Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)

Structure in Dichotomous Preferences

Edith Elkind
University of Oxford, UK
elkind@cs.ox.ac.uk

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

Abstract

Many hard computational social choice problems are known to become tractable when voters' preferences belong to a restricted domain, such as those of single-peaked or single-crossing preferences. However, to date, all algorithmic results of this type have been obtained for the setting where each voter's preference list is a total order of candidates. The goal of this paper is to extend this line of research to the setting where voters' preferences are *dichotomous*, i.e., each voter approves a subset of candidates and disapproves the remaining candidates. We propose several analogues of the notions of single-peaked and single-crossing preferences for dichotomous profiles and investigate the relationships among them. We then demonstrate that for some of these notions the respective restricted domains admit efficient algorithms for computationally hard approval-based multi-winner rules.

1 Introduction

Preference aggregation is a fundamental problem in social choice, which has recently received a considerable amount of attention from the AI community. In particular, an important research question in *computational social choice* is the complexity of computing the output of various preference aggregation procedures. While for most common single-winner rules winner determination is easy, many attractive rules that output a committee (a fixed-size set of winners) or a ranking of the candidates are known to be computationally hard.

There are several ways to circumvent these hardness results, such as using approximate and parameterized algorithms. These standard algorithmic approaches are complemented by an active stream of research that analyzes the computational complexity of voting rules on *restricted preference domains*, such as the classic domains of single-peaked [Black, 1958] or single-crossing [Mirrlees, 1971] preferences. This research direction was popularized by Walsh [2007] and Faliszewski *et al.* [2011], and has lead to a number of efficient algorithms for winner determination under prominent voting rules as well as for manipulation and control, which can be used when voters' preferences belong to one of these restricted domains [Walsh, 2007; Faliszewski *et al.*, 2011; Brandt *et al.*, 2010; Faliszewski *et al.*, 2014; Betzler *et al.*, 2013; Skowron *et al.*, 2015b; Magiera and Faliszewski, 2014].

To the best of our knowledge, this line of work only considers settings where voters' preferences are given by total orders over the set of candidates; indeed, this is perhaps the most widely studied setting in the area of computational social choice. However, computationally complex preference aggregation problems may also arise when voters' preferences are *dichotomous*, i.e., each voter approves a subset of the candidates and disapproves the remaining candidates. Committee selection rules for voters with dichotomous preferences, or *approval-based* rules, have recently attracted some attention from the computational social choice community, and for two prominent such rules (specifically, Proportional Approval Voting (PAV) [Kilgour and Marshall, 2012] and Maximin Approval Voting (MAV) [Brams *et al.*, 2007]) computing the winning committee is known to be NP-hard [Aziz *et al.*, 2014; LeGrand *et al.*, 2007]. It is therefore natural to ask if one could identify a suitable analogue of single-peaked/single-crossing preferences for the the dichotomous setting, and design efficient algorithms for approval-based rules over such restricted dichotomous preference domains.

To address this challenge, in this paper we propose and explore a number of domain restrictions for dichotomous preferences that build on the same intuition as the concepts of single-peakedness and single-crossingness. Some of our restricted domains are defined by embedding voters or candidates into the real line, and requiring that the voters' preferences over the candidates "respect" this embedding; others are obtained by viewing dichotomous preferences as weak orders and requiring them to admit a refinement that has a desirable structural property. Surprisingly, these approaches lead to a large number of concepts that are pairwise non-equivalent and capture different aspects of our intuition about what it means for preferences to be "one-dimensional". We analyze the relationships among these restricted preference domains, (see Figure 3), and discuss the complexity of detecting whether a given dichotomous profile belongs to one of these domains. We then show that these domains are useful from the perspective of algorithm design, by providing polynomial-time and FPT algorithms for PAV and MAV under some of these domain restrictions. The full version of this paper is available on arXiv [Elkind and Lackner, 2015].

2 Basic Definitions

Let $C = \{c_1, \dots, c_m\}$ be a finite set of candidates. A *partial order* $\succ$ over C is a reflexive, antisymmetric and transitive binary relation on C; a partial order $\succ$ is said to be *total* if for each $c, d \in C$ we have $c \succ d$ or $d \succ c$. We say that a partial order $\succ$ over C is a *dichotomous weak order* if C can be partitioned into two disjoint sets C^+ and C^- (one of which may be empty) so that $c \succ d$ for each $c \in C^+, d \in C^-$ and the candidates within C^+ and C^- are incomparable under $\succ$.

An *approval vote* on C is an arbitrary subset of C. We say that an approval vote v is *trivial* if $v = \emptyset$ or $v = C$. A *dichotomous profile* $\mathcal{P} = (v_1, \dots, v_n)$ is a list of n approval votes; we will refer to v_i as the vote of voter i. We write $\overline{v_i} = C \setminus v_i$. We associate an approval vote v_i with the dichotomous weak order $\succ_{v_i}$ that satisfies $c \succ_{v_i} d$ if and only if $c \in v_i$, $d \in \overline{v_i}$. Note that $v_i = \emptyset$ and $v_i = C$ correspond to the same dichotomous weak order, namely the empty one.

A partial order $\succ'$ over C is a *refinement* of a partial order $\succ$ over C if for every $c, d \in C$ it holds that $c \succ d$ implies $c \succ' d$. A profile $\mathcal{P}' = (\succ_1, \dots, \succ_n)$ of total orders is a *refinement* of a dichotomous profile $\mathcal{P} = (v_1, \dots, v_n)$ if $\succ_i$ is a refinement of $\succ_{v_i}$ for each $i = 1, \dots, n$.

Let $\lhd$ be a total order over C. A total order $\succ$ over C is said to be *single-peaked with respect to* $\lhd$ if for any triple of candidates $a, b, c \in C$ with $a \lhd b \lhd c$ or $c \lhd b \lhd a$ it holds that $a \succ b$ implies $b \succ c$. A profile $\mathcal{P}$ of total orders over C is said to be *single-peaked* if there exists a total order $\lhd$ over C such that all orders in $\mathcal{P}$ are single-peaked with respect to $\lhd$.

A profile $\mathcal{P} = (\succ_1, \dots, \succ_n)$ of total orders over C is said to be *single-crossing with respect to the given order of votes* if for every pair of candidates $a, b \in C$ such that $a \succ_1 b$ all votes where a is preferred to b precede all votes where b is preferred to a; $\mathcal{P}$ is *single-crossing* if the votes in $\mathcal{P}$ can be permuted so that it becomes single-crossing with respect to the resulting order of votes.

A profile $\mathcal{P} = (\succ_1, \dots, \succ_n)$ of total orders over C is said to be 1-*Euclidean* if there is a mapping ρ of voters and candidates into the real line such that $c \succ_i d$ if and only if $|\rho(i) - \rho(c)| < |\rho(i) - \rho(d)|$. A 1-Euclidean profile is both single-peaked and single-crossing.

3 Preference Restrictions

We will now define a number of constraints that a dichotomous profile may satisfy. Most of these constraints can be divided into two basic groups: those that are based on ordering voters and/or candidates on the line and requiring the votes to respect this order (this includes VEI, VI, CEI, CI, DE, and DUE), and those that are based on viewing votes as weak orders and asking if there is a single-peaked/single-crossing/1-Euclidean profile of total orders that refines the given profile (this includes PSP, PSC, and PE); we remark that the study of the latter type of constraints was initiated by Lackner [2014]. We will also consider constraints that are based on partitioning voters/candidates (2PART and PART), as well as two constraints (WSC and SSC) that have been introduced in a recent paper of Elkind et al. [2015] in order to understand the best way of extending the single-crossing property to weak orders.

Fix a profile $\mathcal{P} = (v_1, \dots, v_n)$ over C.

1. 2-*partition (2PART)*: We say that $\mathcal{P}$ satisfies 2PART if $\mathcal{P}$ contains only two distinct votes v, v', and $v \cap v' = \emptyset$, $v \cup v' = C$.
2. *Partition (PART)*: We say that $\mathcal{P}$ satisfies PART if C can be partitioned into pairwise disjoint subsets $C_1, \dots, C_\ell$ such that $\{v_1, \dots, v_n\} = \{C_1, \dots, C_\ell\}$ (i.e., each voter in $\mathcal{P}$ approves one of the sets $C_1, \dots, C_\ell$).
3. *Voter Extremal Interval (VEI)*: We say that $\mathcal{P}$ satisfies VEI if the voters in $\mathcal{P}$ can be reordered so that for every candidate c the voters that approve c form a prefix or a suffix of the ordering. Equivalently, both the voters who approve c and the voters who disapprove c form an interval of that ordering. See Figure 1 for an example.
4. *Voter Interval (VI)*: We say that $\mathcal{P}$ satisfies VI if the voters in $\mathcal{P}$ can be reordered so that for every candidate c the voters that approve c form an interval of that ordering. See Figure 2 for an example.
5. *Candidate Extremal Interval (CEI)*: We say that $\mathcal{P}$ satisfies CEI if candidates in C can be ordered so that each of the sets v_i forms a prefix or a suffix of that ordering. Equivalently, both v_i and $\overline{v_i}$ form an interval of that ordering.
6. *Candidate Interval (CI)*: We say that $\mathcal{P}$ satisfies CI if candidates in C can be ordered so that each of the sets v_i forms an interval of that ordering.
7. *Dichotomous Uniformly Euclidean (DUE)*: We say that $\mathcal{P}$ satisfies DUE if there is a mapping ρ of voters and candidates into the real line and a radius r such that for every voter i it holds that $v_i = \{c : |\rho(i) - \rho(c)| \leq r\}$.
8. *Dichotomous Euclidean (DE)*: We say that $\mathcal{P}$ satisfies DE if there is a mapping ρ of voters and candidates into the real line such that for every voter i there exists a radius r_i with $v_i = \{c : |\rho(i) - \rho(c)| \leq r_i\}$.
9. *Possibly single-peaked (PSP)*: We say that $\mathcal{P}$ satisfies PSP if there is a single-peaked profile of total orders $\mathcal{P}'$ that is a refinement of $\mathcal{P}$.
10. *Possibly single-crossing (PSC)*: We say that $\mathcal{P}$ satisfies PSC if there is a single-crossing profile of total orders $\mathcal{P}'$ that is a refinement of $\mathcal{P}$.
11. *Possibly Euclidean (PE)*: We say that $\mathcal{P}$ satisfies PE if there is a 1-Euclidean profile of total orders $\mathcal{P}'$ that is a refinement of $\mathcal{P}$.
12. *Seemingly single-crossing (SSC)*: We say that $\mathcal{P}$ satisfies SSC if the voters in $\mathcal{P}$ can be reordered so that for each pair of candidates $a, b \in C$ it holds that either all votes v_i with $a \in v_i, b \notin v_i$ precede all votes v_j with $a \notin v_j$, $b \in v_j$ or vice versa.
13. *Weakly single-crossing (WSC)*: We say that $\mathcal{P}$ satisfies WSC if the voters in $\mathcal{P}$ can be reordered so that for each pair of candidates $a, b \in C$ it holds that each of the vote sets $V_1 = \{v_i : a \in v_i, b \notin v_i\}$, $V_2 = \{v_i : a \notin v_i, b \in v_i\}$, $V_3 = \{v \in \mathcal{P} : v \notin V_1 \cup V_2\}$ forms an interval of this ordering, with V_3 appearing between V_1 and V_2.

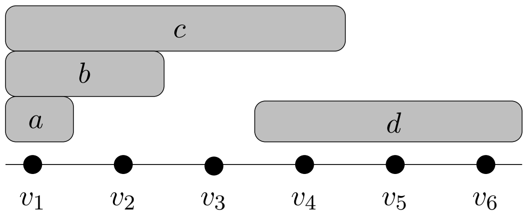

Figure 1: Voter Extremal Interval

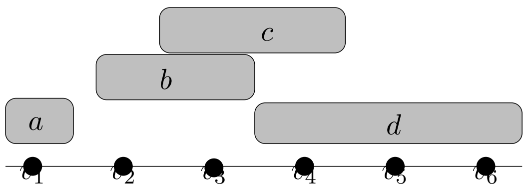

Figure 2: Voter Interval

3.1 Relations

The relationships among the properties defined above are depicted in Figure 3, where arrows indicate containment, i.e., more restrictive notions are at the top. All these containments are strict. We will now sketch proofs of some of the relationships indicated in this diagram; see [Elkind and Lackner, 2015] for the omitted proofs.

The four arrows at the top level of the diagram are immediate: any profile with at most two distinct votes where each candidate is approved in at least one of these votes satisfies VEI, CEI and WSC, and by definition 2PART is a special case of PART.

To understand the arrows in the next level, we first characterize the dichotomous profiles that are weakly single-crossing.

Lemma 1. *A dichotomous profile $\mathcal{P}$ satisfies WSC if and only if there exist three votes u, v, w such that*

(1) for every $v_i \in \mathcal{P}$ it holds that $\succ_{v_i} \in \{\succ_u, \succ_v, \succ_w\}$, and

(2) $\succ_v$ is equal to either $\succ_{u\cap w}$ or $\succ_{u\cup w}$.

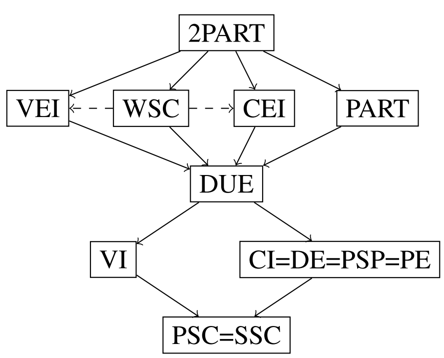

Figure 3: Relations between notions of structure. Dashed lines indicate that the respective containment holds only subject to additional conditions.

Proof sketch. It is easy to check that every profile satisfying (1)–(2) satisfies WSC. For the converse direction, assume without loss of generality that the ordering of the votes $v_1 \sqsubset v_2 \sqsubset \cdots \sqsubset v_n$ witnesses that $\mathcal{P}$ satisfies WSC. Let $u = v_1$, $w = v_n$, and set $C_1 = u \cap w$, $C_2 = u \cap \overline{w}$, $C_3 = \overline{u} \cap w$, $C_4 = \overline{u} \cap \overline{w}$. The WSC property implies that for every $\ell = 1, 2, 3, 4$, every $a, b \in C_\ell$, and every $v_i \in \mathcal{P}$ we have $a \in v_i$ if and only if $b \in v_i$, i.e., candidates in each C_ℓ occur as a block in all votes. Note that $v_1 = u = C_1 \cup C_2$, $v_n = w = C_1 \cup C_3$.

Suppose that $C_1, C_4 \neq \emptyset$. Then $C_1 \subseteq v_i$, $C_4 \subseteq \overline{v_i}$ for all $v_i \in \mathcal{P}$. Thus, if $\mathcal{P}$ contains a vote $v_i \neq u, w$, it has to be the case that $v_i = C_1 = u \cap w$ or $v_i = C_1 \cup C_2 \cup C_3 = u \cup w$; moreover, if both of these votes occur simultaneously and are distinct from each other and u, w (i.e., $C_2, C_3 \neq \emptyset$), the WSC property is violated. When C_1 or C_4 is empty, the analysis is similar; note, however, that trivial votes ($v_i = C$ and $v_i = \emptyset$) may alternate arbitrarily without violating the WSC property (this is why the lemma is stated in terms of weak orders rather than approval votes). $\square$

We can now show that under mild additional conditions (no trivial voters/candidates) WSC implies VEI and CEI.

Proposition 2. *Let $\mathcal{P}$ be a dichotomous profile that either contains only two distinct votes or contains no vote v_i with $v_i = \emptyset$. If $\mathcal{P}$ satisfies WSC, then it satisfies VEI.*

Proof. Assume without loss of generality that $\mathcal{P}$ satisfies WSC with respect to an ordering of voters $v_1 \sqsubseteq \cdots \sqsubseteq v_n$, and let $u = v_1$, $w = v_n$. We will show that $\mathcal{P}$ satisfies VEI with respect to $\sqsubset$. If $\mathcal{P}$ only contains two distinct votes, this claim is immediate, so assume that $\emptyset \notin \mathcal{P}$. Consider a vote $v \in \mathcal{P}$ that is distinct from u and w. Since $\emptyset \notin \mathcal{P}$, by Lemma 1 there exist i, j with $1 < i < j < n$ such that $v_k = u$ for $k < i$, $v_k = v$ for $k = i, \ldots, j$, $v_k = w$ for $k > j$, and $v \in \{u \cup w, u \cap w\}$. Suppose first that $v = u \cap w$. Then candidates in $u \cap w$ are approved by all voters, candidates in $u \setminus w$ are approved by the first $i - 1$ voters, candidates in $w \setminus u$ are approved by the last $n - j$ voters, and the remaining candidates are not approved by anyone. On the other hand, if $v = u \cup w$, then candidates in $u \cap w$ are approved by all voters, candidates in $u \setminus w$ are approved by the first j voters, candidates in $w \setminus u$ are approved by the last $n - i + 1$ voters, and the remaining candidates are not approved by anyone. $\square$

The condition that the profile must not contain $\emptyset$ is necessary: the profile $(\{a, b\}, \emptyset, \{b, c\})$ satisfies WSC, but not VEI.

Proposition 3. *Let $\mathcal{P}$ be a dichotomous profile that either contains only two distinct votes or in which every candidate is approved in at least one vote and disapproved in at least one vote. If $\mathcal{P}$ satisfies WSC, then it satisfies CEI.*

Proof. Suppose that $\mathcal{P}$ is WSC with respect to an ordering of voters $\sqsubset$; let u and w be, respectively, the first and the last vote in this ordering. If $\mathcal{P}$ contains a trivial vote, it contains at most two non-trivial votes, in which case the claim is obvious. Thus, assume that it contains no trivial votes. Then we have $u \cap w = \emptyset$ (any candidate in $u \cap w$ would be approved by all voters) and $\overline{u} \cap \overline{w} = \emptyset$ (any candidate in $\overline{u} \cap \overline{w}$ would be disapproved by all voters). It is now easy to see that ordering the candidates so that all candidates approved by u precede all candidates approved by w witnesses that $\mathcal{P}$ is CEI. $\square$

To see that conditions of Proposition 3 are necessary, consider the profile $(\{a, b\}, \{b, c\})$ over $\{a, b, c, d\}$ and the profile $(\{a, b\}, \{b\}, \{b, c\})$ over $\{a, b, c\}$: both of these profiles satisfy WSC, but not CEI.

Next, we will relate CEI and VEI to DUE.

Proposition 4. *If a dichotomous profile $\mathcal{P}$ satisfies CEI or VEI, then it satisfies DUE.*

Proof. Suppose first that $\mathcal{P}$ satisfies CEI with respect to the ordering $c_1 \lhd \cdots \lhd c_m$ of candidates. Map the candidates into the real line by setting $\rho(c_i) = i$, and let $r = m$. We can now place each voter i to the left or to the right of all candidates at an appropriate distance so that the set of candidates within distance r from him coincides with v_i. For VEI the argument is similar: if $\mathcal{P}$ satisfies VEI with respect to the ordering $v_1 \sqsubseteq \cdots \sqsubseteq v_n$ of voters, we place voters on the real line according to $\rho(i) = i$, let $r = n$, and place each candidate to the left or to the right of all voters at an appropriate distance. $\square$

The proof that WSC implies DUE is also based on our characterization of WSC preferences.

Proposition 5. *If a dichotomous profile $\mathcal{P}$ satisfies WSC, then it satisfies DUE.*

Proof. Clearly empty votes can be ignored when checking whether a profile satisfies DUE, so assume $\mathcal{P}$ contains to empty votes. Then it contains at most three distinct votes u, v, w with $v = u \cap w$ or $v = u \cup w$. Set $\rho(c) = 1$ for $c \in u \setminus w$, $\rho(c) = 2$ for $c \in u \cap w$, $\rho(c) = 3$ for $c \in w \setminus u$, $\rho(c) = 10$ for $c \notin u \cup w$. We set $r = 1$ if $v = u \cap w$ and $r = 2$ if $v = u \cup w$, and position the voters accordingly. $\square$

The last arrow on this level is from PART to DUE: here, the containment is straightforward, as the candidates approved by each voter can be placed as a block on the axis, with the respective voter(s) placed in the center of this block. Also, it is immediate that DUE implies VI and CI. Perhaps more surprisingly, the classes CI, DE, PSP and PE coincide.

Proposition 6. *Let $\mathcal{P}$ be a dichotomous profile. Then the following conditions are equivalent: (a) $\mathcal{P}$ satisfies PE (b) $\mathcal{P}$ satisfies PSP (c) $\mathcal{P}$ satisfies CI (d) $\mathcal{P}$ satisfies DE.*

Proof sketch. Suppose $\mathcal{P}$ satisfies PE, and let $\mathcal{P}'$ be a refinement of $\mathcal{P}$ that, together with a mapping ρ, witnesses this. Then $\mathcal{P}'$ is single-peaked and therefore $\mathcal{P}$ satisfies PSP. If $\mathcal{P}$ satisfies PSP, as witnessed by a refinement $\mathcal{P}'$ and an axis $\lhd$, then $\mathcal{P}$ satisfies CI with respect to $\lhd$. If $\mathcal{P}$ satisfies CI with respect to an order $\lhd$ of candidates, we can map the candidates into the real axis in the order suggested by $\lhd$ so that the distance between every two adjacent candidates is 1. We can then choose an appropriate approval radius and position for each voter. Finally, if $\mathcal{P}$ satisfies DE, as witnessed by a mapping ρ, we can use this mapping to construct a refinement of $\mathcal{P}$; by construction, this refinement is 1-Euclidean (we may have to modify ρ slightly to avoid ties). $\square$

constraint	complexity
2PART	poly (trivial)
PART	poly (trivial)
VEI	poly (CONSECUTIVE 1S)
CEI	poly (CONSECUTIVE 1S)
WSC	poly [Elkind *et al.*, 2015]
DUE	open
VI	poly (CONSECUTIVE 1S)
CI=DE=PSP=PE	poly (CONSECUTIVE 1S)
PSC=SSC	open

Table 1: The complexity of detecting structure in dichotomous profiles

An argument similar to the one used in the proof of Proposition 6 shows that every PE profile is PSC. Interestingly, the converse is not true.

Example 1. Consider the profile $\mathcal{P} = (\{a, b\}, \{a, c\}, \{b, c\})$ over $C = \{a, b, c\}$. It satisfies PSC, as witnessed by the single-crossing refinement $(a \succ b \succ c, c \succ a \succ b, c \succ b \succ a)$. However, in every refinement of $\mathcal{P}$ the first voter ranks c last, the second voter ranks b last, and the third voter ranks a last. Thus, no such refinement can be single-peaked, and, consequently, no such refinement can be 1-Euclidean.

The equivalence between PSC and SSC is not entirely obvious: while it is clear that a profile that violates SSC also violates PSC, to prove the converse one needs to use an argument similar to the proof of Theorem 4 in [Elkind *et al.*, 2015]. We omit the proof of this result, as well as the proof of the following proposition, due to space constraints.

Proposition 7. *If a dichotomous profile $\mathcal{P}$ satisfies VI, it also satisfies SSC.*

For all other pairs of constraints, we have examples (omitted) showing that one does not imply the other.

3.2 Detection

To exploit the constraints defined in Section 3, we have developed algorithms that can decide whether a given profile belongs to one of the restricted domains defined by these constraints. Our results are summarized in Table 1.

Clearly, verifying whether a given profile satisfies 2PART or PART is straightforward. For most of the remaining problems, we can proceed by a reduction to the classic CONSECUTIVE 1S problem [Booth and Lueker, 1976]. This problem asks if the columns of a given 0-1 matrix can be permuted in such a way that in each row of the resulting matrix the 1s are consecutive, i.e., the 1s form an interval in each row; it admits a linear-time algorithm [Booth and Lueker, 1976].

Theorem 8. *Detecting whether a dichotomous profile satisfies CEI, CI, VI or VEI is possible in $\mathcal{O}(m \cdot n)$ time.*

Proof. Let $C = \{c_1, c_2, \ldots, c_m\}$ and $\mathcal{P} = (v_1, v_2, \ldots, v_n)$. We construct an instance of CONSECUTIVE 1S in slightly different ways, depending on the property we want to detect. In all cases, we obtain a "yes"-instance if and only if the given profile has the desired property.

Let us start with CI. For each vote, we create one row of the matrix: for each $i \in [n]$ and $j \in [m]$, the j-th entry of the i-th row is 1 if $c_j \in v_i$ and 0 otherwise. In this way, we obtain an $m \times n$ matrix. Permuting the columns of this matrix so that 1s form an interval in each row is equivalent to permuting candidates so that the set of candidates approved by each voter forms an interval. For CEI, we combine the matrix for CI with its complement, i.e., we add a second row for each vote v_i, so that the j-the entry of that row is 0 of $c_j \in v_i$ and 1 otherwise. A column permutation of the resulting $m \times 2n$ matrix such that 1s form an interval in each row corresponds to permuting candidates so that for each voter both the set of her approved candidates and the set of her disapproved candidates form an interval; this is equivalent to the CEI property. For VI and VEI the argument is similar. □

For WSC, Elkind et al. [2015] provide an algorithm that works for any weak orders (not just dichotomous ones). They leave the complexity of detecting PSC and SSC as an open problem, and we have not been able to resolve it for dichotomous weak orders. The problem of recognizing DUE preferences remains open as well, though it is plausible that a linear-programming based algorithms similar to those of [Doignon and Falmagne, 1994; Knoblauch, 2010; Elkind and Faliszewski, 2014] exist.

4 Algorithms for Committee Selection

In this section, we consider two classic approval-based committee selection rules—Proportional Approval Voting (PAV) and Maximin Approval Voting (MAV)—and argue that we can design efficient algorithms for these rules when voters' preferences belong to some of the domains in our list (for some of the richer domains, we may need to place mild additional restrictions on voters' preferences).

We start by providing formal definitions of these rules.

Definition 1. Every non-increasing infinite sequence of non-negative reals $\mathbf{w} = (w_1, w_2, \dots)$ that satisfies $w_1 = 1$ defines a *committee selection rule* $\mathbf{w}$-PAV. This rule takes a set of candidates C, a dichotomous profile $\mathcal{P} = (v_1, \dots, v_n)$ and a target committee size $k \le |C|$ as its input. For every size-k subset W of C, it computes its $\mathbf{w}$-PAV score as $\sum_{v_i \in \mathcal{P}} u_{\mathbf{w}}(|W \cap v_i|)$, where $u_{\mathbf{w}}(p) = \sum_{j=1}^{p} w_j$, and outputs a size-$k$ subset with the highest $\mathbf{w}$-PAV score, breaking ties arbitrarily. The $\mathbf{w}$-PAV rule with $\mathbf{w} = (1, \frac{1}{2}, \frac{1}{3}, \dots)$ is usually referred to simply as the PAV rule, and we write $u(p) = 1 + \dots + \frac{1}{p}$.

In what follows we assume that the entries of $\mathbf{w}$ are rational and w_i can be computed in time $\mathrm{poly}(i)$.

Definition 2. Given a set of candidates C, a dichotomous profile $\mathcal{P} = (v_1, \dots, v_n)$ and a target committee size $k \le |C|$, the MAV-score of a size-k subset W of C is computed as $\max_{v_i \in \mathcal{P}}(|W \setminus v_i| + |v_i \setminus W|)$. MAV outputs a size-k subset with the lowest MAV score, breaking ties arbitrarily.

The $\mathbf{w}$-PAV rule is defined by Kilgour and Marshall [2012], see also [Kilgour, 2010]. Intuitively, under this rule each voter is assumed to derive a utility of 1 from having exactly one of his approved candidates in the winning set; his marginal utility from having more of his approved candidates in the winning set is non-increasing. The goal of the rule is to maximize the sum of players' utilities. In contrast, MAV [Brams *et al.*, 2007] has an egalitarian objective: for each candidate committee, it computes the dissatisfaction of the least happy voter, and outputs a committee that minimizes the quantity.

Computing the winning committee under MAV and PAV is NP-hard, see, respectively, [LeGrand *et al.*, 2007] and [Skowron *et al.*, 2015a; Aziz *et al.*, 2014]. The hardness result for PAV extends to $\mathbf{w}$-PAV as long as $\mathbf{w}$ satisfies $w_1 > w_2$; moreover, it holds even if each voter approves of at most two candidates or if each candidate is approved by at most three voters.

We will now show that PAV admits an algorithm whose running time is polynomial in the number of voters and the number of candidates if the input profile satisfies CI or VI and, furthermore, each voter approves at most s candidates or each candidate is approved by at most d voters, where s and d are given constants. More specifically, we prove that PAV winner determination for CI and VI preferences is in FPT with respect to parameter s and in XP with respect to parameter d. For simplicity, we state our results for PAV; however, all of them can be extended to $\mathbf{w}$-PAV.

In what follows, we write $[x : y]$ to denote the set $\{z \in \mathbb{Z} : x \le z \le y\}$.

Theorem 9. *Given a dichotomous profile $\mathcal{P} = (v_1, \dots, v_n)$ over a candidate set $C = \{c_1, \dots, c_m\}$ and a target committee size k, if $|v_i| \le s$ for all $v_i \in \mathcal{P}$ and $\mathcal{P}$ satisfies VI, then we can find a winning committee under PAV in time $\mathcal{O}(2^{2s} \cdot k \cdot n)$.*

Proof. Assume that $\mathcal{P}$ satisfies VI with respect to the order of voters $v_1 \sqsubseteq \dots \sqsubseteq v_n$. For each triple (i, A, ℓ), where $i \in [1 : n]$, $A \subseteq v_i$, and $\ell \in [0 : k]$, let $r(i, A, \ell)$ be the maximum utility that the first i voters can obtain from a committee W such that $W \cap v_i = A$, $|W| = \ell$, and $W \subseteq v_1 \cup \dots \cup v_i$.

We have $r(1, A, |A|) = u(|A|)$ for every $A \subseteq v_1$ and $r(1, A, \ell) = -\infty$ for every $A \subseteq v_1$, $\ell \in [0 : k] \setminus \{|A|\}$. To compute $r(i+1, A, \ell)$ for $i \in [1 : n-1]$, $A \subseteq v_{i+1}$ and $\ell \in [0 : k]$, we let $p = |A \setminus v_i|$ and set

$$r(i+1, A, \ell) = \max_{D \subseteq v_i \setminus v_{i+1}} r(i, D \cup (A \cap v_i), \ell - p) + u(|A|).$$

Indeed, every committee W with $|W| = \ell$, $W \cap v_{i+1} = A$, $W \subseteq v_1 \cup \dots \cup v_{i+1}$ contains exactly $\ell - p$ candidates from $v_1 \cup \dots \cup v_i$ and its intersection with v_i is of the form $D \cup (A \cap v_i)$, where candidates in D are approved by v_i, but not v_{i+1}. We output $\max_{A \subseteq v_n} r(n, A, k)$.

This dynamic program has $n \cdot 2^s \cdot (k+1)$ states, and the value of each state is computed using $\mathcal{O}(2^s)$ arithmetic operations. Assuming that basic calculations take constant time, we obtain a total runtime of $\mathcal{O}(2^{2s} \cdot k \cdot n)$. □

A similar dynamic programming algorithm can be used if voters' preferences satisfy CI. We omit its description due to space constraints.

Theorem 10. *Given a dichotomous profile $\mathcal{P} = (v_1, \dots, v_n)$ over a candidate set $C = \{c_1, \dots, c_m\}$ and a target committee size k, if $|v_i| \le s$ for all $v_i \in \mathcal{P}$ and $\mathcal{P}$ satisfies CI, then we can find a winning committee under PAV in time $\mathcal{O}(2^s \cdot n \cdot m)$.*

Our next theorem also considers CI and VI preferences, and deals with the case where no candidate is approved by too many voters. Just as the algorithms in the proofs of Theorems 9 and 10, the algorithms for this case are based on dynamic programming; we omit them due to space constraints.

Theorem 11. *Given a dichotomous profile $\mathcal{P} = (v_1, \dots, v_n)$ over a candidate set $C = \{c_1, \dots, c_m\}$ and a target committee size k, if $|\{i \mid c \in v_i\}| \leq d$ for all $c \in C$ and $\mathcal{P}$ satisfies CI or VI, then we can find a winning committee under PAV in time* $\mathrm{poly}(d, m, n, k^d)$.

The reader may wonder if constraints on s and d in Theorems 9, 10 and 11 are necessary. We conjecture that the answer is yes, i.e., winner determination under PAV remains hard under CI and VI preferences.

However, for "truncated" weight vectors $\mathbf{w}$ we can find $\mathbf{w}$-PAV winners in polynomial time. As the $(1, 0, \dots)$-PAV rule is essentially the classic Chamberlin–Courant rule [Chamberlin and Courant, 1983] for dichotomous preferences, our next result can be seen as an extension of the results of [Betzler *et al.*, 2013] and [Skowron *et al.*, 2015b] for the Chamberlin–Courant rule and single-peaked and single-crossing preferences: while we work on a less expressive domain (dichotomous preferences vs. total orders), we can handle a larger class of rules (all weight vectors with a constant number of non-zero entries rather than just $(1, 0, \dots,))$.

Theorem 12. *Consider a weight vector $\mathbf{w}$ where $w_i = 0$ for $i > i_0$ for some constant i_0. Then given a dichotomous profile $\mathcal{P} = (v_1, \dots, v_n)$ over a candidate set $C = \{c_1, \dots, c_m\}$ and a target committee size k, if $\mathcal{P}$ satisfies VI or CI, we can find a winning committee under $\mathbf{w}$-PAV in polynomial time.*

Moreover, for the more restricted domains, such as VEI, CEI, WSC and PART we can design polynomial-time algorithms for both MAV and PAV, under no additional constraints on preferences (again, our results extend to $\mathbf{w}$-PAV).

Theorem 13. *Given a dichotomous profile $\mathcal{P} = (v_1, \dots, v_n)$ over a candidate set $C = \{c_1, \dots, c_m\}$ and a target committee size k, if $\mathcal{P}$ satisfies VEI, CEI, WSC or PART, we can find a winning committee under MAV and PAV in polynomial time.*

Proof sketch. Consider first VEI. Assume without loss of generality that $\mathcal{P}$ satisfies VEI for voter order $v_1 \sqsubset \dots \sqsubset v_n$. Each candidate in C belongs to one of the following four groups: $C_1 = v_1 \cap v_n$, $C_2 = v_1 \setminus v_n$, $C_3 = v_n \setminus v_1$, and $C_4 = \overline{v_1} \cap \overline{v_n}$; candidates in C_1 are approved by all voters and candidates in C_4 are not approved by any of the voters.

Suppose first that $|C_1 \cup C_2 \cup C_3| < k$. Then there exists an optimal committee for both PAV and MAV that contains all candidates in $C_1 \cup C_2 \cup C_3$ and exactly $k - |C_1 \cup C_2 \cup C_3|$ candidates from C_4. Hence, we can now assume that this is not the case. Then there exist an optimal committee that contains no candidates from C_4.

Now, if $|C_1| \geq k$, an optimal committee for both PAV and MAV consists of k candidates from C_1, and if $|C_1| < k$, there exists an optimal committee that contains all candidates in C_1. It remains to decide how to allocate the remaining places among candidates in C_2 and C_3. To do so, we observe that there is a natural ordering over each of these sets: given a pair of candidates (c, c') in $C_2 \times C_2$ or $C_3 \times C_3$, we write $c \leq c'$ if $\{i : c \in v_i\} \subseteq \{i : c' \in v_i\}$. Note that every two candidates in C_2 are comparable with respect to $\leq$, and so are every two candidates in C_3. It is now easy to see that there exists an optimal committee (for PAV or MAV) that consists of candidates in C_1, top p candidates in C_2 with respect to $\leq$ and top r candidates in C_3 with respect to $\leq$ for some non-negative values of p, r with $p + r + |C_1| = k$. Thus, by considering at most k^2 possibilities for p and r, we can find an optimal committee.

The argument for CEI is similar to the one for VEI: we have to decide how many candidates to select from each end of the candidate ordering witnessing that $\mathcal{P}$ satisfies CEI. For WSC, we can use the characterization in Lemma 1; the problem then boils down to deciding how many candidates to select from each of the sets $u \setminus w$, $u \cap w$ and $w \setminus u$. For PART and PAV, we can show that an optimal committee can be found by a natural greedy algorithm that at each point selects the candidate with the largest "marginal contribution" to the total utility. For PART and MAV, we check, for each $t = 0, \dots, n$, whether there exists a committee whose MAV-score is at most t. This is the case if for each voter $v \in \mathcal{P}$ we can select at least $(|v| + k - t)/2$ candidates from v. Thus, if $v_1, \dots, v_\ell$ are the distinct votes in $\mathcal{P}$, we need to check that $\sum_{i=1}^{\ell} |v_i| \leq \ell t - (\ell - 2)k$. □

5 Conclusions and Open Problems

We have initiated research on analogues of the notions of single-peakedness and single-crossingness for dichotomous preference domains. We have proposed many constraints that capture some aspects of what it means for dichotomous preferences to be single-dimensional, explored the relationship among them, and showed that these constraints can be useful for identifying efficiently solvable special cases of hard voting problems on dichotomous domains. The algorithmic results in Section 4 can be seen as a proof that our approach has merit; however, there is certainly room for improvement there, both in terms of removing restrictions on the sizes of approval sets and number of voters that approve each candidate (for PAV) and in terms of considering larger domains, such as PSC for PAV and CI/VI for MAV.

For many of our constraints, we have provided efficient algorithms for checking whether a given dichotomous profile satisfies that constraint; two notable open cases are DUE and PSC/SSC. In particular, it would be interesting to understand if every profile that satisfies both VI and CI also satisfies DUE; this can be seen as an analogue of the question of whether every single-peaked single-crossing profile of total orders is 1-Euclidean (see discussion in [Doignon and Falmagne, 1994; Elkind *et al.*, 2014]). We can also ask if it is possible to detect if a given dichotomous profile is close to satisfying a structural constraint, and whether such "almost-structured" profiles have useful algorithmic properties; similar issues for profiles of total orders have recently received a lot of attention in the literature [Cornaz *et al.*, 2012; 2013; Bredereck *et al.*, 2013; Erdélyi *et al.*, 2013; Elkind and Lackner, 2014; Faliszewski *et al.*, 2014].

Acknowledgments

The second author has been supported by the Austrian Science Foundation FWF, grant P25518-N23 and Y698.

References

[Aziz *et al.*, 2014] H. Aziz, S. Gaspers, J. Gudmundsson, S. Mackenzie, N. Mattei, and T. Walsh. Computational aspects of multi-winner approval voting. In *AAMAS'15*, 2014.

[Betzler *et al.*, 2013] N. Betzler, A. Slinko, and J. Uhlmann. On the computation of fully proportional representation. *Journal of AI Research*, 47:475–519, 2013.

[Black, 1958] D. Black. *The Theory of Committees and Elections*. Cambridge University Press, 1958.

[Booth and Lueker, 1976] K. Booth and G. Lueker. Testing for the consecutive ones property, interval graphs, and graph planarity using PQ-tree algorithms. *Journal of Computer and System Sciences*, 13(3):335–379, 1976.

[Brams *et al.*, 2007] S. Brams, D. M. Kilgour, and R. M. Sanver. A minimax procedure for electing committees. *Public Choice*, 132(3-4):401–420, 2007.

[Brandt *et al.*, 2010] F. Brandt, M. Brill, E. Hemaspaandra, and L. Hemaspaandra. Bypassing combinatorial protections: Polynomial-time algorithms for single-peaked electorates. In *AAAI'10*, pages 715–722, 2010.

[Bredereck *et al.*, 2013] R. Bredereck, J. Chen, and G. Woeginger. Are there any nicely structured preference profiles nearby? In *IJCAI'13*, pages 62–68, 2013.

[Chamberlin and Courant, 1983] B. Chamberlin and P. Courant. Representative deliberations and representative decisions: Proportional representation and the Borda rule. *Am Pol Sci Rev*, 77(3):718–733, 1983.

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

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

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

[Elkind and Faliszewski, 2014] E. Elkind and P. Faliszewski. Recognizing 1-euclidean preferences: An alternative approach. In *SAGT'14*, pages 146–157, 2014.

[Elkind and Lackner, 2014] E. Elkind and M. Lackner. On detecting nearly structured preference profiles. In *AAAI'14*, pages 661–667, 2014.

[Elkind and Lackner, 2015] E. Elkind and M. Lackner. Structure in dichotomous preferences. Technical report, available on arXiv.org, 2015.

[Elkind *et al.*, 2014] E. Elkind, P. Faliszewski, and P. Skowron. A characterization of the single-peaked single-crossing domain. In *AAAI'14*, pages 654–660, 2014.

[Elkind *et al.*, 2015] E. Elkind, P. Faliszewski, M. Lackner, and S. Obraztsova. The complexity of recognizing incomplete single-crossing preferences. In *AAAI'15*, pages 865–871, 2015.

[Erdélyi *et al.*, 2013] G. Erdélyi, M. Lackner, and A. Pfandler. Computational aspects of nearly single-peaked electorates. In *IJCAI'13*, pages 283–289, 2013.

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

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

[Kilgour and Marshall, 2012] D. M. Kilgour and E. Marshall. Approval balloting for fixed-size committees. In D. S. Felsenthal and M. Machover, editors, *Electoral Systems*, Studies in Choice and Welfare, pages 305–326. Springer, 2012.

[Kilgour, 2010] D. M. Kilgour. Approval balloting for multi-winner elections. In *Handbook on Approval Voting*, chapter 6. Springer, 2010.

[Knoblauch, 2010] V. Knoblauch. Recognizing one-dimensional euclidean preference profiles. *Journal of Mathematical Economics*, 46:1–5, 2010.

[Lackner, 2014] M. Lackner. Incomplete preferences in single-peaked electorates. In *AAAI'14*, pages 742–748, 2014.

[LeGrand *et al.*, 2007] R. LeGrand, E. Markakis, and A. Mehta. Some results on approximating the minimax solution in approval voting. In *AAMAS'07*, pages 1185–1187, 2007.

[Magiera and Faliszewski, 2014] K. Magiera and P. Faliszewski. How hard is control in single-crossing elections? In *ECAI'14*, pages 579–584, 2014.

[Mirrlees, 1971] J. Mirrlees. An exploration in the theory of optimal income taxation. *Review of Economic Studies*, 38:175–208, 1971.

[Skowron *et al.*, 2015a] P. Skowron, P. Faliszewski, and J. Lang. Finding a collective set of items: From proportional multirepresentation to group recommendation. In *AAAI'15*, 2015.

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

[Walsh, 2007] T. Walsh. Uncertainty in preference elicitation and aggregation. In *AAAI'07*, pages 3–8, 2007.