

Multi-Dimensional Single-Peaked Consistency and Its Approximations

Xin Sui

Dept. of Computer Science
University of Toronto
xsui@cs.toronto.edu

Alex Francois-Nienaber

Dept. of Computer Science
University of Toronto
lex@cs.toronto.edu

Craig Boutilier

Dept. of Computer Science
University of Toronto
cebly@cs.toronto.edu

Abstract

Single-peakedness is one of the most commonly used domain restrictions in social choice. However, the extent to which agent preferences are single-peaked in practice, and the extent to which recent proposals for *approximate single-peakedness* can further help explain voter preferences, is unclear. In this article, we assess the ability of both single-dimensional and multi-dimensional approximations to explain preference profiles drawn from several real-world elections. We develop a simple *branch-and-bound* algorithm that finds multi-dimensional, single-peaked axes that best fit a given profile, and which works with several forms of approximation. Empirical results on two election data sets show that preferences in these elections are far from single-peaked in any one-dimensional space, but are nearly single-peaked in two dimensions. Our algorithms are reasonably efficient in practice, and also show excellent anytime performance.

1 Introduction

Social choice deals with the aggregation of individual preferences to determine a consensus decision. However, mechanisms for social choice are complicated by computational, communication and incentive issues. For example, general impossibility results limit the ability to incentivize agents to truthfully reveal their preferences [Gibbard, 1973; Satterthwaite, 1975]. Fortunately, certain *domain restrictions* allow one to circumvent these issues. The widely-studied notion of *single-peaked preferences* [Black, 1948; Moulin, 1980] does just this. When options can be ordered along an axis, and an agent’s preference for an option can be defined using the “distance” along this axis from her most preferred (or *ideal*) point, protocols such as the *median mechanism* reduce communication burden, ensure truthful reporting, and can often ease computational demands. *Multi-dimensional* single-peakedness, using multiple axes, also admits effective mechanisms [Barberà *et al.*, 1993; Barberà, 2010].

While conceptually attractive, single-peakedness is a very strong assumption, one unlikely to hold in realistic settings—e.g., elections with thousands of voters and more than a handful of candidates [Conitzer, 2009; Escoffier *et al.*, 2008]. However, one might hope that preferences are *approximately*

single-peaked, and thereby retain some of the advantages mentioned above. To this end, recent research has begun to investigate computational methods to test single-peakedness [Escoffier *et al.*, 2008], and various forms of approximation (e.g., by deleting outlier candidates, clustering candidates, deleting voters, or adding additional axes) [Escoffier *et al.*, 2008; Faliszewski *et al.*, 2011; Erdélyi *et al.*, 2012; Galand *et al.*, 2012]. These techniques, however, have focused on *one-dimensional (1D)* preferences, and have not been tested empirically to determine if these approximations *can explain observed preferences* in real-world elections or other social choice problems.

We address this issue in two ways. First, we test single-peaked consistency, and several forms of approximation (in isolation and in combination) on two election data sets to see if these approximations have any empirical explanatory power. To do so, we develop a *branch-and-bound (BB)* algorithm to find the *best (1D) axis* given a preference profile, i.e., the ordering of candidates for which the greatest number of voters are single-peaked. The algorithm is easily extended to support various forms of approximation. While the best-axis problem is computationally difficult, our method works well in practice. We show that *none* of the forms of approximation proposed in the literature come close to explaining voter preferences in these elections: the best axis explains under 2.9% of voter preferences in one case and under 0.4% in the other; and even aggressive approximation improves this to only 50% and 25%, respectively. To address this, we extend our algorithm to find the *best multi-dimensional ordering* to explain a preference profile. Focusing on the *two-dimensional (2D)* case, we show that exact 2D-single-peakedness explains the datasets much better—without approximation, the best axis set explains over 65.7% and 47.3% of voter preferences, respectively; and with a very small degree of approximation, the 2D model explains almost all voters.

Apart from our algorithmic developments, our findings suggest that a focus on multi-dimensional rather than 1D models can greatly enhance the applicability of single-peaked models in practice. In so doing, one may be able to apply richer and more effective preference learning and mechanism design techniques to practical problems.

2 Single-peakedness and its Approximations

Our setting is standard in social choice. Let $N = \{1, \dots, n\}$ be a set of agents (or *voters*) and $C = \{1, \dots, z\}$ be a set of z options (or *candidates*). Each voter $i \in N$ has a preference (or strict total order) \succ_i over C , with $a \succ_i b$ meaning i prefers a to b . A *preference profile* $\succ = \{\succ_1, \dots, \succ_n\}$ reflects the joint preferences of all voters.

Single-peaked preferences. An *axis* A is any strict ordering $<_A$ of the candidates C . Intuitively, it represents an ordering of candidates relative to some salient quality, e.g., on the left-right political spectrum. In the 1D case, we assume a single axis. Each voter i has an *ideal point* $\tau_i = \tau(\succ_i)$, which is her most preferred candidate: we say that i 's preference are *single-peaked* w.r.t. axis A if $a \succ_i b$ whenever $b <_A a <_A \tau_i$ or $\tau_i <_A a <_A b$. Informally, the “closer” a candidate is to i 's ideal point τ_i , the more preferred it is. Notice that, given A , i 's ideal point does not uniquely determine its full ranking \succ_i . But when all voter preferences are single-peaked relative to a fixed axis A , mechanisms such as the *median mechanism* can be used, in which voters need only specify their ideal points.

The definition can be extended to m dimensions by assuming m distinct axes. Let $A^m = \langle A_1, \dots, A_m \rangle$ be such an m -dimensional space:

Definition 1 [Barberà et al., 1993] i 's preference order \succ_i is single-peaked w.r.t. A^m if $a \succ_i b$ whenever we have, for each axis A_d , $d \leq m$, either $b <_{A_d} a <_{A_d} \tau_i$ or $\tau_i <_{A_d} a <_{A_d} b$.

In other words, if a lies within the “bounding box” of τ_i and b , then a is preferred to b , since a is “closer” to the peak τ_i than b . For any $a, b \in C$ satisfying this relationship, we refer to $\mathbf{b} = \langle \tau_i, a, b \rangle$ as a *bounding box constraint* on \succ_i , and denote $\mathbf{b}_1 = \tau_i$, $\mathbf{b}_2 = a$ and $\mathbf{b}_3 = b$. This does not restrict i 's preference for a and b if neither lies within the other's bounding box w.r.t. τ_i . As in the 1D case, a profile is *m-dimensional single-peaked* w.r.t. A^m if each voter preference \succ_i is. In such a case, *generalized median mechanisms* can be used to truthfully elicit voter peaks and efficiently select outcomes [Barberà et al., 1993; Barberà, 2010].

Example 1 The following 2D example with five candidates c_1, \dots, c_5 illustrates the concept. Assume two axes, $A_1 = c_4 <_{A_1} c_3 <_{A_1} c_1 <_{A_1} c_2 <_{A_1} c_5$ and $A_2 = c_1 <_{A_2} c_2 <_{A_2} c_4 <_{A_2} c_5 <_{A_2} c_3$. The only bounding box constraint is $\langle c_1, c_2, c_5 \rangle$, which induces two constraints on voter preferences: (a) if $\tau_i = c_1$, then we must have $c_2 \succ_i c_5$; and (b) if $\tau_i = c_5$, then we must have $c_2 \succ_i c_1$.

Single-peaked approximations. While single-peakedness is a powerful concept, preference profiles are unlikely to be single-peaked in practice, especially as the number of voters or candidates becomes large.¹ Several forms of *approximate single-peakedness* have been proposed recently that allow limited violations of the constraints imposed by single-peakedness. We now outline some of these.

¹In problems defined on metric spaces, such as facility location [Procaccia and Tennenholtz, 2009; Lu et al., 2010; Escoffier et al., 2011; Dokow et al., 2012; Sui et al., 2013] single-peakedness is more likely to hold, but even then may be compromised by considerations apart from distance.

Given a profile \succ , one important task is testing whether it is single-peaked. Escoffier et al. [2008] develop an algorithm for this purpose in the 1D setting. It has a running time of $O(nz)$ (recall n and c are the number of agents and candidates, respectively), improving on an earlier $O(n^2z)$ algorithm proposed by Bartholdi and Trick [1986]. It exploits the fact that candidates ranked last in any voter's ranking must lie at the extreme points of the axis. The algorithm finds an axis A that renders the profile (perfectly) single-peaked if one exists; but if no such A exists, it does not find a *best* axis that fits the greatest number of voters, nor does it generalize to multiple dimensions (see below).

Approximation methods attempt to find some minimal change to the profile that would render it single-peaked. Faliszewski et al. [2011] consider the removal of *maverick* voters to render a profile single-peaked (e.g., perhaps certain voters are “irrational” in their declared preferences). The aim is to delete as few mavericks as possible, which measures the quality of the approximation.

Definition 2 [Faliszewski et al., 2011] A profile \succ is k -maverick single-peaked if a profile \succ' obtained by removing at most k voters from \succ is single-peaked.

Erdélyi et al. [2012] consider *local candidate deletion (LCD)*, allowing the deletion of misordered candidates from each voter's preference—the notion is “local” since *different* candidates can be deleted from each \succ_i . The goal is to minimize the (local) number of candidates deleted.

Definition 3 [Erdélyi et al., 2012] A profile \succ is k -local candidate deletion (k -LCD) single-peaked if a profile \succ' obtained by removing at most k candidates from each \succ_i is single-peaked.

One can also approximate single-peakedness by allowing multiple axes, where each voter must be single-peaked w.r.t. at least one of these axes:

Definition 4 [Escoffier et al., 2008] A profile \succ is k -additional axis (k -AA) single-peaked if there are $k + 1$ axes A_1, \dots, A_{k+1} such \succ_i is single-peaked w.r.t. at least one axis, $\forall i \in N$.

It is important to note that k -AA single-peakedness, while it implies $k + 1$ -dimensional single-peakedness, is not equivalent to it. It imposes the stringent requirement that each voter be single-peaked with respect to *one* of the axes, something not needed in true multi-dimensional models.

Several other notions of approximation have been proposed, but these are somewhat weaker than those above, so we do not investigate them. Among these are *k-Dodgson* [Faliszewski et al., 2011], which allows performing at most k swaps of adjacent candidates in each voter's ranking. k -LCD is at least as powerful, since deleting a candidate is at least as effective as swapping two candidates [Erdélyi et al., 2012]. Another is *clustered single-peakedness* [Galand et al., 2012], which allows groups of candidates to be clustered and requires single-peakedness w.r.t. such clusters, with the aim of minimizing maximal cluster size. k -LCD (indeed “global”

candidate deletion) can simulate its effects (though their quality measures are somewhat different).

Other related work. Key research on single-peakedness in the 1D and multi-dimensional cases have been outlined above, as have recent proposals for approximation. See Barbera’s [2010] excellent survey for an overview of how single-peakedness is exploited in voting and mechanism design (and see Schummer and Vohra’s [2007] discussion of computational considerations).

The rich literature on *spatial models* for voter or consumer choice bears a strong relationship as well [Hotelling, 1929; Hinich, 1978; Poole and Rosenthal, 1985]. Spatial models explain voter choice by inferring (from data) distances between voters and candidates, and typically using some form of probabilistic choice based on these distances [Bradley and Terry, 1952; Luce, 1959; Shepard, 1959]. While the model is more restrictive than multi-dimensional single-peakedness in some senses, stochastic choice allows for accommodation of “misorderings,” much like approximations in single-peaked models. Spatial models are typically used to explain choice data rather than full preference rankings (see [Gormley and Murphy, 2007] for an exception).

Finally, work in *behavioral social choice* [Regenwetter *et al.*, 2006] addresses very related questions, observing that standard domain restrictions (such as single-peakedness or Sen’s value restriction) are rarely observed in real-world election data. Probabilistic variants of such conditions are used to explain (among other things) the stability often observed in elections, and specifically the overwhelming lack of majority cycles. We take a different approach, investigating various approximations to the domain restrictions themselves. We extend them to multiple dimensions, and focus on recovering approximate variants of single-peakedness to (ultimately) assist in the design of mechanisms. However, our conceptual stances are very much aligned.

3 One-Dimensional Branch and Bound

To find the best approximations for single-peaked preferences, we first consider the *best axis* problem: given a profile \succ , find a single axis A that explains the preferences of the greatest number of voters (i.e., renders \succ_i single-peaked). Letting the *score* $s(A)$ of an axis be the number of voters whose preferences are consistent with it, we want to find the axis A^* with the greatest score. This is, of course, just the optimization variant of the k -maverick problem, which is NP-complete [Erdélyi *et al.*, 2012]. However, we develop a branch-and-bound algorithm for this problem, and use this as a building block for generating additional axes and for supporting k -LCD approximation. We focus in this section on 1D single-peakedness.

3.1 The Algorithm

Our branch-and-bound algorithm, *1D-SPBB*, is specified in Alg. 1 and adopts ideas from the single-peaked consistency method of Escoffier *et al.* [2008]. Each node in the search tree is labeled by a *partial axis* of the form $A_{p,q} = \{(c_1, \dots, c_p), \dots, (c_q, \dots, c_z)\}$, where p candidates $c_1 <_{A_{p,q}} \dots <_{A_{p,q}} c_p$ are ordered on the left of the

Algorithm 1 1D-SPBB($A_{0,z+1} \leftarrow \emptyset, C_r \leftarrow C, lb^* \leftarrow 0$)

```

1: while There exists an unchecked axis do
2:   Pick two candidates  $c'$  and  $c''$  from  $C_r$ 
3:   Build the axis  $A_{p+1,q-1}$  from  $A_{p,q}$  by locating  $c'$  at
     the position of  $p+1$  and  $c''$  at  $q-1$ 
4:   if  $C_r \setminus \{c', c''\} = \emptyset$  then
5:     Mark the complete axis  $A_{p+1,q-1}$  as checked
6:     Compute the score  $s(A_{p+1,q-1})$ 
7:     if  $s(A_{p+1,q-1}) > lb^*$  then
8:        $lb^* \leftarrow s(A_{p+1,q-1})$  and  $A^* \leftarrow A_{p+1,q-1}$ 
9:   else
10:    Compute the upper bound  $ub$  for  $A_{p+1,q-1}$ 
11:    if  $ub > lb^*$  then
12:      1D-SPBB( $A_{p+1,q-1}, C_r \setminus \{c', c''\}, lb^*$ )
13:    else
14:      Mark the whole branch as checked
15:    Return

```

axis and $z - q + 1$ candidates $c_q <_{A_{p,q}} \dots <_{A_{p,q}} c_z$ are ordered on the right. The remaining candidates $C_r = C \setminus \{c_1, \dots, c_p, c_q, \dots, c_m\}$ are not yet ordered on the axis. An *extension* of partial axis $A_{p,q}$ is any complete axis A that retains the two suborderings and completes the ordering by placing the remaining candidates between the two in some fashion. Let $E(A_{p,q})$ be the set of extensions of $A_{p,q}$. We say a voter i is *consistent* with $A_{p,q}$ if \succ_i is single-peaked for some $A \in E(A_{p,q})$. 1D-SPBB also maintains, at each node, the list of voters who are consistent with that node’s partial axis.

The algorithm starts with an empty axis and extends it from the “outside in.” At each step, 1D-SPBB branches by placing two candidates in C_r at positions $p+1$ and $q-1$ of a partial axis $A_{p,q}$ to form a more complete axis $A_{p+1,q-1}$. It then computes the induced bounding box constraints and the corresponding consistent voters. The number of consistent voters provides a *upper bound* on the score $s(A)$ of any $A \in E(A_{p+1,q-1})$. If the axis is complete, this gives us the exact score $s(A)$ of this axis, and a *lower bound* on $s(A^*)$. In typical fashion, 1D-SPBB maintains a global lower bound lb^* , corresponding to the score of the best complete axis A^* found so far. It cuts the search for extensions of a partial axis $A_{p,q}$ when the upper bound on $A_{p,q}$ falls below lb^* ; and when it terminates, the best axis A^* is the optimal axis.

We now consider several key steps in the algorithm that ensure its practicality despite the theoretical hardness of the problem. First, note that axes of the form $A_{1,z}$ at the first level of the search tree fix only the two extreme points of the axis. Symmetry means that we need not consider any axis with c' at the leftmost position and c at the rightmost, if we have already expanded the partial axis with c leftmost and c' rightmost. This reduces the search tree size by a factor of two, improving efficiency.

A critical component of 1D-SPBB is the identification of consistent voters given a partial (or complete) axis. Given A , Alg. 2 computes an upper bound on the score of any $\tilde{A} \in E(A)$; and if A is complete, it will return $s(A)$. We let $A[j]$ denote the candidate at the j th position of A and $\succ_{i,t}$ the candidate ranked t th in voter i ’s preference \succ_i . The algorithm is based on that of Escoffier *et al.* [2008] for testing single-

Algorithm 2 Compute Score or Upper Bound of (Partial) Axis A

```
1:  $V \leftarrow \{\}$  %Set of consistent voters%
2: for agent  $i \in N$  do
3:   consistent  $\leftarrow$  TRUE
4:    $l \leftarrow 1, r \leftarrow n$  %Left and right pointers%
5:   for  $t$  from  $z$  to 1 do
6:     if  $A[l] = \succ_{i,t}$  or unplaced candidate then
7:        $l \leftarrow l + 1$ 
8:     else if  $A[r] = \succ_{i,t}$  or unplaced candidate then
9:        $r \leftarrow r - 1$ 
10:    else
11:      consistent  $\leftarrow$  FALSE
12:      break
13:    if consistent = TRUE then
14:       $V \leftarrow V \cup \{i\}$ 
```

peaked consistency, exploiting the fact that candidates ranked last in any \succ_i must lie at the extreme ends of the axis.

Since Alg. 2 will be called frequently by 1D-SPBB (Alg. 1), its running time should be slight. Fortunately, it is easy to see that this is the case:

Theorem 1 *Algorithm 2 has a running time of $O(nz)$.*

Good heuristics for selecting branches (i.e., partial axes to expand) can have a dramatic impact on any branch-and-bound algorithm: the ability to increase our lower bound quickly can significantly impact the degree of pruning. Our current heuristic simply expands nodes in descending order of their upper bounds, in the hope that a partial axis with a large upper bound will have some completion with a high score, thereby improving our global lower bound. If additional domain-specific information is available, other heuristics may be used. For instance, if a probabilistic prior distribution over voter preferences is known, then the expected degree of consistency can be used to heuristically score nodes for expansion. Other possibilities include expanding nodes that are “least similar” or most likely to be “correct” given the nodes that have already been expanded.

3.2 Approximation

We use the best axis algorithm 1D-SPBB as the core of more general algorithms to find optimal axes under various forms of approximation, and to estimate the degree to which a preference profile is approximately single-peaked. We propose several extensions of 1D-SPBB for three different notions of approximate single-peaked consistency. In some cases, the algorithms do not guarantee discovery of the optimal approximation (i.e., the minimum k), but they provide both upper and lower bounds on the degree of approximation.

k -maverick consistency. Computing the minimum k for which a profile is k -maverick consistent is precisely what 1D-SPBB does. The best axis A^* found by the algorithm explains $s(A^*)$ consistent voters (and this is the maximum number of voters explainable by any axis). Hence, the remaining $n - s(A^*)$ voters form the maverick set of minimum size. Hence, Alg. 1 can be applied without any change.

k -additional axis consistency. The 1D-SPBB algorithm

can also be used to compute k -AA consistency. We investigate a simple greedy algorithm to approximate the minimum k for which a profile \succ is k -AA single-peaked consistent. The algorithm, 1D-SPBB-AA, works as follows: starting with the full profile, we find the best axis A_1 using 1D-SPBB. We then remove all n_1 voters consistent with A_1 from the profile and rerun 1D-SPBB on the profile of the $n - n_1$ remaining voters. We repeat until the profile is empty. If it terminates after $k + 1$ iterations, 1D-SPBB-AA verifies k -AA consistency.

The value k determined by 1D-SPBB-AA is only an upper bound on the minimum k required for k -AA consistency because of its greedy nature: it may not give the optimal k . The *first* iteration of the algorithm also determines a *lower bound* on k : if the first axis returned explains n_1 voters, then $k \geq \lceil \frac{n}{n_1} \rceil$ is needed to ensure k -AA consistency. We exploit this fact below in analyzing our data sets. Deciding if a profile is k -AA consistent is NP-complete [Erdélyi *et al.*, 2012].

k -local candidate deletion consistency. We can readily adapt 1D-SPBB to work with k -LCD consistency. Specifically, given a fixed value of k , we modify the algorithm to compute the best axis, i.e., the axis that renders the greatest number of voters single-peaked when we allow up to k candidates to be deleted from any voter’s ranking. This is useful if we wish to see how single-peaked a profile is when voters, say, make “mistakes” in their ballots, or fail to distinguish certain candidates from one another. We can also combine this with k -AA consistency, finding the number of additional axes needed when each axis is allowed to explain voter preferences using k -LCD.

1D-SPBB can be used directly for this purpose, and requires only a modification in Alg. 2, when computing the upper bound of a partial axis (or score of a complete axis) for k -LCD. Instead of reporting a violation of single-peakedness when \succ_i is inconsistent with the (partial) axis, it records, for each voter i , the number of inconsistencies detected so far—each inconsistency can be managed by a local deletion. If $k + 1$ violations occur, then i is reported as inconsistent with the (partial) axis.

We use this method to find the best axis for fixed values of k in experiments below. We also adapt the greedy algorithm, 1D-SPBB-AA, to find (approximate) the minimal number of additional axes needed when allowing k -LCD: we call this method 1D-SPBB-AA- k -LCD. Since the number of consistent voters is non-decreasing in k , we can use binary search to find the minimum value of k that ensures k -LCD single-peaked consistency w.r.t. the best axis found by Alg. 1. Since we can always make any profile single-peaked by removing $z - 2$ candidates from each voter’s ranking in the worst case, binary search will take at most $\log_2(z - 1)$ iterations. Like k -AA, this problem is NP-complete [Erdélyi *et al.*, 2012] and the algorithm may not find the *minimum number* of local deletions required: this is due to the fact that when a violation occurs, we simply remove the lower-ranked candidate in \succ_i , whereas a deletion of the higher-ranked candidate may have led to a fewer future deletions for voter i . Thus our method returns only an upper bound of the optimal solution.

	Dublin West	Dublin North
# of consistent voters	109 (2.9%)	14 (0.4%)
Complete axis visited	2	3
Branch out due to bound	9,375	509,202
Running time (in sec.)	0.64	2.92

Figure 1: 1-D branch-and-bound results (best single axis).

3.3 Results from 2002 Irish General Election

We applied our algorithms to two data sets taken from the 2002 Irish general election.² The Dublin West election has 9 candidates and 29,989 votes of the top- t form (for varying values of t), of which 3,800 are complete preference rankings. In Dublin North, there are 12 candidates and 43,942 votes, of which 3,662 are complete.³ Our primary experiments are run on the subset of votes comprising all complete rankings. We first ran 1D-SPBB (Alg. 1), combined with Alg. 2, to compute the best single axis for the two data sets. Fig. 1 shows that the *best* axis explains, assuming single-peaked preferences, a tiny fraction of voter preferences, only 109 of 3,800 (2.87%) and 14 of 3,662 (0.38%) in West and North, respectively. Clearly voter preferences are far from being single-peaked.

Our methods can easily be adopted to partial rankings in the obvious fashion (we omit details here). Preliminary results on the full voting data sets, including the truly partial rankings, show that 6% (resp., 6.5%) of voters are single-peaked w.r.t. the best axis in the West (resp., North) data sets. This larger fraction is not unexpected, since single-peakedness cannot be violated by unranked candidates (so voters with top- t preferences for smaller values of t are consistent with far more axes). Despite this, voter preferences remain far from being single-peaked.

We also see that 1D-SPBB is quite efficient. While the total number of axes are $9!/2 \approx 181K$ and $12!/2 \approx 240M$, respectively, the algorithm needs only 0.64s. (resp., 2.92s) to find the best axis, visiting only two (resp., three) complete axes, and branching out 9,375 (resp., 509,202) times, indicating good pruning due to strong lower bounds.

We also investigate the various approximations described above. The (single) best axis results immediately tell us that k -maverick consistency requires deletion of 97.13% (resp., 99.62%) of voters to ensure the preference profile is single-peaked. We also immediately obtain a lower bound on k -AA consistency: Dublin West needs a minimum of $\lceil \frac{3800}{109} \rceil = 35$ additional axes, while North needs 262 additional axes. We also ran the greedy k -LCD approximation algorithm, 1D-SPBB-AA- k -LCD, for different values of k (when $k = 0$, this is just 1D-SPBB-AA). Fig. 2(a) and (b) show the percentage of voters explained with each additional axis added for different values of k (note the log-scale on the x -axis). Without k -LCD approximation (i.e., when $k = 0$), 447 (resp., 1,452) axes are needed to explain all voter preferences (this is an upper bound on k -AA consistency). k -LCD without multiple axes requires $k = 7$ (resp., $k = 10$)—the maximum possible for each data set. Even with k -LCD for reasonable values of k , many additional axes are needed to explain the data: for

²Data sets obtained from: www.dublincountyreturningofficer.com.

³A ranking has top- t form if a voter ranks only his t most-preferred candidates. If $t < m - 1$, the ranking is incomplete.

instance, 31 axes are needed to explain Dublin West when $k = 2$, while Dublin North, with an aggressive $k = 4$, needs 39 axes. The linear nature of the plots (recall the log-scale) also shows that deletion of maverick voters will not help. This suggests that, even allowing for *combinations of approximations* proposed in the literature, preferences in these data sets are very far from being single-peaked in 1D. This motivates the use of higher-dimensional models, to which we now turn.

4 Two-dimensional Branch and Bound

Since voter preferences in the data sets above are not single-peaked in the one-dimensional sense—even when aggressive approximation is considered—the explanatory power of these proposed approximations in 1D is rather limited. We now extend these techniques to two-dimensional (2D) models. Our extensions generalize beyond two dimensions, but we focus on 2D models for ease of presentation, and also because, as we see below, two dimensions suffice for our data sets.

4.1 The Algorithm

Extending our branch-and-bound algorithm to the 2D case presents several challenges. First, the search space explodes, as we must potentially consider all $O((z!)^2)$ combinations of first and second axes. Second, candidates ranked last in some \succ_i need no longer lie at the extreme point of an axis (e.g., in Ex. 1, $c_4 \succ_i c_3 \succ_i c_1 \succ_i c_5 \succ_i c_2$ is a valid single-peaked preference, but the least preferred candidate c_2 is not extreme on either axis). Third, in two-dimensions, some axes are *dominated* by others—these should be pruned for computational efficiency to the greatest extent possible. We now outline a 2D extension of the 1D-SPBB called *2D-SPBB*. We omit full details due to space limitations and focus on the intuitions behind the method. However, we do explain how to tackle each of the issues above.

To address the combinatorial explosion of possible pairs of axes, instead of considering all candidate permutations as possible first axes, we admit only a relatively small set of potential initial axes, or a limited *sample* of possible axes. For each such (potential) first axis, we fix it as our first dimension, then apply our 1D algorithm 1D-SPBB to compute the second dimension. Our implementation uses 1D-SPBB-AA to find the collection of 1D axes that fully explains the given profile \succ —we use this as our set of potential first dimension axes. This guarantees that each \succ_i is single-peaked consistent w.r.t. at least one of the axes.⁴ This way of structuring 2D-SPBB means any axis searched/expanded in the first dimension is always complete, never partial. Of course, this is simply a heuristic, and may limit our ability to find a good 2D model.

Computing scores and upper bounds. To address the second problem, we develop a new algorithm to compute the upper bound for a pair of partial axes in a 2D space (i.e., maximum number of voters that are consistent with some extension of the partial pair), or the score of the pair of axes if they are complete. This includes variants that incorporate the same forms of approximation as above. One key difference

⁴If a voter is single-peaked w.r.t. one axis A , then she is also single-peaked w.r.t. any 2D-space using A as one of its axes.





Figure 2: 1-D branch-and-bound results, with LCD-approximation: (a) Dublin West; (b) Dublin North.

between 2D and 1D lies in the computation of consistency. In a 2D space, the inconsistency of \succ_i with single-peakedness only occurs with the violation of some bounding box constraint (see Ex. 1): our algorithm checks for such violations. Specifically, for each \succ_i , we compute the set of bounding box constraints \mathbf{B} induced by the partial axes. Recall that each constraint $\mathbf{b} = \langle \mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3 \rangle$, involves three candidates: if $\tau_i = \mathbf{b}_1$ (resp., $\tau_i = \mathbf{b}_3$), then we must have $\mathbf{b}_2 \succ_i \mathbf{b}_3$ (resp., $\mathbf{b}_2 \succ_i \mathbf{b}_1$) to ensure single-peakedness of \succ_i . If no constraints are violated, i joins the set of consistent voters.

As in the 1D case, consistency testing must be fast to ensure that nodes in the branch-and-bound tree are processed quickly. Consistency testing is polynomial time:

Theorem 2 *Given a preference profile \succ , the number of voters consistent with a pair of partial axes A_1, A_2 can be computed in $O(nz^4)$ time.*

Proof: The set of bounding box constraints \mathbf{B} can be computed in $O(z^3)$ time, since each constraint involves candidate triples (of which there are at most $\binom{z}{3}$). Testing a ranking \succ_i against each such constraint (as described above) can be accomplished in $O(z)$ time, and must be done at most once for each of n voters. (generally, substantially fewer at deeper nodes in the tree).⁵ Thus total running time is $O(z^3) + O(nc^4) = O(nz^4)$. ■

We mention two important details regarding the computation of \mathbf{B} . First, it can be done incrementally by inheriting bounding box constraints from nodes higher in the search tree, then adding only the new constraints induced by placing two more candidates on the second axis. Second, for any incomplete axis, apart from “explicit” constraints involving candidates on the axis, we can also compute “implicit” constraints. For example, suppose A_1 is fixed, with $c_1 <_{A_1} c_2 <_{A_1} c_3 <_{A_1} c_4 <_{A_1} c_5 <_{A_1} c_6$, while A_2 is partial, with $A_2 = c_1 <_{A_2} c_6 <_{A_2} \dots <_{A_2} c_5 <_{A_2} c_2$. The only explicit constraint is $\langle c_6, c_5, c_2 \rangle$; but four implicit constraints can be added: $\langle c_1, c_3, c_5 \rangle$, $\langle c_1, c_4, c_5 \rangle$, $\langle c_6, c_4, c_2 \rangle$ and $\langle c_6, c_3, c_2 \rangle$. This allows more precise upper bound computation and more aggressive pruning.

Removing dominated axes. The fact that pairs of axes in 2D give rise to bounding box constraints leads to a form of “domination” that can be exploited to further reduce the combinatorial overhead of searching.

⁵Other efficiencies, e.g., caching consistency tests across voters with identical preference orderings, are also possible.

	Dublin West	Dublin North
$k = 0$	2,498/3800	1,732/3,662
$k = 1$	3,553/3800	2,948/3,662
$k = 2$	3,788/3800	3,436/3,662
$k = 3$	3,800/3800	3,645/3,662

Figure 3: 2-D branch-and-bound: number of consistent voters with single best axis using k -LCD approximation.

Definition 5 *A pair of (partial) axes $A = \langle A_1, A_2 \rangle$ is dominated by $A' = \langle A'_1, A'_2 \rangle$ if the set of bounding box constraints induced by A' is a strict subset of that induced by A .*

Consider $A = \langle A_1, A_2 \rangle$, with complete axis $A_1 = c_1 <_{A_1} c_2 <_{A_1} c_3 <_{A_1} c_4 <_{A_1} c_5$ and partial axis $A_2 = c_1 <_{A_2} c_2 \dots <_{A_2} c_5$. $A = \langle A_1, A_2 \rangle$ is dominated by a different pair $A' = \langle A'_1, A'_2 \rangle$: we obtain strictly fewer bounding box constraints by swapping c_1 with whichever candidate happens to be placed in the second position of A_2 , and c_5 with whichever candidate is placed in the fourth position. As such, assuming A_1 is fixed (as we would in a specific branch of 2D-SPBB), a different axis A'_2 offers strictly more flexibility than A_2 .

We exploit this fact by using an algorithm for removing (some, but not all) dominated axes during 2D-SPBB: detecting this can allow pruning of a large part of the branch-and-bound tree. We test domination by checking whether a swap of two adjacent candidates on any axis can induce a strict subset of original constraints (as in the above example): if yes, the (partial) axis is pruned. This simple test is sound; and while it does not ensure pruning of *all* dominated axes, it improves run-time dramatically.

Approximation. As in the 1D case, 2D-SPBB automatically generates the minimal k required for k -maverick consistency. Of course, if we use sampling to limit the axes that will be considered for the first dimension, we will obtain only an upper bound on k . It can also be applied repeatedly to greedily approximate the minimal set of additional 2D “axis pairs” needed to explain a profile; and we easily incorporate k -LCD approximation into 2D-SPBB using similar modifications to those described in Sec. 3.2. We focus on k -LCD below.

4.2 Results from the 2002 Irish General Election

We use the Dublin West and North data sets to test the effectiveness of 2D-SPBB and specifically the ability of k -LCD approximation to fit the Irish voting data. Fig. 3 shows the fraction of voters that are explained by the *best axis pair* generated using 2D-SPBB, both without approximation ($k = 0$),




Figure 4: 2-D branch-and-bound results, anytime performance: (a) Dublin West; (b) Dublin North.

and allowing k -LCD approximation for $k \leq 3$. The contrast with the 1D fit is notable: even without approximation, the best 2D-axis pair explains 65.7% (resp., 47.3%) of all voters. Allowing 2 out of 9 (resp., 3 out of 12) local candidate deletion provides a near-perfect fit, covering 99.68% (resp., 99.54%) of voters. This strongly suggests that the 2D model carries far more explanatory power for this Irish voting data.

The 2D algorithms are more computationally intensive than their 1D counterparts (though restricting attention to sampled axes in the first dimension helps tremendously). It is instructive to examine the anytime performance of 2D-SPBB to see how quickly it converges to the best 2D model, and how quickly model quality improves for various levels of k -LCD. Figs. 4 (a) and (b) show that, while convergence to the best 2D model can take a considerable amount of time, the anytime performance is very good, allowing the discovery of models that capture most of the (explainable) voters extremely quickly (note the log-scale on the x -axis).

5 Conclusions and Future Work

We have developed a branch-and-bound algorithm designed to determine the degree to which a preference profile can be viewed as single-peaked in both the single- and multi-dimensional senses. It uses, and combines, various forms of approximation. Experiments on two election data sets demonstrate clearly that one-dimensional models, for any reasonable degree of approximation, cannot explain voter preferences in the two data sets we have explored. By contrast, a two-dimensional model provides an excellent fit, using very low degrees of local candidate deletion (as the only form of approximation) to explain the preferences of over 99% of the voters in each data set. Our algorithms are very effective in practice in 1D spaces, and feasible in 2D with strong anytime performance, despite the NP-completeness of these problems. While these findings are preliminary, and need to be corroborated on further election and other preference data sets, they suggest that the extension to two (or additional) dimensions may render the use of single-peaked modeling, or its approximations, more applicable in practice.

A number of interesting directions emerge from this work. One is to further develop the theoretical foundations of single-peaked consistency for different forms of approximations, especially as we move to higher dimensions. Erdélyi *et al.* [2012] show that the 1D-problem is NP-complete for several of the forms of approximation considered here. We

conjecture that the problem remains NP-complete in two-dimensions for the approximations discussed in this paper.

Obviously, multi-dimensional single-peakedness is a much weaker assumption than its 1D counterpart; so while it may fit preference data better, its predictive power is lessened. Developing a deeper understanding of these tradeoffs is vital. An interesting question is, for instance, minimum conditions on profiles that prevent the fit of *any* m -dimensional model (c.f. 1D, where single-peakedness can be violated with only two voters and four candidates, or three voters and three candidates [Ballester and Haeringer, 2011]).

Another direction is the exploration of connections with spatial models [Hinich, 1978; Poole and Rosenthal, 1985], which are more restrictive than multi-dimensional single-peaked models in some ways, and more flexible in others. Interestingly, 1D and 2D models are often considered to be sufficient in practice [Poole and Rosenthal, 1985], corroborating our findings in the more qualitative single-peaked framework.⁶ Related is developing connections to the probabilistic models of behavioral social choice [Regenwetter *et al.*, 2006].

Finally, this work as an important step towards the design of mechanisms for approximately single-peaked preferences. While much attention has been paid to mechanisms that exploit single-peakedness [Moulin, 1980; Barberà *et al.*, 1993; Schummer and Vohra, 2007; Procaccia and Tennenholtz, 2009; Lu *et al.*, 2010; Escoffier *et al.*, 2011; Dokow *et al.*, 2012; Sui *et al.*, 2013; Fotakis and Tzamos, 2010], little work has addressed the impact of approximate single-peakedness on these mechanisms, or the design and analysis of mechanisms specifically for approximate single-peakedness. Having a sense of which forms of approximation best fit real-world preferences can help focus mechanism design efforts on those most likely to have a practical impact.

Acknowledgements: Thanks to the participants of COMSOC-12 for lively discussions and the anonymous reviewers for their comments. This research was supported by NSERC.

References

[Ballester and Haeringer, 2011] M. Angel Ballester and G. Haeringer. A characterization of the single-peaked do-

⁶Our own preliminary experiments with spatial models show that 2D models fit the Irish voting data well.

- main. *Social Choice and Welfare*, 36(2):305–322, 2011.
- [Barberà *et al.*, 1993] S. Barberà, F. Gul, and E. Stacchetti. Generalized median voter schemes and committees. *Journal of Economic Theory*, 61(2):262–289, 1993.
- [Barberà, 2010] S. Barberà. Strategy-proof social choice. In K. J. Arrow, A. K. Sen, and K. Suzumura, eds., *Handbook of Social Choice and Welfare*, volume 2, pp.731–832. North-Holland, Amsterdam, 2010.
- [Bartholdi and Trick, 1986] J. Bartholdi and M. A. Trick. Stable matching with preferences derived from a psychological model. *Operations Research Letters*, 5(4):165–169, 1986.
- [Black, 1948] D. Black. On the rationale of group decision-making. *Journal of Political Economy*, 56(1):23–34, 1948.
- [Bradley and Terry, 1952] R. A. Bradley and M. E. Terry. Rank analysis of incomplete block designs. *Biometrika*, 39:324–345, 1952.
- [Conitzer, 2009] V. Conitzer. Eliciting single-peaked preferences using comparison queries. *Journal of Artificial Intelligence Research*, 35:161–191, 2009.
- [Dokow *et al.*, 2012] E. Dokow, M. Feldman, R. Meir, and I. Nehama. Mechanism design on discrete lines and cycles. In *Proc. 13th ACM Conference on Electronic Commerce (EC’12)*, pp.423–440, Valencia, Spain, 2012.
- [Erdélyi *et al.*, 2012] G. Erdélyi, M. Lackner, and A. Pfandler. The complexity of nearly single-peaked consistency. In *Proc. 4th Intl. Workshop on Computational Social Choice (COMSOC-2012)*, Kraków, Poland, 2012.
- [Escoffier *et al.*, 2008] B. Escoffier, J. Lang, and M. Öztürk. Single-peaked consistency and its complexity. In *Proc. 18th European Conference on Artificial Intelligence (ECAI-08)*, pp.366–370, Patras, Greece, 2008.
- [Escoffier *et al.*, 2011] B. Escoffier, L. Gourvès, T. N. Kim, F. Pascual, and O. Spanjaard. Strategy-proof mechanisms for facility location games with many facilities. In *Proc. 2nd Intl. Conference on Algorithmic Decision Theory (ADT-11)*, pp.67–81, Piscataway, NJ, 2011.
- [Faliszewski *et al.*, 2011] P. Faliszewski, E. Hemaspaandra, and L. A. Hemaspaandra. The complexity of manipulative attacks in nearly single-peaked electorates. In *Proc. 13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-11)*, pp.228–237, Groningen, The Netherlands, 2011.
- [Fotakis and Tzamos, 2010] D. Fotakis and C. Tzamos. Winner-imposing strategyproof mechanisms for multiple facility location games. In *Proc. 6th Intl. Workshop on Internet and Network Economics (WINE-10)*, pp.234–245, Stanford, CA, 2010.
- [Galand *et al.*, 2012] L. Galand, D. Cornaz, and O. Spanjaard. Bounded single-peaked width and proportional representation. In *Proc. 4th Intl. Workshop on Computational Social Choice (COMSOC-2012)*, Kraków, Poland, 2012.
- [Gibbard, 1973] A. Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, 1973.
- [Gormley and Murphy, 2007] I. C. Gormley and T. B. Murphy. A latent space model for rank data. In E. Airoldi, D. Blei, S. Fienberg, A. Goldenberg, E. Xing, and A. Zheng, eds., *Statistical Network Analysis: Models, Issues, and New Directions*, pp.90–102. Springer, 2007.
- [Hinich, 1978] M. J. Hinich. Some evidence on non-voting models in the spatial theory of electoral competition. *Public Choice*, 33(2):83–102, 1978.
- [Hotelling, 1929] H. Hotelling. Stability in competition. *The Economic Journal*, 39(153):41–57, 1929.
- [Lu *et al.*, 2010] P. Lu, X. Sun, Y. Wang, and Z. A. Zhu. Asymptotically optimal strategy-proof mechanisms for two-facility games. In *Proc. 11th ACM Conference on Electronic Commerce (EC’10)*, pp.315–324, Cambridge, MA, 2010.
- [Luce, 1959] R. D. Luce. *Individual Choice Behavior: A Theoretical Analysis*. Wiley, 1959.
- [Moulin, 1980] H. Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4):437–455, 1980.
- [Poole and Rosenthal, 1985] K. T. Poole and H. Rosenthal. A spatial model for legislative roll call analysis. *American Journal of Political Science*, 29(2):357–384, 1985.
- [Procaccia and Tennenholtz, 2009] A. D. Procaccia and M. Tennenholtz. Approximate mechanism design without money. In *Proc. 10th ACM Conference on Electronic Commerce (EC’09)*, pp.177–186, Stanford, CA, 2009.
- [Regenwetter *et al.*, 2006] M. Regenwetter, B. Grofman, A. A. J. Marley, and I. Tsetlin. *Behavioral Social Choice: Probabilistic Models, Statistical Inference, and Applications*. Cambridge University Press, Cambridge, 2006.
- [Satterthwaite, 1975] M. A. Satterthwaite. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975.
- [Schummer and Vohra, 2007] J. Schummer and R. V. Vohra. Mechanism design without money. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, eds., *Algorithmic Game Theory*, pp.243–265. Cambridge University Press, 2007.
- [Shepard, 1959] R. N. Shepard. Stimulus and response generalization: A stochastic model relating generalization to distance in psychological space. *Psychometrika*, 22(4):325–345, 1959.
- [Sui *et al.*, 2013] X. Sui, C. Boutilier, and T. Sandholm. Analysis and optimization of multi-dimensional percentile mechanisms. In *Proc. 23rd Intl. Joint Conference on Artificial Intelligence (IJCAI-13)*, Beijing, 2013. to appear.