pp. 577–584. IFAAMAS.
- Baumeister, D., & Rothe, J. (2012). Taking the final step to a full dichotomy of the possible winner problem in pure scoring rules. *Information Processing Letters*, 112(5), 186 – 190.
- Betzler, N., Slinko, A., & Uhlmann, J. (2013). On the computation of fully proportional representation. *Journal of Artificial Intelligence Research*, 47, 475–519.
- Betzler, N., Brederbeck, R., & Niedermeier, R. (2014). Theoretical and empirical evaluation of data reduction for exact Kemeny rank aggregation. *Autonomous Agents and Multi-Agent Systems*, 28(5), 721–748.
- Betzler, N., & Dorn, B. (2010). Towards a dichotomy for the possible winner problem in elections based on scoring rules. *Journal of Computer and System Sciences*, 76(8), 812–836.
- Betzler, N., Hemmann, S., & Niedermeier, R. (2009). A multivariate complexity analysis of determining possible winners given incomplete votes. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI 2009)*, pp. 53–58.
- Black, D. (1948). On the rationale of group decision-making. *Journal of Political Economy*, 56(1), 23–34.
- Black, D. (1958). *The Theory of Committees and Elections*. Cambridge University Press.
- Booth, K., & Lueker, G. (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.
- Boutilier, C., & Rosenschein, J. (2016). Incomplete information and communication in voting. In Brandt, F., Conitzer, V., Endriss, U., Lang, J., & Procaccia, A. (Eds.), *Handbook of Computational Social Choice*. Cambridge University Press.
- Brandt, F., Brill, M., Hemaspaandra, E., & Hemaspaandra, L. (2015). Bypassing combinatorial protections: Polynomial-time algorithms for single-peaked electorates. *Journal of Artificial Intelligence Research*, 53, 439–496.
- Bredereck, R., Chen, J., & Woeginger, G. (2013). A characterization of the single-crossing domain. *Social Choice and Welfare*, 41(4), 989–998.
- Bredereck, R., Chen, J., & Woeginger, G. J. (2016). Are there any nicely structured preference profiles nearby?. *Mathematical Social Sciences*, 79, 61–73.

- Cantala, D. (2004). Choosing the level of a public good when agents have an outside option. *Social Choice and Welfare*, 22(3), 491–514.
- Chen, J., & Finnendahl, U. (2018). On the number of single-peaked narcissistic or single-crossing narcissistic preference profiles. *Discrete Mathematics*, 341(5), 1225–1236.
- Chen, J., Pruhs, K. R., & Woeginger, G. J. (2017). The one-dimensional euclidean domain: finitely many obstructions are not enough. *Social Choice and Welfare*, 48(2), 409–432.
- Conitzer, V. (2009). Eliciting single-peaked preferences using comparison queries. *Journal of Artificial Intelligence Research*, 35, 161–191.
- Cornaz, D., Galand, L., & Spanjaard, O. (2013). Kemeny elections with bounded single-peaked or single-crossing width. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI 2013)*, pp. 76–82. IJCAI/AAAI Press.
- Cornaz, D., Galand, L., & Spanjaard, O. (2012). Bounded single-peaked width and proportional representation. In *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI 2012)*, pp. 270–275. IOS Press.
- De Gemmis, M., Iaquinta, L., Lops, P., Musto, C., Narducci, F., & Semeraro, G. (2010). Learning preference models in recommender systems. In *Preference Learning*, pp. 387–407. Springer.
- Dery, L. N., Kalech, M., Rokach, L., & Shapira, B. (2014). Reaching a joint decision with minimal elicitation of voter preferences. *Information Sciences*, 278, 466–487.
- Doignon, J.-P., & Falmagne, J.-C. (1994). A polynomial time algorithm for unidimensional unfolding representations. *Journal of Algorithms*, 16(2), 218–233.
- Dom, M. (2009). Algorithmic aspects of the consecutive-ones property. In *Bulletin of the European Association for Theoretical Computer Science*, Vol. 98, pp. 27–59. EATCS.
- Dummett, M., & Farquharson, R. (1961). Stability in voting. *Econometrica*, 29(1), 33–43.
- Durand, S. (2003). Finding sharper distinctions for conditions of transitivity of the majority method. *Discrete applied mathematics*, 131(3), 577–595.
- Dwork, C., Kumar, R., Naor, M., & Sivakumar, D. (2001). Rank aggregation methods for the web. In *Proceedings of the 10th International World Wide Web Conference (WWW 2001)*, pp. 613–622. ACM Press.
- Elkind, E., Faliszewski, P., Lackner, M., & Obraztsova, S. (2015). The complexity of recognizing incomplete single-crossing preferences. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI 2015)*, pp. 865–871. AAAI Press.
- Elkind, E., Faliszewski, P., & Slinko, A. (2012). Clone structures in voters’ preferences. In *Proceedings of the 13th ACM conference on electronic commerce (EC 2012)*, pp. 496–513. ACM.
- Elkind, E., & Ismaili, A. (2015). OWA-based extensions of the chamberlin–courant rule. In *Proceedings of the 4th International Conference on Algorithmic Decision Theory (ADT 2015)*, pp. 486–502. Springer.
- Elkind, E., & Lackner, M. (2015). Structure in dichotomous preferences. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI 2015)*, pp. 2019–2025.

- Elkind, E., Lackner, M., & Peters, D. (2016). Preference restrictions in computational social choice: Recent progress. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI 2016)*, pp. 4062–4065. IJCAI/AAAI Press.
- Elkind, E., Lackner, M., & Peters, D. (2017). Structured preferences. In Endriss, U. (Ed.), *Trends in Computational Social Choice*, chap. 10, pp. 187–207. AI Access.
- Erdélyi, G., Lackner, M., & Pfandler, A. (2017). Computational aspects of nearly single-peaked electorates. *Journal of Artificial Intelligence Research*, 58, 297–337.
- Escoffier, B., Lang, J., & Öztürk, M. (2008). Single-peaked consistency and its complexity. In *Proceedings of the 18th European Conference on Artificial Intelligence (ECAI 2008)*, pp. 366–370.
- Fagin, R., Kumar, R., Mahdian, M., Sivakumar, D., & Vee, E. (2006). Comparing partial rankings. *SIAM Journal on Discrete Mathematics*, 20(3), 628–648.
- Fagin, R., Kumar, R., & Sivakumar, D. (2003). Comparing top k lists. *SIAM Journal on Discrete Mathematics*, 17(1), 134–160.
- Faliszewski, P., Hemaspaandra, E., & Hemaspaandra, L. (2014). The complexity of manipulative attacks in nearly single-peaked electorates. *Artificial Intelligence*, 207, 69–99.
- Faliszewski, P., Hemaspaandra, E., Hemaspaandra, L., & Rothe, J. (2011). The shield that never was: Societies with single-peaked preferences are more open to manipulation and control. *Information and Computation*, 209(2), 89–107.
- Fishburn, P. (1973). *The Theory of Social Choice*. Princeton University Press.
- Fitzsimmons, Z., & Hemaspaandra, E. (2015). Complexity of manipulative actions when voting with ties. In *Proceedings of the 4th International Conference on Algorithmic Decision Theory (ADT 2015)*, pp. 103–119. Springer.
- Fitzsimmons, Z. (2015). Single-peaked consistency for weak orders is easy. In *Proceedings of the 15th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2015)*, pp. 127–140.
- Fitzsimmons, Z., & Hemaspaandra, E. (2016). Modeling single-peakedness for votes with ties. In *Proceedings of the 8th European Starting AI Researcher Symposium (STAIRS 2016)*, pp. 63–74. IOS Press.
- Fulkerson, D., & Gross, G. (1965). Incidence matrices and interval graphs. *Pacific Journal of Mathematics*, 15(5), 835–855.
- Gaertner, W. (2002). Domain restrictions. In Arrow, K., Sen, A., & Suzumura, K. (Eds.), *Handbook of Social Choice and Welfare, Volume 1*, chap. 3, pp. 131–170. Elsevier.
- Garey, M., & Johnson, D. (1979). *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company.
- Gibbard, A. (1973). Manipulation of voting schemes. *Econometrica*, 41(4), 587–601.
- Ginsburgh, V., & Noury, A. (2008). The Eurovision song contest. Is voting political or cultural?. *European Journal of Political Economy*, 24(1), 41–52.
- Habib, M., McConnell, R., Paul, C., & Viennot, L. (2000). Lex-BFS and partition refinement, with applications to transitive orientation, interval graph recognition and consecutive ones testing. *Theoretical Computer Science*, 234(1), 59–84.

- Hemaspaandra, E., Hemaspaandra, L., & Rothe, J. (2016). The complexity of manipulative actions in single-peaked societies. In Rothe, J. (Ed.), *Economics and Computation*, pp. 327–360. Springer.
- Hemaspaandra, E., Hemaspaandra, L. A., & Rothe, J. (1997). Exact analysis of Dodgson elections: Lewis Carroll’s 1876 voting system is complete for parallel access to NP. *Journal of the ACM*, 44(6), 806–825.
- Hemaspaandra, E., Spakowski, H., & Vogel, J. (2005). The complexity of Kemeny elections. *Theoretical Computer Science*, 349(3), 382–391.
- Hsu, W.-L. (2002). A simple test for the consecutive ones property. *Journal of Algorithms*, 43(1), 1–16.
- Konczak, K., & Lang, J. (2005). Voting procedures with incomplete preferences. In *Proceedings of the 1st Multidisciplinary Workshop on Advances in Preference Handling (MPref 2005)*, pp. 124–129.
- Lackner, M.-L., & Lackner, M. (2017). On the likelihood of single-peaked preferences. *Social Choice and Welfare*, 48(4), 717–745.
- Lackner, M. (2014). Incomplete preferences in single-peaked electorates. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI 2014)*, pp. 742–748. AAAI Press.
- LeGrand, R., Markakis, E., & Mehta, A. (2007). Some results on approximating the minimax solution in approval voting. In *AAMAS’07*, pp. 1185–1187.
- Liu, H., & Guo, J. (2016). Parameterized complexity of winner determination in minimax committee elections. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pp. 341–349. International Foundation for Autonomous Agents and Multiagent Systems.
- Magiera, K., & Faliszewski, P. (2019). Recognizing top-monotonic preference profiles in polynomial time. *Journal of Artificial Intelligence Research*, 66, 57–84.
- Masthoff, J. (2015). Group recommender systems: Aggregation, satisfaction and group attributes. In Francesco Ricci, Lior Rokach, B. S. (Ed.), *Recommender Systems Handbook*. Springer.
- Mattei, N., & Walsh, T. (2013). Preflib: A library of preference data. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT 2013)*. Springer.
- McConnell, R. M. (2004). A certifying algorithm for the consecutive-ones property. In *Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 768–777. Society for Industrial and Applied Mathematics.
- Meidanis, J., Porto, O., & Telles, G. P. (1998). On the consecutive ones property. *Discrete Applied Mathematics*, 88(1), 325–354.
- Menon, V., & Larson, K. (2017). Computational aspects of strategic behaviour in elections with top-truncated ballots. *Autonomous Agents and Multiagent Systems*, 31(6), 1506–1547.
- Menon, V., & Larson, K. (2016). Reinstating combinatorial protections for manipulation and bribery in single-peaked and nearly single-peaked electorates. In Schuurmans, D., & Wellman, M. P. (Eds.), *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA.*, pp. 565–571. AAAI Press.

- Moulin, H. (1991). *Axioms of Cooperative Decision Making*. Cambridge University Press.
- Moulin, H. (1984). Generalized Condorcet-winners for single peaked and single-plateau preferences. *Social Choice and Welfare*, 1(2), 127–147.
- Opatrny, J. (1979). Total ordering problem. *SIAM Journal on Computing*, 8(1), 111–114.
- Peters, D. (2017). Recognising multidimensional Euclidean preferences. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI 2017)*, pp. 642–648. AAAI Press.
- Peters, D. (2018). Single-peakedness and total unimodularity: New polynomial-time algorithms for multi-winner elections. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI 2018)*, pp. 1169–1176. AAAI Press.
- Peters, D., & Lackner, M. (2017). Preferences single-peaked on a circle. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI 2017)*, pp. 649–655. AAAI Press.
- Pini, M., Rossi, F., Venable, K., & Walsh, T. (2011). Incompleteness and incomparability in preference aggregation: Complexity results. *Artificial Intelligence*, 175(7-8), 1272–1289.
- Raffinot, M. (2011). Consecutive ones property testing: Cut or swap. In *Conference on Computability in Europe*, pp. 239–249. Springer.
- Rothe, J., Spakowski, H., & Vogel, J. (2003). Exact complexity of the winner problem for Young elections. *Theory of Computing Systems*, 36(4), 375–386.
- Satterthwaite, M. (1975). Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2), 187–217.
- Schafer, J. B., Frankowski, D., Herlocker, J., & Sen, S. (2007). Collaborative filtering recommender systems. In *The adaptive web*, pp. 291–324. Springer.
- Walsh, T. (2007). Uncertainty in preference elicitation and aggregation. In *Proceedings of the 22nd National Conference on Artificial Intelligence (AAAI 2007)*, pp. 3–8. AAAI Press.
- Xia, L., & Conitzer, V. (2011). Determining possible and necessary winners under common voting rules given partial orders. *Journal of Artificial Intelligence Research*, 41(2), 25–67.