Single-peakedness Based on the Net Preference Matrix: Characterization and Algorithms

Olivier Spanjaard and Paul Weng

Abstract

In this paper, we propose a new relaxation of the single-peaked property by constraining the net preference matrix instead of the preference profile. More precisely, this new domain restriction holds for a given profile $\mathcal{P}$ if there exists a single-peaked profile $\mathcal{P}'$ with the same net preference matrix as $\mathcal{P}$. We provide a characterization of net preference matrices that can be "implemented" by a single-peaked profile, and propose a polynomial time algorithm to recognize these matrices. In cases where preferences are not net single-peaked, we recall a dynamic programming approach for obtaining an axis over candidates such that the net preference matrix respects as much as possible (in some sense) our net single-peakedness condition. We also formulate an integer linear program for partial net single-peakedness. Finally we present the results of numerical tests on real election data.

1 Introduction

Research in social choice aims at studying collective decision problems from the viewpoint of axiomatic properties and very often at promoting fairness and strategy-proofness. The standard input data in a social choice problem are preference profiles, i.e., multisets of rankings (representing the individual preferences of voters) over alternatives (candidates). The main issue is then to provide rules able to aggregate the input rankings into a unique "collective" ranking synthesizing the individual preferences. One important result in the field is Arrow's theorem establishing the impossibility to design an aggregation rule satisfying a small set of desirable properties [20]. A popular way to circumvent this impossibility result is to drop the universality property, stating that any preference profile is acceptable as input, and instead requiring that the input preferences satisfy specific consistency conditions.

The most famous domain restriction is the single-peakedness condition introduced by Black [3]. Intuitively, preferences are single-peaked if 1) all voters agree on a left-right axis on the alternatives, and 2) the preferences of all voters decrease along the axis when moving away from their preferred alternative to the right or left. If preferences are single-peaked, then it prevents the appearance of cycles in the majority relation (i.e., a candidate is preferred to another candidate if she is preferred by a majority of voters). Note this property naturally holds if the alternatives are for instance taxation levels: the preferences are very likely to be single-peaked with respect to the order over reals.

However, as emphasized by Feld and Grofman [13], it becomes an extremely strong condition if the alternatives are candidates in an election. It requires indeed on the one hand that all voters agree on a common axis over the candidates, on the other hand that no individual preference -even slightly- deviates from the single-peaked condition. Given a left-right axis A, the number of rankings consistent with A (i.e., such that condition 2 holds) is 2^{m-1}, over $m!$ possible rankings in total, where m is the number of alternatives. The proportion of consistent rankings within all possible rankings thus quickly becomes tiny (for instance, $2^{m-1}/m! \approx 0.01$ for $m = 7$), as well as the likelihood that no voter deviates from this subset of preferences.

This observation is corroborated by the experimentations carried out by Sui et al. [23] on 2002 Irish General Election data in Dublin West and Dublin North, where the best axes explain only 2.9% and 0.4% of voter preferences respectively. In other words, respectively 97.1% and 99.6% maverick voters[1] have to be removed to make the profile single-peaked. This can be viewed as a relaxation of single-peakedness, where the number k of removed voters measures the quality of the approximation. It was named *k-maverick single-peakedness* by Faliszewski et al. [12], and is also known as *partial single-peakedness* in economics [21]. Several other relaxations of the single-peaked domain have been proposed in the computational social choice community [11, 10, 12], but even combinations of them with loose approximation parameters are not able to explain more than 50% of voter preferences at best. To adress this issue, Sui et al. proposed to consider multi-dimensional single-peaked consistency [1] to explain a preference profile, a relaxation of single-peakedness involving several axes over the candidates (each axis representing one political dimension). Focusing on the two-dimensional case and considering the same Irish election data sets, their experiments show that it explains 65.7% and 47.3% of voter preferences respectively.

We propose here a new relaxation of single-peakedness that we call *net single-peakedness* and that involves only one dimension. We say that a preference profile $\mathcal{P}$ is *net single-peaked* if there exists another profile $\mathcal{P}'$ that is single-peaked and exhibits the same net preference matrix as $\mathcal{P}$. The net preference matrix of $\mathcal{P}$ is the skew-symmetric matrix $M = (M_{ij})$ where M_{ij} is the number of voters in $\mathcal{P}$ that prefers candidate i to j minus the number of voters that prefers j to i. The key idea is that we do not require *all* individual preferences to be single-peaked w.r.t. a common axis, merely that group majority choices are made *as if* the preferences were single-peaked w.r.t. a common axis. In other words, we require *aggregate* consistency regarding an axis and not *individual* consistency of *each* individual.

Our work has been partly inspired by results of Feld and Grofman [13] who have examined conditions on the net preferences over *triples* of candidates (i.e., the frequency of ordering $i \succ j \succ k$ minus the frequency of ordering $k \succ j \succ i$, where i, j, k denote candidates). Besides the fact that their conditions differ from the one we propose (we consider net preferences over *pairs* of candidates), their goal is different: they did not aim at identifying an axis over the candidates, but at guaranteeing the transitivity of the majority relation. The idea of considering conditions on *aggregate* preferences instead of *individual* preferences was also present in the notion of semi single-peakedness introduced by Rasch [22]. The semi single-peakedness condition consists in relaxing condition 2 above by only requiring that the preferences of a majority of the voters decrease along the axis on the left of a given candidate c, and the preferences of a majority (but not necessarily the same voters) decrease on the right of c. Note that our net single-peakedness condition is much weaker.

We provide a characterization of net preference matrices M that can be “implemented” by a single-peaked profile (in the sense that there exists a single-peaked profile whose corresponding net preference matrix is M). Thanks to this characterization result, we propose an algorithm to recognize net single-peakedness. For the sake of completeness, we also briefly recall an already known dynamic programming procedure to compute best axes regarding several relaxations of net single-peakedness and we present a method based on integer linear programming for computing the best axis for partial net singled-peakedness. These procedures will allow us to perform numerical tests on real election data to evaluate to which extent the net single-peakedness assumption holds (in a sense to be made precise later). The main motivation behind this work is to try to identify a domain restriction (here, net single-peakedness) that could be more realistic than the impartial culture assumption (stating that all input rankings are equally likely). This would be of great interest, for instance, in the average-case complexity analysis of winner determination or manipulation algorithms.

[1]The *maverick* voters do not vote according to a societal axis, but according to their own biases.

2 Background and Notations

For any $i, j \in \mathbb{N}$, let $[\![i, j]\!]$ denote the set $\{i, i+1, \ldots, j\}$ if $i \leq j$, and the empty set if $i > j$. Let $V = \{v_1, v_2, \ldots, v_n\}$ be a set of n voters and $C = \{c_1, c_2, \ldots, c_m\}$ a set of m candidates. Let $\mathcal{P}$ denote the (multi-)set of preference relations $\succ_v$ over C for every $v \in V$. This set is called a *preference profile* if the preference relations $\succ_v$ are linear orders.

For a permutation σ of index set $[\![1, m]\!]$, we denote by $(\sigma_1, \ldots, \sigma_m)$ the sequence $(\sigma^{-1}(1), \ldots, \sigma^{-1}(m))$. We call this sequence the *axis* of σ (or σ*-axis* for short).

A preference profile $\mathcal{P}$ is *single-peaked* for a σ-axis if for all $v \in V$ there exists $p \in [\![1, m]\!]$ such that: if $p > j > i$ or $i > j > p$, then $c_{\sigma_p} \succ_v c_{\sigma_j} \succ_v c_{\sigma_i}$. For any voter v, index σ_p is called the *peak* of the σ-axis for v. A preference profile $\mathcal{P}$ is *single-peaked* if it is single-peaked for some σ-axis.

For a matrix $M \in \mathcal{M}_{m\times m}(\mathbb{Z})$, we denote by M^t its transpose, and by M_{ij} its entry at position (i, j). Let $A(\mathcal{P}) \in \mathcal{M}_{m\times m}(\mathbb{Z})$ be the matrix, called the *weighted majority matrix*, such that $A(\mathcal{P})_{ij}$ is equal to the number of voters preferring candidate c_i to candidate c_j in profile $\mathcal{P}$. The *net preference matrix* $B(\mathcal{P})$ of profile $\mathcal{P}$ is defined by $B(\mathcal{P}) = A(\mathcal{P}) - A(\mathcal{P})^t$. Note that this matrix is skew-symmetric. A matrix $M \in \mathcal{M}_{m\times m}(\mathbb{Z})$ is said to be a net preference matrix if there exists a preference profile $\mathcal{P}$ such that $M = B(\mathcal{P})$.

We recall a result by Debord [8], which characterizes net preference matrices and will help prove Theorem 2:

Theorem 1 (Debord 1987) *Let $M \in \mathcal{M}_{m\times m}(\mathbb{Z})$ be a skew-symmetric matrix. The following two propositions are equivalent:*

- *Matrix M is a net preference matrix.*
- *All off-diagonal entries of M are of the same parity.*

In the next section, we present our relaxed condition of single-peakedness, called net single-peakedness.

3 Net Single-peakedness

We say that a profile $\mathcal{P}$ is *net single-peaked for a σ-axis* if there exists a profile $\mathcal{P}'$, single-peaked for this σ-axis, such that $\mathcal{P}$ and $\mathcal{P}'$ have the same net preference matrix, i.e., $B(\mathcal{P}) = B(\mathcal{P}')$. In other words, in net single-peakedness, one does not take into account pairwise preferences that are outvoted. Furthermore, a profile is said to be *net single-peaked* if it is net single-peaked for some axis.

Example 1 *Let $m = 4$ and $n = 5$. The profile $\mathcal{P}$ can be read in Table 1 where each column represents the preference relation of a voter (decreasing preference from the top). The net preference matrix is written:*

$$B(\mathcal{P}) = \begin{pmatrix} 0 & 1 & 1 & -1 \\ -1 & 0 & 3 & -1 \\ -1 & -3 & 0 & -1 \\ 1 & 1 & 1 & 0 \end{pmatrix}$$

Noting that three distinct candidates are ranked in the last positions for at least one voter, we conclude that the profile is not single-peaked (at most two candidates can be ranked in the last position in a single-peaked profile [11]). However, $\mathcal{P}$ is net single-peaked as $B(\mathcal{P}) = B(\mathcal{P}')$ for profile $\mathcal{P}'$ defined in Table 1, which is single-peaked for axis $(1, 4, 2, 3)$.

v_1	v_2	v_3	v_4	v_5
c_1	c_2	c_3	c_4	c_1
c_4	c_3	c_4	c_2	c_2
c_2	c_4	c_1	c_1	c_4
c_3	c_1	c_2	c_3	c_3

v'_1	v'_2	v'_3
c_1	c_2	c_4
c_4	c_3	c_1
c_2	c_4	c_2
c_3	c_1	c_3

Table 1: Profiles $\mathcal{P}$ (left) and $\mathcal{P}'$ (right) in Example 1.

In order to give a characterization of net single-peakedness, we introduce the class $\mathcal{M}_{\leq}$ of skew-symmetric integer matrices whose entries above the main diagonal are nondecreasing when moving to the right or down in the matrix, i.e., $\mathcal{M}_{\leq} = \{M \in \mathcal{M}_{m\times m}(\mathbb{Z}) \,|\, \forall i < j < k, M_{ij} \leq M_{ik} \text{ and } M_{ik} \leq M_{jk}\}$.

For any permutation σ and matrix M, M^σ denotes the matrix whose rows and columns have been permuted according to the σ-axis, i.e., $M^\sigma_{i,j} = M_{\sigma_i,\sigma_j}$.

A profile is *canonical* if it consists of one single preference relation. For any $x, y \in \mathbb{Z}$ such that $x < y$, let $\mathcal{M}_{\leq}(x, y)$ be the subset of $\mathcal{M}_{\leq}$ consisting of matrices whose off-diagonal entries are either x or y. We first give a characterization of single-peaked canonical profiles.

Lemma 1 *A canonical preference profile $\mathcal{P}$ is single-peaked for a given σ-axis if and only if $B(\mathcal{P})^\sigma \in \mathcal{M}_{\leq}(-1, 1)$.*

We do not present its proof due to lack of space.

Example 2 *Let $m = 3$ and σ be the identity permutation. The net preference matrices of canonical profiles that are single-peaked for the σ-axis are:*

$$\begin{pmatrix} 0 & 1 & 1 \\ -1 & 0 & 1 \\ -1 & -1 & 0 \end{pmatrix} \quad \begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & 1 \\ -1 & -1 & 0 \end{pmatrix}$$
$$\mathcal{P} = \{c_1 \succ c_2 \succ c_3\} \quad \mathcal{P} = \{c_2 \succ c_1 \succ c_3\}$$

$$\begin{pmatrix} 0 & -1 & -1 \\ 1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix} \quad \begin{pmatrix} 0 & -1 & -1 \\ 1 & 0 & -1 \\ 1 & 1 & 0 \end{pmatrix}$$
$$\mathcal{P} = \{c_2 \succ c_3 \succ c_1\} \quad \mathcal{P} = \{c_3 \succ c_2 \succ c_1\}$$

Our characterization result can now be formulated.

Theorem 2 *A preference profile $\mathcal{P}$ is* net single-peaked *for a given σ-axis if and only if $B(\mathcal{P})^\sigma \in \mathcal{M}_{\leq}$.*

Proof. Assume that $\mathcal{P}$ is net single-peaked for σ. By definition, there exists a profile $\mathcal{P}'$ single-peaked for σ such that $B(\mathcal{P}) = B(\mathcal{P}')$. Profile $\mathcal{P}'$ can be seen as the union of canonical profiles $\mathcal{P}_1, \ldots, \mathcal{P}_n$, all single-peaked for σ. Therefore, $B(\mathcal{P}') = \sum_i B(\mathcal{P}_i)$. By Lemma 1, all matrices $B(\mathcal{P}_i)^\sigma$ are in $\mathcal{M}_{\leq}(-1, 1) \subseteq \mathcal{M}_{\leq}$. As the addition operation is closed in $\mathcal{M}_{\leq}$, we deduce that $B(\mathcal{P}')^\sigma \in \mathcal{M}_{\leq}$, and thus $B(\mathcal{P})^\sigma \in \mathcal{M}_{\leq}$ because $B(\mathcal{P})^\sigma = B(\mathcal{P}')^\sigma$.

Conversely, suppose that for a profile $\mathcal{P}$, we have $B(\mathcal{P})^\sigma \in \mathcal{M}_{\leq}$ for a given permutation σ. Let $M = B(\mathcal{P})^\sigma$. Note that the skew-symmetric matrix I_1 (resp. I_{-1}) consisting of 1's (resp. -1) in the strictly upper triangular part belongs to $\mathcal{M}_{\leq}(-1, 1)$ and that adding I_1 to all matrices of $\mathcal{M}_{\leq}(-1, 1)$ yields $\mathcal{M}_{\leq}(0, 2)$. By Theorem 1, the off-diagonal entries of M are of the same parity. W.l.o.g., assume that they are even (otherwise add I_{-1}). It is easy to see that the matrices in $\mathcal{M}_{\leq}(0, 2)$ define a generating set of the subset of matrices in $\mathcal{M}_{\leq}$ whose off-diagonal entries are even. In other words, matrix M writes $M = \sum_i a_i M_i$, where $a_i \in \mathbb{N}$ and $\{M_1, M_2, \ldots\} = \mathcal{M}_{\leq}(0, 2)$. We therefore have $M = \sum_i a_i(M_i + I_{-1}) + \sum_i a_i I_1$ because $I_{-1} + I_1$ is the zero matrix. Note that $M_i + I_{-1} \in \mathcal{M}_{\leq}(-1, 1)$ and $I_1 \in \mathcal{M}_{\leq}(-1, 1)$. Let $\mathcal{P}_i$ denote the canonical profile such that $B(\mathcal{P}_i) = M_i + I_{-1}$, and $\mathcal{P}_0$ denote the canonical profile such that $B(\mathcal{P}_0) = I_1$. The profile $\mathcal{P}'$ defined as the union of $\sum_i a_i$ profiles $\mathcal{P}_0$, a_1

profiles $\mathcal{P}_1$, a_2 profiles $\mathcal{P}_2$, ... is single-peaked for σ because $\mathcal{P}_0, \mathcal{P}_1, \dots$ are all single-peaked for σ. Furthermore, $B(\mathcal{P}')^\sigma = M = B(\mathcal{P})^\sigma$ and therefore $B(\mathcal{P}') = B(\mathcal{P})$. $\square$

Example 3 *(Example 1 cont'd) For the σ-axis $(1, 4, 2, 3)$, one can check that*

$$B(\mathcal{P})^\sigma = \begin{pmatrix} 0 & -1 & 1 & 1 \\ 1 & 0 & 1 & 1 \\ -1 & -1 & 0 & 3 \\ -1 & -1 & -3 & 0 \end{pmatrix}$$

belongs to $\mathcal{M}_\leq$, and therefore $\mathcal{P}$ is net single-peaked by Theorem 2. Thus, our characterization result makes it possible to check if a profile $\mathcal{P}$ is net single-peaked without searching for a single-peaked profile $\mathcal{P}'$ such that $B(\mathcal{P}) = B(\mathcal{P}')$.

Furthermore, our characterization result is related to a statement made (without justification) in an article by Greenberg [14] dealing with *1-Euclidean preferences*, also called *Coombs' unfolding model* [6, 7]. Preferences are *1-Euclidean* if the voters and candidates can be mapped to points on the real line so that each voter prefers a candidate that is closer to her to the one that is further away. The statement in the article by Greenberg is that *if* preferences are 1-Euclidean, *then* $B(\mathcal{P})^\sigma \in \mathcal{M}_\leq$ for permutation σ corresponding to the ordering of the candidates on the real line. Our result is stronger in two ways:

- the validity domain is enlarged because 1-Euclidean preferences are clearly single-peaked,
- the sufficient condition is turned into a *necessary and sufficient* condition.

Having a characterization result, a natural subsequent question is now to investigate its algorithmic implications in order to recognize a net single-peaked profile.

4 Algorithms for Net Single-Peakedness

4.1 Recognizing Net Single-peakedness

Our characterization result states that recognizing a net single-peaked profile $\mathcal{P}$ is equivalent to detecting if there exists a permutation σ such that the permuted net preference matrix $B(\mathcal{P})^\sigma$ belongs to $\mathcal{M}_\leq$. This problem resembles the *seriation problem.* We first describe a known solution method for the seriation problem before explaining how it can be adapted in order to recognize net single-peakedness.

Given a symmetric $m \times m$ matrix M whose entries are nonnegative (and whose diagonal entries are null), the seriation problem asks if there exists a permutation σ such that M^σ is *anti-Robinson*, i.e., the entries of M^σ are monotonically nondecreasing when moving away from the main diagonal (along a row or a column).

A polynomial-time algorithm for the seriation problem consists in considering all balls $\mathrm{Ball}_\alpha(i) = \{j : M_{ij} \leq \alpha\}$ for $i \in [\![1, m]\!]$. The idea is that M^σ is anti-Robinson if and only if every ball consists of consecutive elements in permutation σ. (Note that it is sufficient to only consider balls on rows because the symmetry of the matrix implies that the property also holds on columns if it holds on rows.) This enables to solve the seriation problem by reducing it to the *consecutive ones problem* [19].

The consecutive ones problem is stated as follows: given a binary matrix A, does there exist a permutation of its columns such that the consecutive ones property holds, i.e., the 1's are consecutive in every row? In a seminal paper, Booth and Lueker [4] provides a *polynomial time* algorithm to decide if a binary matrix A has the consecutive ones property and, if yes, compute a *PQ-tree* for A. A PQ-tree is a concise data structure which gives an implicit representation of *all* the consecutive-ones orderings of the columns of A.

From a symmetric matrix M, one generates a binary matrix A as follows. For each index $j \in [\![1, m]\!]$, the matrix A contains a corresponding column. For each ball $\text{Ball}_\alpha(i)$, the matrix A contains a corresponding row with value 1 in column j if $j \in \text{Ball}_\alpha(i)$. The consecutive-ones orderings of the columns of A (if it exists) correspond to the permutations σ for which M^σ is anti-Robinson.

The PQ-tree algorithm by Booth and Lueker solves the consecutive ones problem in $O(x + y + z)$ time, where x and y are the numbers of columns and rows, and z is the total number of 1's in matrix A. For each $i \in [\![1, m]\!]$, there are at most m distinct nonempty balls and therefore rows in A, involving at most $m(m + 1)/2$ values 1. In total, A has $x = m$ columns, at most $y = m^2$ rows and at most $z = m^2(m + 1)/2$ values 1. Hence, an anti-Robinson matrix can be recognized in $O(m + m^2 + m^2(m + 1)/2) = O(m^3)$ time. By using the previous transformation, a concise representation of *all* permutations σ for which M^σ is anti-Robinson can therefore be recovered in $O(m^3)$ time.

We now describe how to adapt this approach for recognizing a net preference matrix that belongs to $\mathcal{M}_\leq$. Clearly, a skew-symmetric matrix M^σ belongs to $\mathcal{M}_\leq$ if and only if the entries are monotonically nondecreasing when moving away from the main diagonal[2] *along a row.* (By skew-symmetry of the matrix, if this property holds on rows, then the entries are monotonically nonincreasing when moving away from the main diagonal along a column.) To adapt the solution method used for the seriation problem, the only requirement is to redefine a ball so that it always includes the main diagonal entry. Note that this property mechanically holds in the seriation problem because $M_{ii} = 0$ and $M_{ij} \geq 0 \ \forall j$, which is not the case for a net preference matrix. We therefore consider all balls $\text{Ball}_\alpha(i) = \{i\} \cup \{j : M_{ij} \leq \alpha\}$ for $i \in [\![1, m]\!]$. Similarly to the seriation problem, the idea is that M^σ belongs to $\mathcal{M}_\leq$ if and only if every ball consists of consecutive elements in permutation σ. It is therefore possible to obtain a concise representation of *all* permutations σ for which M^σ belongs to $\mathcal{M}_\leq$ in $O(m^3)$ time.

Example 4 *(Example 1 cont'd) Consider the net preference matrix of Example 1:*

$$B(\mathcal{P}) = \begin{pmatrix} 0 & 1 & 1 & -1 \\ -1 & 0 & 3 & -1 \\ -1 & -3 & 0 & -1 \\ 1 & 1 & 1 & 0 \end{pmatrix}$$

When generating matrix A from $B(\mathcal{P})$, the balls that contain all indices in $[\![1, 4]\!]$ can be omitted, as well as those that are singletons, because they do not matter for the consecutive ones property to hold. The only balls that matter are thus $Ball_{-1}(1) = \{1, 4\}$, $Ball_{-1}(2) = \{1, 2, 4\}$ and $Ball_{-3}(3) = \{2, 3\}$. The obtained binary matrix A is indicated below, where the first (resp. second, third) row corresponds to $Ball_{-1}(1)$ (resp. $Ball_{-1}(2)$, $Ball_{-3}(3)$):

$$A = \begin{pmatrix} 1 & 0 & 0 & 1 \\ 1 & 1 & 0 & 1 \\ 0 & 1 & 1 & 0 \end{pmatrix}$$

There are four permutations (two different ones and their symmetric) σ of the columns of A for which the 1's are consecutive in each row, namely $(1, 4, 2, 3)$, $(4, 1, 2, 3)$, $(3, 2, 4, 1)$ and $(3, 2, 1, 4)$. One can easily check that $B(\mathcal{P})^\sigma$ belongs to $\mathcal{M}_\leq$ for those permutations (see Example 3).

When a profile P is not net single-peaked, a natural question to ask is whether it is close to a profile that is net single-peaked. We present two approaches for measuring how close a profile is from net single-peakedness.

[2] Regardless of the main diagonal entry itself.

4.2 Reorganizing the Net Preference Matrix

Even if the preferences are not net single-peaked, it can be informative to reorganize the net preference matrix (i.e., permute the indices) so that the entries are "as much as possible" nondecreasing within a row moving to the right from the main diagonal, and "as much as possible" nonincreasing within a column moving up from the main diagonal. This can be formalized by defining an objective function that reflects these gradient conditions.

For several natural objective functions, an optimal reorganization (i.e., permutation) can be performed by a dynamic programming procedure proposed by Hubert and Golledge [16] (see also Hubert et al. [17]). For the paper to be self-contained, we describe here this procedure. As noted by Hahsler et al. [15], the typical objective function is written as follows for a permutation σ of the candidates:

$$f(\sigma) = \sum_{i<k<j} d(M^{\sigma}_{i,k}, M^{\sigma}_{i,j}) + \sum_{i<k<j} d(M^{\sigma}_{i,j}, M^{\sigma}_{k,j}) \tag{1}$$

where $d(x, y)$ is a function which defines how a violation or satisfaction of a gradient condition for a triple (i, k, j) is counted. Hubert et al. [17] suggest two possible counting functions d_1 and d_2:

$$\begin{aligned} d_1(x,y) &= \begin{cases} 1 & \text{if} \quad x < y \\ 0 & \text{if} \quad x = y \\ -1 & \text{if} \quad x > y \end{cases} \\ d_2(x,y) &= y - x \end{aligned}$$

Function d_1 counts the number of satisfactions of the gradient conditions minus the number of violations. Function d_2 weighs each satisfaction or violation by its magnitude given by the absolute difference between the values. An optimal reorganization is obtained for a permutation σ^* such that $f(\sigma^*) = \max_\sigma f(\sigma)$. The optimal value of the objective function can be computed by using the following recursion:

$$\begin{aligned} g(\emptyset) &= 0, \\ g(C') &= \max_{k \in C'} [g(C' \setminus \{k\}) + d_r(C', k) + d_c(C', k)] \end{aligned}$$

for $C' \subseteq C$, where $g(C')$ is the contribution to the final value of the objective function for the candidates in C', placed in some order in the first $|C'|$ rows and columns, and:

$$\begin{aligned} d_r(C', k) &= \textstyle\sum_{i \in C' \setminus \{k\}} \sum_{j \in \overline{C'}} d(M_{i,k}, M_{i,j}), \\ d_c(C', k) &= \textstyle\sum_{j \in \overline{C'}} \sum_{i \in C' \setminus \{k\}} d(M_{i,j}, M_{k,j}). \end{aligned}$$

where $\overline{C'} = C \setminus C'$ and d is set to d_1 or d_2.

Note that $g(C) = f(\sigma^*)$. The principle of the recursion is illustrated in Figure 1. Horizontal and vertical symbols "$\leq$" indicate the tests of gradient conditions that have already been carried out in previous recursions. While computing $d_r(C', k) + d_c(C', k)$, only the gradient conditions *between* the grey blocks are considered, corresponding to the question marks in the figure.

Given that the candidates in $C' \setminus \{k\}$ (resp. $\overline{C'}$) are placed in the first $|C'| - 1$ (resp. last $|\overline{C'}|$) rows and columns, the marginal contribution to the final value for the objective function of candidate k, placed at the $|C'|^{\text{th}}$ position, depends neither on how the candidates in $C' \setminus \{k\}$ are ordered nor on how the candidates in $\overline{C'}$ are ordered. This is the key property that enables the dynamic programming solution procedure.

There are 2^m subsets C' of C, and the computation of each $g(C')$ requires to look up $|C'|$ values $g(C' \setminus \{k\})$ (note that $|C'| \leq m$), with $O(m^2)$ time for computing each $d_r(C', k) + d_c(C', k)$. The procedure is therefore clearly not polynomial time, more precisely

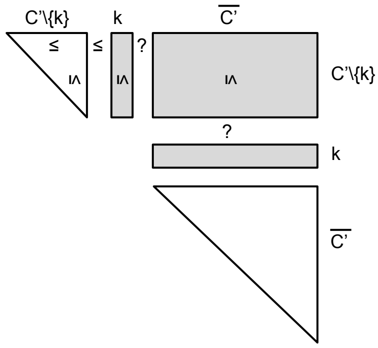

Figure 1: Illustration of the dynamic programming procedure.

it runs in $O(2^m m^3)$ time. The number m of candidates is nevertheless often small in political election contexts, thus the procedure is efficient enough to quickly compute the best permutations in our experiments on real election data.

4.3 Partial Net Single-peakedness

Another approach to define nearly net single-peakedness is based on determining the largest subset of voters in the profile that induce a net single-peaked matrix. The largest subset of a profile $\mathcal{P}$ for a permutation σ of candidates can be found by solving an integer linear program (ILP):

$$\begin{array}{lll} \max & \sum_v x_v & \\ \text{s.t.} & \sum_v B(v)^\sigma x_v \in \mathcal{M}_\leq & \\ & x_v \in \{0,1\} & \forall v \end{array} \tag{2}$$

where boolean variable x_v states if voter v is in the maximal subset or not and $B(v)$ denotes $B(\{\succ_v\})$ the net preference matrix of the canonical profile containing only the preference relation of voter v.

The largest subset of a profile $\mathcal{P}$ that induces a net single-peaked matrix can be found by considering every permutation of candidates. As this computation may be quite CPU intensive, two techniques can be used to make it faster. First, we can add the following constraint to the ILP: $\sum_v x_v \geq b$ where b is the best value found so far. Second, we can start the search for the best axis with a potentially good axis. We propose to start with a permutation σ that minimizes $h(\sigma) = \min_{M \in \mathcal{M}_\leq} \sum_{i,j} |B(\mathcal{P})^\sigma_{i,j} - M_{i,j}|$. This requires to solve $m!$ linear programs (LP). Given a permutation σ, one can indeed compute $h(\sigma)$ by solving the following LP:

$$\begin{array}{ll} \min & \sum_{i,j} e^+_{i,j} + e^-_{i,j} \\ \text{s.t.} & \left(B(\mathcal{P})^\sigma_{i,j} + e^+_{i,j} - e^-_{i,j}\right)_{1 \leq i \leq m, 1 \leq j \leq m} \in \mathcal{M}_\leq \\ & e^+_{i,j} \geq 0 \quad e^-_{i,j} \geq 0 \quad \forall i,j \end{array} \tag{3}$$

5 Numerical Experiments

We carried out experiments on data sets taken from the 2007 Glasgow City Council election as well as the 2002 Irish general election. These data sets are available in the PrefLib library, that collects preference data assembled by Mattei and Walsh [18].

Both the 2007 Glasgow election and the 2002 Irish election were separated by voting districts: 21 wards for the Glasgow election; 42 constituencies for the Irish election, among which we investigate here only the 3 constituencies where electronic voting machines were

voting district	votes	cand	axes	PNSP	unweighted		weighted	
					opt	prob	opt	prob
2007 Glasgow City Council election								
1. Linn	494	10/11	72	34.9693%	0.1535%	0.1047	0.0117%	0.0084
2. Newlands	648	9/9	24	31.0016%	0.0502%	0.012	0.075%	0.0178
3. GreaterPollock	818	9/9	72	32.0463%	0.0022%	0.0016	0.0028%	0.002
4. Craigton	718	10/10	144	26.7806%	0.0028%	0.004	0.1253%	0.1652
5. Govan	411	10/11	144	36.7246%	0.0031%	0.0045	0.0057%	0.0081
6. Pollokshields	767	7/9	12	33.6927%	0.2778%	0.0329	0.119%	0.0142
7. Langside	1040	8/8	24	35.6902%	0.2679%	0.0624	0.0942%	0.0224
8. SouthsideCentral	726	9/9	72	–	0.0006%	0.0004	0.0006%	0.0004
9. Calton	363	9/10	96	36.4903%	3.2595%	0.9585	5.7881%	0.9967
10. Anderston	593	8/9	24	34.6416%	0.2083%	0.0488	0.7589%	0.1672
11. Hillhead	630	9/10	96	–	0.2072%	0.1806	0.075%	0.0695
12. PartickWest	962	9/9	48	36.2460%	0.1207%	0.0563	0.0022%	0.0011
13. Garscadden	559	10/10	144	46.5580%	0.0497%	0.0691	0.0557%	0.0771
14. Drumchapel	556	10/10	144	41.1009%	0.022%	0.0312	0.0319%	0.0448
15. Maryhill	1071	8/8	24	24.6914%	0.0694%	0.0165	0.0248%	0.0059
16. Canal	419	10/11	144	–	0.013%	0.0185	0.1019%	0.1365
17. Springburn	365	10/10	48	36.1345%	0.1238%	0.0577	2.6594%	0.7258
18. EastCentre	284	12/13	1728	–	0.0107%	0.1682	–	–
19. Shettleston	405	11/11	144	39.3484%	0.0355%	0.0498	0.1368%	0.179
20. Baillieston	535	11/11	576	47.3282%	0.0001%	0.0005	0.0172%	0.0945
21. NorthEast	690	9/10	48	30.8846%	0.0457%	0.0217	0.0006%	0.0003
2002 Irish general election								
1. Dublin North	4259	9/12	12	56.3765%	5.1279%	0.4683	8.6172%	0.6609
2. Dublin West	4810	8/9	4	–	9.6379%	0.3333	11.3641%	0.3828
3. Meath	3166	9/14	36	49.7937%	0.0006%	0.0002	0.0006%	0.0002

Table 2: Results of numerical tests on data sets taken from 2007 Glasgow City Council election and 2002 Irish general election. Column "voting district" contains the name of each considered ward or constituency, column "votes" (resp. "axes") the number of complete ranking ballots (resp. the number of Wikipedia axes). Couple x/y in column "cand" gives the number x of candidates affiliated to a party and the total number y of candidates (including independent candidates). Symbol "–" means that the computation exceeded the time limit or ran out of memory.

used (Dublin North, Dublin West and Meath). Each ward (resp. constituency) involved different candidates and voters, and elected 3 or 4 councillors (resp. between 3 and 5 deputies) using the single transferable vote system. This implies that some political parties had several candidates for the same voting district. A ballot consists in the k most preferred candidates of a voter, for varying values of k. In order to fit the data with our setting, we restricted ourselves to the ballots for which $k = m$ (complete rankings of the candidates).

Partial net single-peakedness. The first question that comes to mind while carrying out numerical tests is to evaluate if net single-peakedness is significantly more likely than classical single-peakedness. The first notable result is that for none of the tested data sets the preference profile is net single-peaked. In order to deepen the analysis, we evaluated the maximum percentage of voters that constitutes a single-peaked electorate. The results are given in column "PNSP" (that stands for "partial net single-peakedness") of Table 2. It can be observed that the percentage obtained in the Irish general election are much greater than those obtained with the classical single-peakedness assumption (we recall that, for the classical single-peakedness assumption, the best axes explain only 2.9% and 0.4% of voter preferences in Dublin West and Dublin North).

Comparison with a reference axis. Another concern is to determine how a reference left-right axis over the candidates (obtained from an external source) compares with permutation σ^* optimizing objective function f. In order to build a reference axis that can be recovered by any experimenter, we used Wikipedia as external source. The free encyclopedia provides indeed a political position (of course debatable) for each political party (e.g., left wing, right wing, centre, centre right, etc.). We assumed that the political position of an affiliated candidate corresponds to that of the belonging party, and we built an axis over the candidates based on these positions. Actually, the "Wikipedia axis" is not unique since several parties can share the same political position. For instance, a Wikipedia axis reads $((1,3),2,(4,5))$, where the numbers are the indices of candidates and candidates $\{1,3\}$ as well as $\{4,5\}$ share the same political position. This corresponds to the following set of $2 \times 2 = 4$ axes: $(1,3,2,4,5)$, $(1,3,2,5,4)$, $(3,1,2,4,5)$, $(3,1,2,5,4)$. For this reason, we talk of Wikipedia *axes* in the following. Furthermore, we excluded the independent candidates from the data sets because we were not able to define a political position for them.

The numerical results are synthesized in Table 2. We compared the feasible axes by using the objective function in Equation 1 with counting functions d_1 (columns labeled by "unweighted") and d_2 (columns labeled by "weighted"). For both counting functions, it rarely happens that one of the Wikipedia axes is optimal for the corresponding objective function (twice for d_1, three times for d_2). To have further insights on the quality of the Wikipedia axes regarding the objective function, we also measured the percentage of axes (called *top axes* hereafter) that are better or equivalent to the best Wikipedia axis with respect to the considered objective function (column "top"). This can be done efficiently by slightly adapting the dynamic programming procedure described in Section 4.2: the computation took less than a minute (often a few seconds) on a modern personal computer for each voting district, except for ward 18 in Glasgow for which the computation took much longer (resp. ran out of memory) for d_1 (resp. d_2). To assess the significance of this percentage, we provide the probability that a randomly drawn sample of axes includes (at least) one of the top axes (column "prob"), where the size of the sample is exactly the number of Wikipedia axes. This probability is equal to $1 - \Pi_{i=0}^{a-1}(1 - \frac{t}{m!/2-i})$, where a is the number of Wikipedia axes (value in column "axes"), t is the number of top axes (obtained from column "top") and m is the number of affiliated candidates (first value in column "cand").

The conclusions that can be drawn from the results using one or the other of the objective functions are very similar. For almost all wards in the Glasgow election there is a Wikipedia axis that belongs to the very best axes regarding the objective function, and this is significant given the very low probability that it occurs, while for the Irish general election in Dublin North and Dublin West there is no significantly good axis among the Wikipedia axes. These latter results are consistent with those obtained by Sui et al. [23], that found very tiny single-peaked subelectorates for these constituencies. Overall, the numerical results seem to show that the net preference matrix provides some good hints about the way the electorate views the relative positions of the candidates on the left-right axis, but the nature of the information handled (net preferences) does not make it possible to fully learn an axis from it.

6 Conclusion

We introduced in this paper a new domain restriction, that we call *net single-peakedness*, derived from single-peakedness and based on the net preference matrix. We showed that a net single-peaked profile can be recognized in polynomial time thanks to a characterization result we established, and we performed numerical tests to assess the occurrence of net single-peakedness on real election data. While none of the tested profiles are net single-peaked, it appears that one can find axes that are compatible with a significant percentage of the voters. Furthermore, for almost all data sets, there exists a plausible left-right axis that is among the top axes regarding an objective function reflecting the fulfillment of the net single-peakedness conditions (the further from net single-peakedness, the worse the value of the objective function).

Nevertheless, as indicated above, it seems that the net preference matrix does not make it possible to fully learn an axis, while on the contrary the brute information given by the preference profile is "too rich", as shown by the results of Sui et al. [23]. An interesting research direction would be to study another type of aggregated information from which to try to learn a political axis.

Acknowledgements. This work was supported by the ANR-14-CE24-0007-01 project "CoCoRICo-CoDec".

References

[1] Barberà, S.; Gul, F.; and Stacchetti, E. 1993. Generalized median voter schemes and committees. *Journal of Economic Theory* 61(2):262289.

[2] Bartholdi, J.; Trick, M. 1986. Stable matching with preferences derived from a psychological model. *Operations Research Letters* 5(4):165-169.

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

[4] Booth, K., and Lueker, G. 1976. Testing for the consecutive ones property, interval graphs, and graph planarity using PQ-tree algorithms. *J. Comput. Syst. Sci.* 13(3):335–379.

[5] Clearwater, A.; Puppe, C.; and Slinko, A. 2015. Generalizing the single-crossing property on lines and trees to intermediate preferences on median graphs. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI 2015).*

[6] Coombs, C. 1950. Psychological scaling without a unit of measurement. *Psychological Review* 57:145–158.

[7] Coombs, C. 1964. *A theory of data.* Wiley.

[8] Debord, B. 1987. Caractérisation des matrices des préférences nettes et méthodes d'agrégation associées. *Mathématiques et Sciences humaines* 97:5–17.

[9] Demange, G. 1982. Single-peaked orders on a tree. *Mathematical Social Sciences* 3(4):389–396.

[10] Erdélyi, G.; Lackner, M.; and Pfandler, A. 2013. Computational aspects of nearly single-peaked electorates. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI 2013).*

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

[12] Faliszewski, P.; Hemaspaandra, E.; and Hemaspaandra, L. 2014. The complexity of manipulative attacks in nearly single-peaked electorates. *Artif. Intell.* 207:69–99.

[13] Feld, S., and Grofman, B. 1986. Research note partial single-peakedness: An extension and clarification. *Public Choice* 51(1):71–80.

[14] Greenberg, M. 1965. A method of successive cumulations for the scaling of pair-comparison preference judgments. *Psychometrika* 30(4):441–448.

[15] Hahsler, M.; Hornik, K.; and Buchta, C. 2008. Getting things in order: An introduction to the `R` package `seriation`. *Journal of Statistical Software* 25(3).

[16] Hubert, L., and Golledge, R. 1981. Matrix reorganization and dynamic programming: applications to paired comparisons and unidimensional seriation. *Psychometrika* 46(4):429–441.

[17] Hubert, L.; Arabie, P.; and Meulman, J. 2001. *Combinatorial Data Analysis.* Society for Industrial and Applied Mathematics.

[18] Mattei, N., and Walsh, T. 2013. Preflib: A library of preference data. In *Proceedings of Third International Conference on Algorithmic Decision Theory (ADT 2013)*, Lecture Notes in Artificial Intelligence. Springer.

[19] Mirkin, B., and Rodin, S. 1984. *Graphs and Genes.* Springer-Verlag.

[20] Moulin, H. 1988. *Axioms of cooperative decision making.* Econometric Society Monographs. Cambridge University Press, Cambridge.

[21] Niemi, R. 1969. Majority decision-making with partial unidimensionality. *The American Political Science Review* 63(2):488–497.

[22] Rasch, B. 1987. Manipulation and strategic voting in the norwegian parliament. *Public Choice* 52(1):57–73.

[23] Sui, X.; Francois-Nienaber, A.; and Boutilier, C. 2013. Multi-dimensional single-peaked consistency and its approximations. In *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013.*

[24] Trick, M. 1989. Recognizing single-peaked preferences on a tree. *Mathematical Social Sciences* 17(3):329–334.

Olivier Spanjaard
Sorbonne Universités, UPMC Univ Paris 06
Laboratoire d'Informatique de Paris 6
4 Place Jussieu
75005 Paris, France
Email: `olivier.spanjaard@lip6.fr`

Paul Weng
SYSU-CMU Joint Institute of Engineering
Office 217B
Guangzhou Higher Education Mega Center
132 East Waihuan Road
Guangzhou, 510006, P.R. China
Email: `paweng@cmu.edu`