

Spatial Voting with Incomplete Voter Information

Aviram Imber¹, Jonas Israel², Markus Brill^{2,3}, Hadas Shachnai¹, Benny Kimelfeld¹

¹Technion – Israel Institute of Technology, Haifa, Israel

²Research Group Efficient Algorithms, TU Berlin, Germany

³Department of Computer Science, University of Warwick, UK

aviram.imber@cs.technion.ac.il, j.israel@tu-berlin.de, markus.brill@warwick.ac.uk, hadas@cs.technion.ac.il, bennyk@technion.ac.il

Abstract

We consider spatial voting where candidates are located in the Euclidean d -dimensional space, and each voter ranks candidates based on their distance from the voter’s ideal point. We explore the case where information about the location of voters’ ideal points is incomplete: for each dimension, we are given an interval of possible values. We study the computational complexity of finding the possible and necessary winners for positional scoring rules. Our results show that we retain tractable cases of the classic model where voters have partial-order preferences. Moreover, we show that there are positional scoring rules under which the possible-winner problem is intractable for partial orders, but tractable in the one-dimensional spatial setting. We also consider approval voting in this setting. We show that for up to two dimensions, the necessary-winner problem is tractable, while the possible-winner problem is hard for any number of dimensions.

1 Introduction

In the spatial model of voting (Enelow and Hinich 1984; Miller 1995), both candidates and voters are associated with points in the d -dimensional Euclidean space $\mathbb{R}^d$. It is assumed that the locations of candidates and voters correspond to their “ideal points” and that each voter’s preferences over the candidates can be inferred from the Euclidean distance between the candidates’ and the voter’s ideal points. For example, the location of a candidate or voter in $\mathbb{R}^d$ could reflect the stance (or opinion) of the candidate or voter regarding d different *issues* that are relevant for the election. In the social choice literature, preferences with this structure are often referred to as *(d-)Euclidean preferences* (Bogomolnaia and Laslier 2007; Elkind, Lackner, and Peters 2022).

We consider a setting where only partial information about the preferences of voters is available. In such a setting, the exact preference order of a voter is unknown but assumed to come from a known space of possible preference orders. Each combination of possible preference orders is a possible voting profile that may result in different sets of winners (given a fixed voting rule). Natural computational tasks that arise in such scenarios ask about the *possible winners* (who win in at least one possible profile) and the *necessary winners* (who win in every possible profile) (Lang

2020). A prominent manifestation of this idea is the seminal framework of Konczak and Lang (2005), in which voter preferences are specified as partial orders and a possible profile is obtained by extending each partial order into a total preference order. A thorough picture of the complexity of the possible and necessary winner problems has been established in a series of studies (Betzler and Dorn 2010; Xia and Conitzer 2011; Baumeister and Rothe 2012). For example, under every positional scoring rule in the setting of partial orders, the necessary winners can be found in polynomial time, yet it is NP-complete to decide whether a candidate is a possible winner (assuming a regularity condition), except for the tractable cases of the plurality and veto rules.

We study the complexity of the computational problems $PW\langle d \rangle$ and $NW\langle d \rangle$, where the goal is to find the possible and necessary winners, respectively, when we have incomplete information about voters’ ideal points in spatial voting with d dimensions. More precisely, instead of the ideal points, we are given — for each voter and dimension — an interval of possible values for the voter’s opinion. Hence, each voter is associated with a space of possible ideal points. Different points in this space may induce different preference orders over the candidates (whose locations are assumed to be known precisely). We thus get a mechanism for defining a space of possible total orders that is different¹ from the classical partial-order setting (Konczak and Lang 2005). We refer to our setting as *partial spatial voting*.

We first focus on the class of positional scoring rules and compare the computational complexity of the possible and necessary winner problems to the classic model of partial orders. We investigate the following questions: (1) Is the necessary-winner problem still tractable for all positional scoring rules? (2) Is the possible-winner problem still tractable for plurality and veto? (3) Are there positional scoring rules where the possible-winner problem is tractable for partial spatial voting but not for partial orders? We answer all three questions positively. For some of our results, we uncover and exploit an interesting relationship between the possible-winner problem and *scheduling*.

We then consider spatial *approval* voting with incomplete information. We provide an efficient algorithm for computing necessary winners in one and two dimensions and prove

¹In Section 3, we show that the two settings are incomparable.

Problem	Positional scoring rules	Approval voting
NW $\langle d \rangle$	in P [Theorem 2]	in P for $d \leq 2$ [Theorem 8]
PW $\langle 1 \rangle$	in P for all two-valued rules [Theorems 3 and 5], weighted veto, and $F(k, t)$ whenever $k > t$ [Theorem 4]	NP-complete [Theorem 9]
PW $\langle d \rangle$	in P for plurality and veto [Theorem 5]; NP-complete for k -approval with $d \geq 2$ and $k \geq 3$ [Theorem 6]	NP-complete [Theorem 9]

Table 1: Complexity results for computing the necessary and possible winners in the studied models of uncertainty.

that computing possible winners is intractable, even for one dimension. Our results are summarized in Table 1. Omitted proofs can be found in the full version of this paper (Imber et al. 2023).

Related work. Spatial voting in one dimension is intuitively similar to assuming *single-peaked preferences* (Black 1948; Arrow 1951). Yet, there are considerable differences, as single-peaked preferences do not impose any restrictions on the comparison between candidates on different sides of the peak. Walsh (2007) showed hardness results for possible and necessary winner questions under single-peaked (but not necessarily 1-Euclidean) preferences for STV and polynomial-time results for Condorcet-consistent rules.


Bogomolnaia and Laslier (2007) showed that *every* (complete) preference profile can be represented in the spatial model, by choosing the dimension d to be sufficiently large. Given a preference profile, it can be efficiently decided whether the profile can be represented as a one-dimensional spatial model (Doignon and Falmagne 1994; Knoblauch 2010; Elkind and Faliszewski 2014); for higher dimensions, the problem becomes intractable (Peters 2017). Jamieson and Nowak (2011) studied the problem of learning the ranking of candidates in spatial voting using pairwise comparisons. They established a bound on the number of possible rankings; we use this bound in Section 4. Barrot et al. (2013) and Imber et al. (2022) study the possible-winner and necessary-winner problems for approval voting in single-winner and multi-winner elections, respectively, using preference models similar to partial orders.

The problems considered in this paper also relate to manipulation and control problems that involve reasoning about a space of possibilities of profiles. Lu et al. (2019) study control where a party can select a subset of issues to focus on. Estornell et al. (2020) study manipulation of spatial voting where the issues are weighted and a malicious attacker can change the weights. Wu et al. (2022) study manipulation where the adversary can change the position of a candidate.

2 Preliminaries

We first introduce the basic concepts and notation that we use throughout the paper.

Voting profiles. Let $C = \{c_1, \dots, c_m\}$ be a set of $m \geq 2$ candidates and $V = \{v_1, \dots, v_n\}$ a set of voters. A *ranking profile* $\mathbf{R} = (R_1, \dots, R_n)$ consists of n linear orders over C . Each R_i represents the preference order of voter v_i .

Figure 1: Example of a spatial voting profile with $d = 2$: a set $C = \{c_1, c_2, c_3\}$ of candidates and a single voter v .

In the *spatial voting* setting (Enelow and Hinich 1984), each candidate is associated with a d -dimensional vector corresponding to its positions (opinions) on d issues, denoted by $\mathbf{c}_i = \langle c_{i,1}, \dots, c_{i,d} \rangle \in \mathbb{R}^d$. For simplicity, we assume that all candidates have distinct positions, that is, there are no perfect clones.

A *spatial voting profile* $\mathbf{T} = (T_1, \dots, T_n)$ consists of a vector $T_j = \langle T_{j,1}, \dots, T_{j,d} \rangle$ for each voter v_j , representing the voter's positions on the d issues. Given a spatial profile $\mathbf{T}$, we can construct a ranking profile $\mathbf{R}_{\mathbf{T}} = (R_{T_1}, \dots, R_{T_n})$ where each voter v_j ranks candidates in C according to their Euclidean distance $\|T_j - \mathbf{c}_i\|_2$ from v_j . The closest candidate is ranked first, and the farthest is ranked last (position m) in v_j 's preferences. We break ties by a linear order over the candidates, which is given as part of the input for each voter. An example of a spatial voting profile and its associated ranking profile is illustrated in Figure 1. By a slight abuse of terminology, we may identify voters with their points in $\mathbb{R}^d$, and we use the terms *dimension* and *issue* interchangeably.

Voting rules. A *voting rule* is a function that maps each ranking profile to a set of *winners*. A *positional scoring rule* r is a series $\{\vec{s}_m\}_{m \geq 2}$ of m -dimensional score vectors $\vec{s}_m = (\vec{s}_m(1), \dots, \vec{s}_m(m))$ of natural numbers, where $\vec{s}_m(1) \geq \dots \geq \vec{s}_m(m)$ and $\vec{s}_m(1) > \vec{s}_m(m)$. We assume that $\vec{s}_m(j)$ is computable in polynomial time in m , and the scores in each $\vec{s}_m$ are mutually co-prime (i.e., their greatest common divisor is one). Some examples of positional scoring rules include the *plurality* rule $(1, 0, \dots, 0)$, the *k-approval* rule $(1, \dots, 1, 0, \dots, 0)$ that begins with k ones, the *veto* rule $(1, \dots, 1, 0)$, and the *k-veto* rule that ends with k zeros. A positional scoring rule is *pure* if every $\vec{s}_{m+1}$ can be obtained from $\vec{s}_m$ by inserting a score value at some position (while satisfying $\vec{s}_{m+1}(1) \geq \dots \geq \vec{s}_{m+1}(m+1)$).

Given a ranking profile $\mathbf{R} = (R_1, \dots, R_n)$ a positional scoring rule r , the score $s_r(R_i, c)$ that the voter v_i contributes to the candidate c is $\vec{s}_m(j)$, where j is the position of

c in R_i . The score of c in $\mathbf{R}$ is $s_r(\mathbf{R}, c) = \sum_{i=1}^n s_r(R_i, c)$, which we also denote as $s(\mathbf{R}, c)$ if r is clear from context. A candidate c is a *winner* if $s_r(\mathbf{R}, c) \geq s_r(\mathbf{R}, c')$ for all candidates c' . The set $r(\mathbf{R})$ contains all winners.

A positional scoring rule r is *two-valued* if there are only two values in each $\vec{s}_m$. For such rules, we assume, w.l.o.g., that $\vec{s}_m$ consists only of zeros and ones, and hence is of the form $\vec{s}_m = (1, \dots, 1, 0, \dots, 0)$. Thus, we can denote any two-valued rule as k -approval, where $k = k(m)$ may depend on m . For example, $(m-2)$ -approval is the same as 2-veto.

For k -approval, we can naturally convert a ranking profile $\mathbf{R} = (R_1, \dots, R_n)$ to an *approval profile* $\mathbf{A} = (A_1, \dots, A_n)$, where each $A_i \subseteq C$ consists of the first k candidates in the order R_i . In other words, A_i denotes the set of candidates that the voter v_i “approves.” The score $s(A_i, c)$ that the voter v_i contributes to the candidate c is one if $c \in A_i$ and zero otherwise. The winners then are the candidates with the maximal score $s(\mathbf{A}, c) = \sum_{i=1}^n s(A_i, c)$.

Incomplete profiles. Throughout this paper, we study problems where voter preferences are incompletely specified, and we are interested in “possible” and “necessary” winners. Abstractly speaking, an incomplete voting profile is simply a set $\mathcal{R}$ of ranking profiles. Given $\mathcal{R}$, a candidate c is called a *possible winner* w.r.t. a voting rule r if c is a winner in at least one profile $\mathbf{R} \in \mathcal{R}$, and a *necessary winner* w.r.t. r if c is a winner in every profile $\mathbf{R} \in \mathcal{R}$. In contrast to possible winners, necessary winners may not exist.

Incomplete profiles give rise to challenging computational problems when they are represented in a compact manner. For example, Konczak and Lang (2005) use a partial order over the candidates to represent the set of linear extensions, as we explain next. In the following section, we introduce another compact representation and compare it to their one.

Partial order profiles. A *partial order profile* $\mathbf{P} = (P_1, \dots, P_n)$ consists of n partial orders (reflexive, anti-symmetric, transitive relations) on the set C of candidates, where each P_i represents the incomplete preferences of voter v_i . A *ranking completion* of $\mathbf{P}$ is a ranking profile $\mathbf{R} = (R_1, \dots, R_n)$ where each R_i is a completion (i.e., linear extension) P_i . As said above, a candidate c is a necessary winner if c is a winner in every ranking completion $\mathbf{R}$ of $\mathbf{P}$, and c is a possible winner if there exists a ranking completion $\mathbf{R}$ of $\mathbf{P}$ where c is a winner.

For a positional scoring rule r , the decision problems PW_{po} and NW_{po} (where “po” stands for “partial order”) are those of determining, given a set C of candidates, a partial order profile $\mathbf{P}$ and a candidate $c \in C$, whether c is a possible winner and a necessary winner, respectively. A classification of the complexity of these problems for positional scoring rules was established in a sequence of publications.

Theorem 1 Betzler and Dorn (2010); Xia and Conitzer (2011); Baumeister and Rothe (2012). NW_{po} can be solved in polynomial time for every positional scoring rule. PW_{po} is solvable in polynomial time for plurality and veto; for all other pure positional scoring rules, PW_{po} is NP-complete.

3 The Model of Partial Spatial Voting

We introduce a model of incompleteness for spatial voting. A *partial spatial profile* $\mathbf{P} = (P_1, \dots, P_n)$ consists of a vector $P_j = \langle [\ell_{j,1}, u_{j,1}], \dots, [\ell_{j,d}, u_{j,d}] \rangle$ for every voter v_j . Each pair $[\ell_{j,i}, u_{j,i}]$ represents a closed interval of possible values for the position of v_j on issue i : $\ell_{j,i}$ is a lower bound, and $u_{j,i}$ is an upper bound. Note that the positions of the voters are incompletely specified, but those of the candidates are known precisely. Let $[n]$ denote $\{1, \dots, n\}$. A spatial voting profile $\mathbf{T} = (T_1, \dots, T_n)$ is a *spatial completion* of $\mathbf{P}$ if $T_{j,i} \in [\ell_{j,i}, u_{j,i}]$ for every voter v_j and issue $i \in [d]$. We can then compute a ranking profile $\mathbf{R}_{\mathbf{T}}$ as before.²

We call a ranking profile $\mathbf{R}$ a *ranking completion* of $\mathbf{P}$ if there exists a spatial completion $\mathbf{T}$ such that $\mathbf{R} = \mathbf{R}_{\mathbf{T}}$. For k -approval, it will be useful to convert the ranking profile to an approval profile, as described in Section 2. We call an approval profile $\mathbf{A}$ an *approval completion* of $\mathbf{P}$ if there exists a spatial profile $\mathbf{T}$ such that $\mathbf{R}_{\mathbf{T}}$ is converted to $\mathbf{A}$.

Again, given a partial spatial profile $\mathbf{P}$, a candidate is a necessary winner if it is a winner in every ranking completion $\mathbf{R}$ of $\mathbf{P}$, and a possible winner if there exists a ranking completion $\mathbf{R}$ of $\mathbf{P}$ where it is a winner. For a positional scoring rule r and dimension d , we consider the decision problems where we are given a set C of d -dimensional candidates, a partial spatial profile $\mathbf{P}$, and a candidate $c \in C$, and we need to determine whether c is a possible or a necessary winner. We denote these two problems by $\text{PW}\langle d \rangle$ and $\text{NW}\langle d \rangle$, respectively. Note that the number of dimensions is fixed and not part of the input for the problem.

Partial spatial vs. partial order profiles. Before we move on, we make a note on the expressiveness of partial spatial voting compared to partial-order profiles. We say that a partial profile $\mathbf{P}$ (in one of the two models) can be *expressed* by the other model if there is a partial profile $\mathbf{P}'$ in the other model with the same set of ranking completions. In the case of full information, every (complete) profile can be expressed by a spatial profile with $d \leq \min\{m, n\}$ dimensions (Bogomolnaia and Laslier 2007).

For every number d of issues, we can easily come up with partial-order profiles (and even complete ranking profiles) that cannot be expressed as (partial) spatial profiles, simply by using the property that in spatial voting all voters must respect the positions of the candidates, where in partial orders each voter can have a completely different structure. For example, if $d = 1$ then preferences will be single-peaked. Moreover, a partial order can have strictly more ranking completions than the upper bound of a partial spatial order. Indeed, while a partial order can have $\Omega(m!)$ completions, Jamieson and Nowak (2011) showed that a partial spatial voter has $\mathcal{O}(dm^{2d})$ completions (see Lemma 1).

On the other hand, consider an instance with three candidates $C = \{c_1, c_2, c_3\}$, one-dimensional positions $c_1 = 1$, $c_2 = 2$, $c_3 = 3$, and a single voter with $P = [1, 3]$. The

²We can model the scenario where we do not know anything about voter v_j ’s position regarding issue i by setting $\ell_{j,i} = -\infty$ and $u_{j,i} = +\infty$. Our algorithms can handle this case efficiently through minor modifications.

voter has four ranking completions: (c_1, c_2, c_3) , (c_2, c_1, c_3) , (c_2, c_3, c_1) , and (c_3, c_2, c_1) . It is easy to verify that there is no partial order with this set of ranking completions.

We conclude that complexity results on possible or necessary winners for the partial-order model do not immediately imply results for the partial spatial model, and vice versa, since neither of the two models generalizes the other.

4 Computing Necessary Winners

In this section, we show that the necessary-winner problem can be solved in polynomial time, similarly to NW_{po} (as stated in Theorem 1), for every positional scoring rule and for every fixed number of dimensions.

Theorem 2. *Let $d \geq 1$ be fixed. $NW\langle d \rangle$ is solvable in polynomial time for every positional scoring rule.*

The remainder of this section is devoted to proving Theorem 2. To determine whether a candidate c is a necessary winner for a given partial spatial profile $\mathbf{P}$, we use the same concept from the algorithm for NW_{po} given partial orders (Xia and Conitzer 2011): a candidate c is *not* a necessary winner if and only if there exists another candidate c' and a ranking completion $\mathbf{R}$ where $s(\mathbf{R}, c') > s(\mathbf{R}, c)$. To this end, we iterate over every other candidate c' and compute the maximal score difference $s(\mathbf{R}, c') - s(\mathbf{R}, c)$ among the ranking completions $\mathbf{R}$ of $\mathbf{P}$. Observe that it suffices to consider each voter $v_j \in V$ separately and compute the maximal score difference $s(R_j, c') - s(R_j, c)$ among the ranking completions R_j of P_j , since we can sum these values to obtain the maximal value of $s(\mathbf{R}, c') - s(\mathbf{R}, c)$. Then, c is not a necessary winner if and only if the maximal score difference is positive for some candidate c' .

The difference between our algorithm and the one for partial orders is in the way we compute the maximal score difference for each voter. We show that for a partial spatial voter, we can enumerate all ranking completions in polynomial time and compute the score difference in each ranking. This is impossible for partial orders, where the number of ranking completions can be exponential in m .

Next, we prove that we can enumerate the ranking completions of a single voter in polynomial time. We will use a geometric concept from the proof of Jamieson and Nowak (2011) of a polynomial bound on the number of possible complete rankings of a given vote.

Lemma 1 Jamieson and Nowak 2011. *At most $\mathcal{O}(m^{2d})$ rankings can be constructed from spatial votes over a given sequence $c_1, \dots, c_m$ of d -dimensional candidates.*

The proof of Lemma 1 relies on the following main idea. Every pair c, c' of candidates corresponds to a $(d-1)$ -dimensional hyperplane partitioning $\mathbb{R}^d$ into two d -faces (halfspaces): points closer to c , and points closer to c' . The set of all (at most) $\binom{m}{2} = m(m-1)/2$ hyperplanes then partitions $\mathbb{R}^d$ into a set Φ of regions, as illustrated in Figure 2. There is a one-to-one correspondence between these regions and the possible rankings: a face $\varphi \in \Phi$ consists of exactly those points in $\mathbb{R}^d$ where the ranking of candidates in C according to distance does not change, and no two of these faces correspond to the same ranking of candidates.




Figure 2: An illustration of the proof of Lemma 1 for $d = 2$ and $C = \{c_1, c_2, c_3\}$. A voter can be positioned at any point in the rectangle $[\ell_1, u_1] \times [\ell_2, u_2]$. Each line $H_{i,j}$ partitions $\mathbb{R}^2$ into 2 regions: the points closer to c_i , and the points closer to c_j . The top-left region — above $H_{1,2}$ and to the left of $H_{2,3}$ and $H_{1,3}$ — corresponds to the possible positions of the voter where the preference ranking equals $c_1 \succ c_2 \succ c_3$.

Lemma 1 implies that the number of ranking completions of a partial spatial vote is at most $\mathcal{O}(m^{2d})$. However, the bound itself is not enough for proving Theorem 2; we need to explicitly construct (and not just count) all feasible completions in order to parse them in the computation of the maximal score difference. Next, we explain how this is done.

Lemma 2. *Let C be a set of m d -dimensional candidates and $P = \langle [\ell_1, u_1], \dots, [\ell_d, u_d] \rangle$ a d -dimensional partial spatial voter. The set of ranking completions of P can be enumerated in polynomial time.*

Proof. The enumeration algorithm uses the geometric interpretation described above. A pseudocode of the algorithm is given in the full version of the paper. Given the candidates C , we compute the corresponding set H of at most $m(m-1)/2$ hyperplanes. To represent the arrangement of these hyperplanes, i.e., of the geometric relation of the (d) -faces spanned by the hyperplanes, we construct an incidence graph $G(H)$. It consists of a node for each face of the arrangement, i.e., a node for each point, line(-segment), plane(-segment), etc., where two or more hyperplanes intersect. Furthermore, two nodes are connected by an edge if the corresponding faces are incident, i.e., one is contained in the other. Using an algorithm of Edelsbrunner, O'Rourke, and Seidel (1986), $G(H)$ can be constructed in $\mathcal{O}(m^{2d})$ time.

We iterate over the nodes in $G(H)$ that correspond to d -faces of the arrangement. By Lemma 1, there are at most $\mathcal{O}(m^{2d})$ such nodes. For each node x be such a node. By considering all $(d-1)$ -faces incident to x , as well as $\mathbf{P}$, we represent the intersection of the d -face with $\mathbf{P}$ as a set of $\mathcal{O}(m^2 + d)$ linear inequalities with d variables. We can then determine whether the d -face and $\mathbf{P}$ intersect, by checking the feasibility of a linear program with the aforementioned set of constraints, which can be done in polynomial time. If there is a feasible solution (i.e., a point in the intersection), we compute the ranking from that point. $\square$

This completes the proof of Theorem 2. Note that the running time of our algorithm is exponential in d , since we enumerate all ranking completions of a voter. We can show that

for plurality and veto, one can find the necessary winners in time polynomial in d (and in n and m). We do so by a reduction to NW_{po} , which is tractable (Betzler and Dorn 2010). Hence, the necessary winners can be found in polynomial time for plurality and veto, even if d is part of the input.

5 Computing Possible Winners

We now turn to the problem of computing possible winners and show that, for some positional scoring rules, this problem is closely related to a scheduling problem.

5.1 The Case of a Single Dimension

We start with the one-dimensional case ($d = 1$) and study the complexity of $\text{PW}\langle 1 \rangle$. In this case, every candidate c is associated with a single real value c . We assume without loss of generality that $c_1 < \dots < c_m$. A partial profile $\mathbf{P}$ consists of a pair $P_j = [\ell_j, u_j]$ for every voter v_j .

For partial orders, finding the possible winners is NP-complete except for plurality and veto (see Theorem 1). In spatial voting, we are able to provide efficient algorithms for multiple well-studied classes of scoring rules. We begin our investigation with two-valued scoring rules, and then prove tractability for other classes of positional scoring rules.

Two-Valued rules. We begin with two-valued (one/zero) scoring rules. The simplest and most well-known rules in this class are plurality and veto.³

We show that $\text{PW}\langle 1 \rangle$ is tractable for any two-valued rule. Recall that we denote a two-valued rule as k -approval for $k = k(m)$. We introduce an alternative definition for partial spatial profiles for k -approval, in the case of a single dimension. Let $\mathbf{P} = (P_1, \dots, P_n)$ be a partial spatial profile where every voter v_j is associated with a pair $P_j = [\ell_j, u_j]$. Since we assume $c_1 < \dots < c_m$, the set of candidates that v_j possibly approves in a completion of P_j is a sequence $(c_{i_j}, c_{i_j+1}, \dots, c_{i_j+t})$ of consecutive candidates. (We can find this sequence in polynomial time using Lemma 2.) In every completion, the candidates that v_j approves form a substring of length k of $(c_{i_j}, c_{i_j+1}, \dots, c_{i_j+t})$; moreover, every such substring is the approval set of some completion. Hence, we can define a partial spatial profile $\mathbf{P} = (P_1, \dots, P_n)$ for k -approval as follows. Each voter v_j is associated with a sequence $P_j = (c_\ell, c_{\ell+1}, \dots, c_u)$ of at least k consecutive candidates. In an approval completion $\mathbf{A} = (A_1, \dots, A_n)$ of $\mathbf{P}$, the set A_j is a substring of length k of P_j . We then use $\mathbf{A}$ to compute the scores of the candidates and select the winners.

With this definition, we solve $\text{PW}\langle 1 \rangle$ in polynomial time for k -approval. We use a reduction to scheduling with arrival times and deadlines, defined as follows.

Definition 1 Non-preemptive multi-machine scheduling with arrival times and deadlines. We are given a set $M = \{M_1, \dots, M_t\}$ of identical machines and a set $\mathbf{J} = \{J_1, \dots, J_n\}$ of n jobs. Each job J_j has an arrival time a_j , a deadline d_j , and processing time p_j . We assume that

³We later show that for plurality and veto, $\text{PW}\langle d \rangle$ is solvable in polynomial time for every fixed $d \geq 1$. In particular, the tractability of $\text{PW}\langle 1 \rangle$ follows from both Theorem 3 and Theorem 5.

$a_j, d_j, p_j \in \mathbb{N}$. A feasible schedule is a mapping $f: \mathbf{J} \rightarrow \mathbb{R} \times M$ that maps each $J_j \in \mathbf{J}$ to a pair $f(J_j) = (s_j, h_j)$ such that the following properties hold:

1. Every job is processed between its arrival time and deadline: $a_j \leq s_j \leq d_j - p_j$ for all $j \in [n]$.
2. Each machine runs at most one job at any time: if $h_i = h_j$ for $i, j \in [n]$ with $i \neq j$ and $s_i \leq s_j$, then $s_j \geq s_i + p_i$.

Since the arrival times, deadlines, and processing times are all integers, we may assume w.l.o.g. that the starting time of every job in a feasible schedule is also an integer.⁴ We now present the algorithm for $\text{PW}\langle 1 \rangle$.

Theorem 3. $\text{PW}\langle 1 \rangle$ is solvable in polynomial time under every two-valued positional scoring rule.

Proof sketch. Let $k = k(m)$. We are given a set C of candidates, a candidate $c \in C$, and a partial spatial profile $\mathbf{P}$. We start by constructing another partial spatial profile $\mathbf{P}'$ with the following properties: (i) In $\mathbf{P}'$, each voter either necessarily approves c or never approves c ; and (ii) c is a possible winner in $\mathbf{P}'$ if and only if c is a possible winner in $\mathbf{P}$.

Next, we show a reduction from deciding whether c is a possible winner in $\mathbf{P}'$ to multi-machine scheduling where all jobs have the same processing time k . Deciding whether a feasible schedule exists in this setting is solvable in polynomial time (Vakhania 2012). Let V_c be the set of voters who necessarily approve c in $\mathbf{P}'$ (and the voters of $V \setminus V_c$ never approve c). Note that $s(\mathbf{A}, c) = |V_c|$ in every approval completion $\mathbf{A}$ of $\mathbf{P}'$. In the reduction, the number of machines is $|V_c|$. For each voter $v_j \in V$ with $P_j = (c_\ell, \dots, c_u)$, define a job J_j with arrival time ℓ , deadline $u+1$, and processing time k . It can be shown that $\mathbf{P}'$ can be constructed in polynomial time, and that c is a possible winner in $\mathbf{P}'$ w.r.t. k -approval if and only if there exists a feasible schedule of all jobs. $\square$

Beyond two-valued rules. We consider two families of rules with more than two values. We refer to the first family as *weighted veto rules*. These rules are of the form $\vec{s}_m = (\alpha, \dots, \alpha, \beta_1, \dots, \beta_k)$ for $\alpha > \beta_1 \geq \dots \geq \beta_k$ and $k < m/2$. The condition $k < m/2$ implies that each voter assigns the highest score α to more than half of the candidates. Moreover, for two positive integers k and t , we denote by $F(k, t)$ the three-valued rule with scoring vector $\vec{s}_m = (2, \dots, 2, 1, \dots, 1, 0, \dots, 0)$ that begins with k occurrences of two and ends with t zeros. For example, the scoring vector for $F(2, 1)$ is $\vec{s}_m = (2, 2, 1, \dots, 1, 0)$.

Theorem 4. $\text{PW}\langle 1 \rangle$ is solvable in polynomial time under every weighted veto rule, and under $F(k, t)$ whenever $k > t$.

5.2 The Case of Multiple Dimensions

For $d > 1$, we show that the tractable cases of possible winners for the partial orders model, namely plurality and veto, are also tractable for the spatial model.

Theorem 5. $\text{PW}\langle d \rangle$ is solvable in polynomial time for plurality and veto.

⁴Otherwise, we can iterate over the jobs, sorted from the smallest starting time to the largest, and change the starting time s_i of job J_i to $\lfloor s_i \rfloor$, without harming the feasibility.




Figure 3: An example of two voters in a completion of the partial profile $\mathbf{P}$ from the proof of Theorem 6. The voter v represents a job of length k , and approves the k candidates closest to it among $c_1, \dots, c_{\tilde{d}}$. The voter v' represents a job of length $k - 1$ and approves c^* and the $k - 1$ candidates closest to it among $c_1, \dots, c_{\tilde{d}}$.

To prove Theorem 5, we reduce $\text{PW}\langle d \rangle$ to PW_{po} , which is solvable in polynomial time (Betzler and Dorn 2010), similar to the proof of Theorem 2 for the special case of plurality and veto (discussed at the end of Section 4). In fact, we get again a running time that is polynomial in n, m and d . Hence, the possible winners can be found in polynomial time for plurality and veto even if d is part of the input.

For k -approval, the problem becomes intractable for $d \geq 2$, and every $k \geq 3$.

Theorem 6. *Let $d \geq 2$ and $k \geq 3$ be fixed. $\text{PW}\langle d \rangle$ is NP-complete for k -approval.*

Proof sketch. We focus on $d = 2$, since this is a special case of any $d > 2$. We show a reduction from non-preemptive scheduling with arrival times and deadlines (from Definition 1) where we have a single machine and every processing time satisfies $p_j \in \{k, k - 1\}$. Deciding whether a feasible schedule exists is strongly NP-complete for every $k > 2$ (Elffers and de Weerd 2014); then, we may assume that the maximal deadline $d_{\max}$ is polynomial in the number of jobs n . We also assume that the minimal arrival time is 1.

Let $\mathbf{J}_k$ and $\mathbf{J}_{k-1}$ be the sets of jobs with processing times k and $k - 1$, respectively. Each job J_j has an arrival time a_j and a deadline d_j . Let $\tilde{d}$ be the smallest multiple of k that is greater or equal to $d_{\max} - 1$. In the reduction, illustrated in Figure 3, the candidates are $C = \{c^*, c_1, \dots, c_{\tilde{d}}\}$. The positions are $c^* = \langle 0, 3\tilde{d} \rangle$ for c^* and $c_i = \langle i + 0.5, 0 \rangle$ for every c_i . For every job $J_j \in \mathbf{J}_k$ we introduce a voter who can approve any substring of length k of $(c_{a_j}, \dots, c_{d_j-1})$. For every job $J_j \in \mathbf{J}_{k-1}$ we introduce a voter who necessarily approves c^* , and can approve any substring of length $k - 1$ of $(c_{a_j}, \dots, c_{d_j-1})$. We also have $(|\mathbf{J}_{k-1}| - 1) \cdot \tilde{d}/k$ voters without uncertainty such that every candidate among $c_1, \dots, c_{\tilde{d}}$ is approved by exactly $|\mathbf{J}_{k-1}| - 1$ voters. It can be shown that c^* is a possible winner of $\mathbf{P}$ if and only if there exists a feasible schedule. $\square$

6 Partial Spatial Approval Voting

We now turn our attention to approval voting (AV), where each voter partitions the set of candidates into “approved” and “unapproved” candidates and the candidate with the




Figure 4: Example of an arrangement for $d = 2$, candidate set $\{c_1, c_2, c_3\}$, and a voter v . The possible approval sets of v are $\emptyset, \{c_1\}, \{c_2\}, \{c_3\}$, and $\{c_1, c_2\}$, depending on the actual position of v inside the rectangle. The red points are all event points from the sweep line algorithm in Theorem 7.

highest number of approvals is selected (Brams and Fishburn 1983). In contrast to the positional scoring rule k -approval, the number of approved candidates is not fixed under AV. In the spatial framework, it is often assumed that each voter v_j is associated with an *approval radius* $\rho_j \in \mathbb{R}$ and approves all candidates that are at a distance of at most ρ_j from voter v_j 's position (e.g., Godziszewski et al. 2021).

More formally, a *partial spatial approval profile* consists of a partial spatial profile $\mathbf{P}$ and a radius $\rho_j \in \mathbb{R}$ for every voter v_j . Each spatial completion $\mathbf{T} = (T_1, \dots, T_n)$ of $\mathbf{P}$ gives rise to an approval profile $\mathbf{A}_{\mathbf{T}} = (A_{T_1}, \dots, A_{T_n})$, where $A_{T_j} = \{c \in C : \|T_j - c\|_2 \leq \rho_j\}$ is the approval set of voter v_j . We call $\mathbf{A}_{\mathbf{T}}$ an *approval completion* of $\mathbf{T}$ (keeping in mind that, in contrast to the approval completions in Section 3, the approval sets can have different sizes). The definitions of possible and necessary winners and their respective decision problems remain unchanged.

First, for $d \leq 2$ dimensions, we show an upper bound on the number of different approval completions a single voter can have, and how to enumerate them in polynomial time. This is crucial for our results for necessary-winner problems. The regions in $\mathbb{R}^d$ where a voter has the same approval completion are bounded by (intersections of) d -dimensional *spheres* (rather than hyperplanes as in Section 4).

We consider the connection between a partial spatial approval profile, for a single voter, and an arrangement of d -spheres (with the d -dimensional rectangle of the voter). Instead of assuming the voter has a radius ρ around its position and approves all candidates inside this sphere, we can equivalently assume that we have a sphere with radius ρ around each *candidate's* position and the voter approves all candidates whose sphere contains the voter's position. See Figure 4 for a graphical depiction in two dimensions.

Now, recall that the voter's possible positions are described by a d -dimensional rectangle. Thus, if this rectangle intersects a sphere associated with a candidate c , then there is a completion where the voter approves c . Moreover, if there is a point inside the rectangle that is contained in the spheres of, say, the candidates c_1, c_2 , and c_3 , and in none of the other spheres, then there is a completion of the profile where the voter's approval set is exactly $\{c_1, c_2, c_3\}$: the voter's approval set is $\{c_1, c_2, c_3\}$ in exactly the completions where the position of the voter is in the intersection of the rectangle of the voter and the three spheres around c_1, c_2 and

c_3 . Thus, the possible completions of the voter's approval profile correspond to all (maximal) regions of the arrangement of m d -spheres with radius ρ (each around the position of one of the candidates), and the d -dimensional rectangle described by $\langle [\ell_{j,1}, u_{j,1}], \dots, [\ell_{j,d}, u_{j,d}] \rangle$ (where we only consider those regions that lie inside the rectangle).

Theorem 7. Fix $d \leq 2$ and let C be a set of d -dimensional candidates and P a single d -dimensional partial vote with radius $\rho \in \mathbb{R}_{\geq 0}$. Then, the set of approval completions of P can be enumerated in polynomial time.

Proof sketch. We focus on $d = 2$. To enumerate the regions of the arrangement as described above, we use a standard tool in computational geometry called a *sweep line* algorithm (e.g., see Halperin and Sharir 2017). The idea is to go over the arrangement from left to right, focusing on *event points* where something in the arrangement changes. The event points are all m leftmost points of the circles corresponding to the candidates and the bottom-left point of the rectangle corresponding to the voter. Additionally, every intersection point of the circles among themselves and with the rectangle constitutes an event point. See Figure 4 for an example. It is easy to verify that the total number of event points is at most $\mathcal{O}(m^2)$, and all these points can be computed in polynomial time. The sweep line algorithm considers the event points from left to right (from lowest to highest x -coordinate, breaking ties by their y -coordinates) and introduces the regions neighboring the current event point.

We store the regions as a list of $(m+1)$ -dimensional vectors. The i 'th entry of the vector is 1 if the region lies inside the sphere of c_i , and 0 otherwise. The $(m+1)$ st entry is 1 if the region is inside the rectangle. Given the list of all regions, we can iterate over those where the last entry is 1 and get the list of possible approval completions. In the full proof, we argue why the aforementioned event points exactly correspond to the moments in the sweep procedure where a new region might be introduced, and describe how we introduce such a region, i.e., update the data at an event point. $\square$

We can use the above result to find the necessary winners for AV if $d \leq 2$, similarly to the proof of Theorem 2. Note that we run the sweep line algorithm described above once for each voter v_j , using that voters radius ρ_j .⁵

Theorem 8. Let $d \leq 2$. $\text{NW}\langle d \rangle$ is solvable in polynomial time for AV.

On the other hand, we show that the possible-winner problem is hard, even for $d = 1$, by a reduction from non-preemptive scheduling with arrival times and deadlines (see Definition 1) where we have a single machine but arbitrary job lengths. Deciding whether a feasible schedule exists is strongly NP-complete (Garey and Johnson 1979). This is the first negative result we have for partial spatial voting (using approval or ranked preferences) in the one-dimensional case.

Theorem 9. $\text{PW}\langle d \rangle$ is NP-complete for any $d \geq 1$ for AV.

⁵Note that every two d -spheres intersect in a $(d-1)$ -sphere. For $d \geq 3$, it is thus not clear how to define the corresponding event points for a sweep plane algorithm in d dimensions.

We now briefly discuss the impact of these results on approval-based *multi-winner* elections.

Approval-based committee voting. In approval-based committee (ABC) voting (Lackner and Skowron 2022), a *committee* (i.e., subset of candidates) of a fixed size k needs to be selected based on approval ballots. Adapting the notions of Imber et al. (2022) to our setting, a set of candidates W is a *possible committee* if there is a completion of the partial spatial approval profile where W is selected, and a *necessary committee* if W wins in every completion. We parametrize both problems by the dimension d and the committee size k . The hardness of possible winners of the single-winner AV rule (Theorem 9) can be extended to ABC: we can embed any single-winner instance into an ABC instance by simply adding $k-1$ dummy candidates and appropriately many voters (without uncertainty) who ensure that the dummy candidates are in every winning committee. Then, solving the possible committee problem in this ABC instance would also solve the possible-winner problem for the single-winner instance. Note that this extends to any ABC voting rule that is equivalent to AV for $k = 1$. Conversely, using Theorem 7 and techniques used by Imber et al. (2022) to find a necessary committee under partial orders, one can show that the necessary committee problem in our setting is tractable for all *ABC scoring rules* (a large class of rules including all Thiele rules) for $d \leq 2$.

7 Conclusions

We introduced the framework of partial spatial voting, where candidates and voters are positioned in a geometrical space, but voters can have intervals of possible values in each dimension/issue. For positional scoring rules, we recovered the tractable cases of necessary and possible winners in the model of partial orders, for every fixed number of issues. In particular, we showed that the possible winners can be found in polynomial time for the plurality and veto rules, and that the necessary winners can be found in polynomial time for every positional scoring rule. We identified cases where the possible-winner problem is hard for partial orders but not for partial spatial voting. Specifically, this holds for two-valued rules other than plurality and veto, such as k -approval and k -veto for $k > 1$. We showed that the possible-winner problem may become intractable when the number of issues increases to a higher fixed number. For partial spatial approval ballots, we showed that the possible-winner problem under AV is intractable for all $d \geq 1$, and we gave an efficient algorithm for the necessary-winner problem under AV for $d \leq 2$.

It is left for future work to complete the picture of complexity for all positional scoring rules (see Table 1). Especially interesting is the case of Borda for $d = 1$. It is also left to study the complexity of non-positional voting rules, as done in the case of the necessary and possible winners for model of partial orders (Xia and Conitzer 2011). It would also be interesting to study the implications of voter uncertainty on the strategic considerations of candidates. For instance, in a model where probabilistic information about voters' locations is available, candidates may want to relocate in order to increase their winning probability.

Acknowledgments

The work of Markus Brill and Jonas Israel was supported by the Deutsche Forschungsgemeinschaft under grant BR 4744/2-1. The work of Aviram Imber and Benny Kimelfeld was supported by the Deutsche Forschungsgemeinschaft under project 412400621.

References

- Arrow, K. J. 1951. *Social Choice and Individual Values*. New Haven: Cowles Foundation, 1st edition. 2nd edition 1963.
- Barrot, N.; Gourvès, L.; Lang, J.; Monnot, J.; and Ries, B. 2013. Possible Winners in Approval Voting. 57–70. Springer-Verlag.
- Baumeister, D.; and Rothe, J. 2012. Taking the final step to a full dichotomy of the possible winner problem in pure scoring rules. *Information Processing Letters*, 112(5): 186–190.
- Betzler, N.; and Dorn, B. 2010. Towards a dichotomy for the Possible Winner problem in elections based on scoring rules. *Journal of Computer and System Sciences*, 76(8): 812–836.
- Black, D. 1948. On the Rationale of Group Decision-making. *Journal of Political Economy*, 56(1): 23–34.
- Bogomolnaia, A.; and Laslier, J.-F. 2007. Euclidean preferences. *Journal of Mathematical Economics*, 43(2): 87–98.
- Brams, S. J.; and Fishburn, P. C. 1983. *Approval Voting*. Boston: Birkhäuser.
- Doignon, J.-P.; and Falmagne, J.-C. 1994. A polynomial time algorithm for unidimensional unfolding representations. *Journal of Algorithms*, 16(2): 218–233.
- Edelsbrunner, H.; O’Rourke, J.; and Seidel, R. 1986. Constructing arrangements of lines and hyperplanes with applications. *SIAM Journal on Computing*, 15(2): 341–363.
- Elffers, J.; and de Weerd, M. 2014. Scheduling with two non-unit task lengths is NP-complete. *arXiv preprint arXiv:1412.3095*.
- Elkind, E.; and Faliszewski, P. 2014. Recognizing 1-Euclidean preferences: An alternative approach. In *Proceedings of the 7th International Symposium on Algorithmic Game Theory (SAGT)*, 146–157. Springer.
- Elkind, E.; Lackner, M.; and Peters, D. 2022. Preference restrictions in computational social choice: A survey. *arXiv preprint arXiv:2205.09092*.
- Enelow, J. M.; and Hinich, M. J. 1984. *The spatial theory of voting: An introduction*. Cambridge University Press.
- Estornell, A.; Das, S.; Elkind, E.; and Vorobeychik, Y. 2020. Election Control by Manipulating Issue Significance. In *Proceedings of Machine Learning Research*, 340–349. AUAI Press.
- Garey, M. R.; and Johnson, D. S. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman.
- Godziszewski, M.; Batko, P.; Skowron, P.; and Faliszewski, P. 2021. An Analysis of Approval-Based Committee Rules for 2D-Euclidean Elections. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, 5448–5455. AAAI Press.
- Halperin, D.; and Sharir, M. 2017. Arrangements. In *Handbook of discrete and computational geometry*, 723–762. Chapman and Hall/CRC.
- Imber, A.; Israel, J.; Brill, M.; and Kimelfeld, B. 2022. Approval-Based Committee Voting under Incomplete Information. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, 5076–5083. AAAI Press.
- Imber, A.; Israel, J.; Brill, M.; Shachnai, H.; and Kimelfeld, B. 2023. Finding Possible and Necessary Winners in Spatial Voting with Partial Information. *arXiv preprint arXiv:2302.08929*.
- Jamieson, K. G.; and Nowak, R. 2011. Active ranking using pairwise comparisons. *Proceedings of the 24th International Conference on Neural Information Processing Systems*, 2240–2248.
- Knoblauch, V. 2010. Recognizing one-dimensional Euclidean preference profiles. *Journal of Mathematical Economics*, 46(1): 1–5.
- Konczak, K.; and Lang, J. 2005. Voting procedures with incomplete preferences. In *Proceedings of the IJCAI-05 Multi-disciplinary Workshop on Advances in Preference Handling*.
- Lackner, M.; and Skowron, P. 2022. *Multi-Winner Voting with Approval Preferences*. Springer.
- Lang, J. 2020. Collective Decision Making under Incomplete Knowledge: Possible and Necessary Solutions. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI) Survey Track*, 4885–4891.
- Lu, J.; Zhang, D. K.; Rabinovich, Z.; Obratsova, S.; and Vorobeychik, Y. 2019. Manipulating Elections by Selecting Issues. In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 529–537. IFAAMAS.
- Miller, N. R. 1995. *Committees, Agendas and Voting*. Harwood Academic.
- Peters, D. 2017. Recognising multidimensional Euclidean preferences. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*. AAAI Press.
- Vakhania, N. 2012. Branch less, cut more and minimize the number of late equal-length jobs on identical machines. *Theoretical Computer Science*, 465: 49–60.
- Walsh, T. 2007. Uncertainty in Preference Elicitation and Aggregation. In *Proceedings of the 21st AAAI Conference on Artificial Intelligence (AAAI)*, 3–8. AAAI Press.
- Wu, J.; Estornell, A.; Kong, L.; and Vorobeychik, Y. 2022. Manipulating Elections by Changing Voter Perceptions. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, 557–563.
- Xia, L.; and Conitzer, V. 2011. Determining Possible and Necessary Winners Given Partial Orders. *Journal of Artificial Intelligence Research*, 41: 25–67.