Expected Frequency Matrices of Elections: Computation, Geometry, and Preference Learning

Niclas Boehmer
Algorithmics and Computational Complexity
Technische Universität Berlin
niclas.boehmer@tu-berlin.de

Robert Bredereck
TU Clausthal
robert.bredereck@tu-clausthal.de

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

Piotr Faliszewski
AGH University
faliszew@agh.edu.pl

Stanisław Szufa
AGH University
Jagiellonian University
szufa@agh.edu.pl

Abstract

We use the "map of elections" approach of Szufa et al. (AAMAS-2020) to analyze several well-known vote distributions. For each of them, we give an explicit formula or an efficient algorithm for computing its frequency matrix, which captures the probability that a given candidate appears in a given position in a sampled vote. We use these matrices to draw the "skeleton map" of distributions, evaluate its robustness, and analyze its properties. Finally, we develop a general and unified framework for learning the distribution of real-world preferences using the frequency matrices of established vote distributions.

1 Introduction

Computational social choice is a research area at the intersection of social choice (the science of collective decision-making) and computer science, which focuses on the algorithmic analysis of problems related to preference aggregation and elicitation (Brandt et al., 2013). Many of the early papers in this field were primarily theoretical, focusing on establishing the worst-case complexity of winner determination and strategic behavior under various voting rules—see, e.g., the papers of Hemaspaandra et al. (1997), Dwork et al. (2001), and Conitzer et al. (2007)—but more recent work often combines theoretical investigations with empirical analysis. For example, formal bounds on the running time and/or approximation ratio of a winner determination algorithm can be complemented by experiments that evaluate its performance on realistic instances; see, e.g., the works of Conitzer (2006), Betzler et al. (2014), Faliszewski et al. (2018) and Wang et al. (2019).

However, performing high-quality experiments requires the ability to organize and understand the available data. One way to achieve this is to form a so-called "map of elections," recently introduced by Szufa et al. (2020) and extended by Boehmer et al. (2021b). The idea is as follows. First, we fix a distance measure between elections. Second, we sample a number of elections from various distributions and real-life datasets—e.g., those collected in PrefLib (Mattei & Walsh, 2013)—and measure the pairwise distances between them. Third, we embed these elections into the 2D plane, mapping each election to a point so that the Euclidean distances between points are approximately equal to the distances between the respective elections. Finally, we plot these points, usually coloring them to indicate their origin (e.g., the distribution from which a given election was sampled); see Figure 2 later in the paper for an example of such a map. A location of an election on a map provides useful information about its properties. For example, Szufa et al. (2020) and Boehmer et al. (2021a,b) have shown that it can be used to predict (a) the Borda score of the winner of the election, (b) the

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

running time of ILP solvers computing the winners under the Harmonic-Borda multiwinner voting rule, or (c) the robustness of Plurality and Borda winners. Moreover, real-world elections of the same type (such as the ones from politics, sports, or surveys) tend to cluster in the same areas of the map; see also the positions on the map of the datasets collected by Boehmer & Schaar (2022). As such, the map has proven to be a useful framework to analyze the nature of elections and to visualize experimental results in a non-aggregate fashion.

Unfortunately, extending the map to incorporate additional examples and distributions is a challenging task, as the visual representation becomes cluttered and, more importantly, the embedding algorithms, which map elections to points in 2D, find it more difficult to preserve pairwise distances between points as the number of points increases. It is therefore desirable to reduce the number of points in a way that preserves the key features of the framework.

We address this challenge by drawing a map of *distributions* rather than individual elections, which we call the *skeleton map*. That is, instead of sampling 20–30 points from each distribution and placing them all on the map, as Szufa et al. (2020) and Boehmer et al. (2021b) do (obtaining around 800 points in total), we create a single point for each distribution. This approach is facilitated by the fact that prior work on the "map of elections" framework represented elections by their *frequency matrices*, which capture their essential features. The starting point of our work is the observation that this representation extends to distributions in a natural way. Thus, if we can compute the frequency matrix of some distribution $\mathcal{D}$, then, instead of sampling elections from $\mathcal{D}$ and creating a point on the map for each sample, we can create a single point for $\mathcal{D}$ itself.

Our Contribution. We provide three sets of results. First, for a number of prominent vote distributions, we show how to compute their frequency matrices, by providing an explicit formula or an efficient algorithm. Second, we draw the map of distributions (the *skeleton map*) and argue for its credibility and robustness. Finally, we use our results to estimate the parameters of the distributions that are closest to the real-world elections considered by Boehmer et al. (2021b). In more detail, we work in the setting of preference learning, where we are given an election and we want to learn the parameters of some distribution, so as to maximize the similarity of the votes sampled from this distribution and the input election. For example, we may be interested in fitting the classic model of Mallows (1957). This model is parameterized by a central vote v and a dispersion parameter ϕ, which specifies how likely it is to generate a vote at some distance from the central one (alternatively, one may use, e.g., the Plackett–Luce model). Previous works on preference learning typically proposed algorithms to learn the parameters of one specific (parameterized) vote distribution (see, e.g., the works of Lu & Boutilier (2014); Mandhani & Meilă (2009); Meila & Chen (2010); Vitelli et al. (2017); Murphy & Martin (2003); Awasthi et al. (2014) for (mixtures of) the Mallows model and the works of Guiver & Snelson (2009); Hunter (2004); Minka (2004); Gormley & Murphy (2008) for (mixtures of) the Plackett–Luce model). Using frequency matrices, we offer a more general approach. Indeed, given an election and a parameterized vote distribution whose frequency matrix we can compute, the task of learning the distribution's parameters boils down to finding parameters that minimize the distance between the election and the matrices of the distribution. While this minimization problem may be quite challenging, our approach offers a uniform framework for dealing with multiple kinds of distributions at the same time. We find that for the case of the Mallows distribution, our approach learns parameters very similar to those established using maximum likelihood-based approaches. Omitted proofs and discussions are in the appendix. The source code used for the experiments is available in a GitHub repository[1].

2 Preliminaries

Given an integer t, we write $[t]$ to denote the set $\{1, \dots, t\}$. We interpret a vector $x \in \mathbb{R}^m$ as an $m \times 1$ matrix (i.e., we use column vectors as the default).

Preference Orders and Elections. Let C be a finite, nonempty set of candidates. We refer to total orders over C as *preference orders* (or, equivalently, *votes*), and denote the set of all preference orders over C by $\mathcal{L}(C)$. Given a vote v and a candidate c, by $\mathrm{pos}_v(c)$ we mean the position of c in v (the top-ranked candidate has position 1, the next one has position 2, and so on). If a candidate a is ranked above another candidate b in vote v, we write $v\colon a \succ b$. Let $\mathrm{rev}(v)$ denote the *reverse* of

[1] github.com/Project-PRAGMA/Expected-Frequency-Matrices-NeurIPS-2022

vote v. An *election* $E = (C, V)$ consists of a set $C = \{c_1, \dots, c_m\}$ of candidates and a collection $V = (v_1, \dots, v_n)$ of votes. Occasionally we refer to the elements of V as voters rather than votes.

Frequency Matrices. Consider an election $E = (C, V)$ with $C = \{c_1, \dots, c_m\}$ and $V = (v_1, \dots, v_n)$. For each candidate c_j and position $i \in [m]$, we define $\#\mathrm{freq}_E(c_j, i)$ to be the fraction of the votes from V that rank c_j in position i. We define the column vector $\#\mathrm{freq}_E(c_j)$ to be $(\#\mathrm{freq}_E(c_j, 1), \dots, \#\mathrm{freq}_E(c_j, m))$ and matrix $\#\mathrm{freq}(E)$ to consist of vectors $\#\mathrm{freq}_E(c_1), \dots, \#\mathrm{freq}_E(c_m)$. We refer to $\#\mathrm{freq}(E)$ as the *frequency matrix of election* E. Frequency matrices are bistochastic, i.e., their entries are nonnegative and each of their rows and columns sums up to one.

Example 2.1. Let $E = (C, V)$ be an election with candidate set $C = \{a, b, c, d, e\}$ and four voters, v_1, v_2, v_3, and v_4. Below, we show the voters' preference orders (on the left) and the election's frequency matrix (on the right).

$$
\begin{aligned}
v_1 &\colon a \succ b \succ c \succ d \succ e,\\
v_2 &\colon c \succ b \succ d \succ a \succ e,\\
v_3 &\colon d \succ e \succ c \succ b \succ a,\\
v_4 &\colon b \succ c \succ a \succ d \succ e.
\end{aligned}
\qquad
\begin{array}{c}
\\ 1\\ 2\\ 3\\ 4\\ 5
\end{array}
\begin{array}{c}
\begin{array}{ccccc} a & b & c & d & e \end{array}\\
\begin{bmatrix}
1/4 & 1/4 & 1/4 & 1/4 & 0\\
0 & 1/2 & 1/4 & 0 & 1/4\\
1/4 & 0 & 1/2 & 1/4 & 0\\
1/4 & 1/4 & 0 & 1/2 & 0\\
1/4 & 0 & 0 & 0 & 3/4
\end{bmatrix}
\end{array}
$$

Given a vote v, we write $\#\mathrm{freq}(v)$ to denote the frequency matrix of the election containing this vote only; $\#\mathrm{freq}(v)$ is a permutation matrix, with a single 1 in each row and in each column. Thus, for an election $E = (C, V)$ with $V = (v_1, \dots, v_n)$ we have $\#\mathrm{freq}(E) = \frac{1}{n} \cdot \sum_{i=1}^{n} \#\mathrm{freq}(v_i)$.

Compass Matrices. For even m, Boehmer et al. (2021b) defined the following four $m \times m$ "compass" matrices, which appear to be extreme on the "map of elections":

1. The *identity matrix*, ID_m, has ones on the diagonal and zeroes everywhere else (it corresponds to an election where all voters agree on a single preference order).
2. The *uniformity matrix*, UN_m, has all entries equal to $1/m$ (it corresponds to lack of agreement; each candidate is ranked at each position equally often).
3. The *stratification matrix*, ST_m, is partitioned into four quadrangles, where all entries in the top-left and bottom-right quadrangles are equal to $2/m$, and all other entries are equal to zero (it corresponds to partial agreement; the voters agree which half of the candidates is superior, but disagree on everything else).
4. The *antagonism matrix*, AN_m, has values $1/2$ on both diagonals and zeroes elsewhere (it captures a conflict: it is a matrix of an election where half of the voters rank the candidates in one way and half of the voters rank them in the opposite way).

Below, we show examples of these matrices for $m = 4$:

$$
\mathrm{UN}_4 = \begin{bmatrix} 1/4 & 1/4 & 1/4 & 1/4\\ 1/4 & 1/4 & 1/4 & 1/4\\ 1/4 & 1/4 & 1/4 & 1/4\\ 1/4 & 1/4 & 1/4 & 1/4 \end{bmatrix}, \mathrm{ID}_4 = \begin{bmatrix} 1 & 0 & 0 & 0\\ 0 & 1 & 0 & 0\\ 0 & 0 & 1 & 0\\ 0 & 0 & 0 & 1 \end{bmatrix}, \mathrm{ST}_4 = \begin{bmatrix} 1/2 & 1/2 & 0 & 0\\ 1/2 & 1/2 & 0 & 0\\ 0 & 0 & 1/2 & 1/2\\ 0 & 0 & 1/2 & 1/2 \end{bmatrix}, \mathrm{AN}_4 = \begin{bmatrix} 1/2 & 0 & 0 & 1/2\\ 0 & 1/2 & 1/2 & 0\\ 0 & 1/2 & 1/2 & 0\\ 1/2 & 0 & 0 & 1/2 \end{bmatrix}.
$$

We omit the subscript in the names of these matrices if its value is clear from the context or irrelevant.

EMD. Let $x = (x_1, \dots, x_n)$ and $y = (y_1, \dots, y_n)$ be two vectors with nonnegative real entries that sum up to 1. Their *Earth mover's distance*, denoted $\mathrm{EMD}(x, y)$, is the cost of transforming x into y using operations of the form: Given indices $i, j \in [n]$ and a positive value δ such that $x_i \geq \delta$, at the cost of $\delta \cdot |i - j|$, replace x_i with $x_i - \delta$ and x_j with $x_j + \delta$ (this corresponds to moving δ amount of "earth" from position i to position j). $\mathrm{EMD}(x, y)$ can be computed in polynomial time by a standard greedy algorithm.

Positionwise Distance (Szufa et al., 2020). Let $A = (a_1, \dots, a_m)$ and $B = (b_1, \dots, b_m)$ be two $m \times m$ frequency matrices. Their *raw positionwise distance* is $\mathrm{rawPOS}(A, B) = \min_{\sigma \in S_m} \sum_{i=1}^{m} \mathrm{EMD}(a_i, b_{\sigma(i)})$, where S_m denotes the set of all permutations over $[m]$. We will normalize these distances by $\frac{1}{3}(m^2 - 1)$, which Boehmer et al. (2021b, 2022) proved to be the maximum distance between two $m \times m$ frequency matrices and the distance between ID_m and

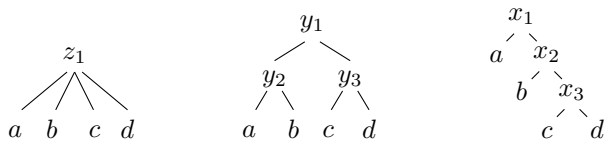

(a) $\mathcal{T}_1$ (b) $\mathcal{T}_2$ (c) $\mathcal{T}_3$

Figure 1: Three examples of clone decomposition trees.

UN_m: $\mathrm{nPOS}(A, B) = \frac{\mathrm{rawPOS}(A,B)}{\frac{1}{3}(m^2-1)}$. For two elections E and F with equal-sized candidate sets, their positionwise distance, raw or normalized, is defined as the positionwise distance between their frequency matrices.

Paths Between the Compass Matrices. Let X and Y be two compass matrices. Boehmer et al. (2021b) showed that if we take their affine combination $Z = \alpha X + (1-\alpha)Y$ ($0 \leq \alpha \leq 1$) then $\mathrm{nPOS}(X, Z) = (1-\alpha)\mathrm{nPOS}(X, Y)$ and $\mathrm{nPOS}(Z, Y) = \alpha\mathrm{nPOS}(X, Y)$. Such affine combinations form direct paths between the compass matrices; they are also possible between any two frequency matrices of a given size, not just the compass ones, but may require shuffling the matrices' columns (Boehmer et al., 2021b).

Structured Domains. We consider two classes of structured elections, single-peaked elections (Black, 1958), and group-separable elections (Inada, 1964). For a discussion of these domains and the motivation behind them, see the original papers and the overviews by Elkind et al. (2017, 2022).

Intuitively, an election is single-peaked if we can order the candidates so that, as each voter considers the candidates in this order (referred to as the *societal axis*), his or her appreciation first increases and then decreases. The axis may, e.g., correspond to the left-right political spectrum.

Definition 2.2. Let v be a vote over C and let $\lhd$ be the societal axis over C. We say that v is *single-peaked with respect to* $\lhd$ if for every $t \in [|C|]$ its t top-ranked candidates form an interval within $\lhd$. An election is *single-peaked with respect to* $\lhd$ if all its votes are. An election is *single-peaked (SP)* if it is single-peaked with respect to some axis.

Note that the election from Example 2.1 is single-peaked with respect to the axis $a \lhd b \lhd c \lhd d \lhd e$.

We also consider *group-separable elections*, introduced by Inada (1964). For our purposes, it will be convenient to use the tree-based definition of Karpov (2019). Let $C = \{c_1, \ldots, c_m\}$ be a set of candidates, and consider a rooted, ordered tree $\mathcal{T}$ whose leaves are elements of C. The *frontier* of such a tree is the preference order that ranks the candidates in the order in which they appear in the tree from left to right. A preference order is *consistent* with a given tree if it can be obtained as its frontier by reversing the order in which the children of some nodes appear.

Definition 2.3. An election $E = (C, V)$ is *group-separable* if there is a rooted, ordered tree $\mathcal{T}$ whose leaves are members of C, such that each vote in V is consistent with $\mathcal{T}$.

The trees from Definition 2.3 form a subclass of *clone decomposition trees*, which are examples of PQ-trees (Elkind et al., 2012; Booth & Lueker, 1976).

Example 2.4. Consider candidate set $C = \{a, b, c, d\}$, trees $\mathcal{T}_1$, $\mathcal{T}_2$, and $\mathcal{T}_3$ from Figure 1, and votes $v_1 \colon a \succ b \succ c \succ d$, $v_2 \colon c \succ d \succ b \succ a$, and $v_3 \colon b \succ d \succ c \succ a$. Vote v_1 is consistent with each of the trees, v_2 is consistent with $\mathcal{T}_2$ (reverse the children of y_1 and y_2), and v_3 is consistent with $\mathcal{T}_3$ (reverse the children of x_1 and x_3).

3 Frequency Matrices for Vote Distributions

We show how to compute frequency matrices for several well-known distributions over votes.

3.1 Setup and Interpretation

A *vote distribution* for a candidate set C is a function $\mathcal{D}$ that assigns a probability to each preference order over C. Formally, we require that for each $v \in \mathcal{L}(C)$ it holds that $\mathcal{D}(v) \geq 0$ and

$\sum_{v \in \mathcal{L}(C)} \mathcal{D}(v) = 1$. We say that a vote v is in the *support* of $\mathcal{D}$ if $\mathcal{D}(v) > 0$. Given such a distribution, we can form an election by repeatedly drawing votes according to the specified probabilities. For example, we can sample each element of $\mathcal{L}(C)$ with equal probability; this distribution, which is known as *impartial culture (IC)*, is denoted by $\mathcal{D}_{\mathrm{IC}}$ (we omit the candidate set from our notation as it will always be clear from the context). The *frequency matrix of a vote distribution* $\mathcal{D}$ over a candidate set C is $\#\mathrm{freq}(\mathcal{D}) = \sum_{v \in \mathcal{L}(C)} \mathcal{D}(v) \cdot \#\mathrm{freq}(v)$. For example, we have $\#\mathrm{freq}(\mathcal{D}_{\mathrm{IC}}) = \mathrm{UN}$. One interpretation of $\#\mathrm{freq}(\mathcal{D})$ is that the entry for a candidate c_j and a position i is the probability that a vote v sampled from $\mathcal{D}$ has c_j in position i (which we denote as $\mathbb{P}[\mathrm{pos}_v(c_j) = i]$). Another interpretation is that if we sample a large number of votes then the resulting election's frequency matrix would be close to $\#\mathrm{freq}(\mathcal{D})$ with high probability. More formally, if we let $\mathcal{M}_n$ be a random variable equal to the frequency matrix of an n-voter election generated according to $\mathcal{D}$, then it holds that $\lim_{n \to \infty} \mathbb{E}(\mathcal{M}_n) = \#\mathrm{freq}(\mathcal{D})$.

3.2 Group-Separable Elections

We first consider sampling group-separable votes. Given a rooted tree $\mathcal{T}$ whose leaves are labeled by elements of $C = \{c_1, \dots, c_m\}$, let $\mathcal{D}^{\mathcal{T}}_{\mathrm{GS}}$ be the distribution assigning equal probability to all votes consistent with $\mathcal{T}$, and zero probability to all other votes; one can think of $\mathcal{D}^{\mathcal{T}}_{\mathrm{GS}}$ as impartial culture restricted to the group-separable subdomain defined by $\mathcal{T}$. To sample from $\mathcal{D}^{\mathcal{T}}_{\mathrm{GS}}$, we can toss a fair coin for each internal node of $\mathcal{T}$, reversing the order of its children if this coin comes up heads, and output the frontier of the resulting tree. We focus on the following types of trees:

1. $\mathrm{Flat}(c_1, \dots, c_m)$ is a tree with a single internal node, whose children, from left to right, are $c_1, c_2, \dots, c_m$. There are only two preference orders consistent with this tree, $c_1 \succ \cdots \succ c_m$ and its reverse.
2. $\mathrm{Bal}(c_1, \dots, c_m)$ is a perfectly balanced binary tree with frontier $c_1, \dots, c_m$ (hence we assume the number m of candidates to be a power of two).
3. $\mathrm{CP}(c_1, \dots, c_m)$ is a binary caterpillar tree: it has internal nodes $x_1, \dots, x_{m-1}$; for each $j \in [m-2]$, x_j has c_j as the left child and x_{j+1} as the right one, whereas x_{m-1} has both c_{m-1} and c_m as children.

The first tree in Figure 1 is flat, the second one is balanced, and the third one is a caterpillar tree. If $\mathcal{T}$ is a caterpillar tree, then we refer to $\mathcal{D}^{\mathcal{T}}_{\mathrm{GS}}$ as the *GS/caterpillar distribution*. We use a similar terminology for the other trees.

Theorem 3.1. *Let F be the frequency matrix of distribution $\mathcal{D}^{\mathcal{T}}_{\mathrm{GS}}$. If $\mathcal{T}$ is flat then $F = \mathrm{AN}$, and if it is balanced then $F = \mathrm{UN}$. If $\mathcal{T}$ is a caterpillar tree $\mathrm{CP}(c_1, \dots, c_m)$, then for each candidate c_j the probability that c_j appears in a position $i \in [m]$ in a random vote v sampled from $\mathcal{D}^{\mathcal{T}}_{\mathrm{GS}}$ is:*

$$\tfrac{1}{2^j}\tbinom{j-1}{i-1} \cdot \mathbb{1}_{i \le j} + \tfrac{1}{2^j}\tbinom{j-1}{(i-1)-(m-j)} \cdot \mathbb{1}_{i > m-j}.$$

Proof. The cases of flat and balanced trees are immediate, so we focus on caterpillar trees. Let $\mathcal{T} = \mathrm{CP}(c_1, \dots, c_m)$ with internal nodes $x_1, \dots, x_{m-1}$, and consider a candidate c_j and a position $i \in [m]$. Let v be a random variable equal to a vote sampled from $\mathcal{D}^{\mathcal{T}}_{\mathrm{GS}}$. We say that a node x_ℓ, $\ell \in [m-1]$, is *reversed* if the order of its children is reversed. Note that for $\ell < r$ it holds that c_r precedes c_ℓ in the frontier if and only if x_ℓ is reversed. Suppose that x_j is not reversed. Then v ranks c_j above each of $c_{j+1}, \dots, c_m$. This means that for c_j to be ranked exactly in position i, it must be that $j \ge i$ and exactly $i-1$ nodes among $x_1, \dots, x_{j-1}$ are not reversed. If $j \ge i$, the probability that x_j and $i-1$ nodes among $x_1, \dots, x_{j-1}$ are not reversed is $\frac{1}{2^j} \cdot \binom{j-1}{i-1}$. On the other hand, if x_j is reversed, then v ranks candidates $c_{j+1}, \dots, c_m$ above c_j. As there are $m-j$ of them, for c_j to be ranked exactly in position i it must hold that $i > m-j$ and exactly $(i-1)-(m-j)$ nodes among $x_1, \dots, x_{j-1}$ are not reversed. This happens with probability $\frac{1}{2^j} \cdot \binom{j-1}{(i-1)-(m-j)}$. $\square$

Regarding distributions $\mathcal{D}^{\mathcal{T}}_{\mathrm{GS}}$ not handled in Theorem 3.1, we still can compute their frequence matrices efficiently.

Theorem 3.2. *There is an algorithm that given a tree $\mathcal{T}$ computes $\#\mathrm{freq}(D^{\mathcal{T}}_{\mathrm{GS}})$ using polynomially many arithmetic operations with respect to the number of nodes in $\mathcal{T}$.*

3.3 From Caterpillars to Single-Peaked Preferences.

There is a relationship between GS/caterpillar votes and single-peaked ones, which will be very useful when computing one of the frequency matrices in the next section.

Theorem 3.3. *Given a ranking v over $C = \{c_1, \dots, c_m\}$, let $\widehat{v}$ be another ranking over C such that, for each $j \in [m]$, if c_j is ranked in position i in v then c_i is ranked in position $m - j + 1$ in $\widehat{v}$. Suppose that v is in the support of $\mathcal{D}_{\mathrm{GS}}^{\mathcal{T}}$, where $\mathcal{T} = \mathrm{CP}(c_1, \dots, c_m)$. Then $\widehat{v}$ is single-peaked with respect to $c_1 \lhd \cdots \lhd c_m$.*

There are exactly 2^{m-1} votes in the support of $\mathcal{D}_{\mathrm{GS}}^{\mathcal{T}}$ (this follows by simple counting) and there are 2^{m-1} votes that are single-peaked with respect to $c_1 \lhd \cdots \lhd c_m$. As $u \neq v$ implies $\widehat{u} \neq \widehat{v}$, it follows that the mapping $v \mapsto \widehat{v}$ is a bijection between all votes in the support of $\mathcal{D}_{\mathrm{GS}}^{\mathcal{T}}$ and all votes that are single-peaked with respect to $c_1 \lhd \cdots \lhd c_m$.

3.4 Single-Peaked Elections

We consider two models of generating single-peaked elections, one due to Walsh (2015) and one due to Conitzer (2009). Let us fix a candidate set $C = \{c_1, \dots, c_m\}$ and a societal axis $c_1 \lhd \cdots \lhd c_m$. Under the Walsh distribution, denoted $\mathcal{D}_{\mathrm{SP}}^{\mathrm{Wal}}$, each vote that is single-peaked with respect to $\lhd$ has equal probability (namely, $\frac{1}{2^{m-1}}$), and all other votes have probability zero. By Theorems 3.1 and 3.3, we immediately obtain the frequency matrix for the Walsh distribution (in short, it is the transposed matrix of the GS/caterpillar distribution).

Corollary 3.4. *Consider a candidate set $C = \{c_1, \dots, c_m\}$ and an axis $c_1 \lhd \cdots \lhd c_m$. The probability that candidate c_j appears in position i in a vote sampled from $\mathcal{D}_{\mathrm{SP}}^{\mathrm{Wal}}$ is: $\frac{1}{2^{m-i+1}}\binom{m-i}{j-1} \cdot \mathbb{1}_{j \leq m-i+1} + \frac{1}{2^{m-i+1}}\binom{m-i}{j-i} \cdot \mathbb{1}_{j > i-1}$.*

To sample a vote from the Conitzer distribution, $\mathcal{D}_{\mathrm{SP}}^{\mathrm{Con}}$ (also known as the *random peak distribution*), we pick some candidate c_j uniformly at random and rank him or her on top. Then we perform $m - 1$ iterations, where in each we choose (uniformly at random) a candidate directly to the right or the left of the already selected ones, and place him or her in the highest available position in the vote.

Theorem 3.5. *Let $c_1 \lhd \cdots \lhd c_m$ be the societal axis, where m is an even number, and let v be a random vote sampled from $\mathcal{D}_{\mathrm{SP}}^{\mathrm{Con}}$ for this axis. For $j \in [\frac{m}{2}]$ and $i \in [m]$ we have:*

$$\mathbb{P}[\mathrm{pos}_v(c_j) = i] = \begin{cases} 2/2m & \text{if } i < j, \\ (j+1)/2m & \text{if } i = j, \\ 1/2m & \text{if } j < i < m - j + 1, \\ (m-j+1)/2m & \text{if } i = m - j + 1, \\ 0 & \text{if } i + j > m. \end{cases}$$

Further, for each candidate $c_j \in C$ and each position $i \in [m]$ we have $\mathbb{P}[\mathrm{pos}_v(c_j) = i] = \mathbb{P}[\mathrm{pos}_v(c_{m-j+1}) = i]$.

3.5 Mallows Model

Finally, we consider the classic Mallows distribution. It has two parameters, a central vote v^* over m candidates and a dispersion parameter $\phi \in [0, 1]$. The probability of sampling a vote v from this distribution (denoted $\mathcal{D}_{\mathrm{Mal}}^{v^*,\phi}$) is: $\mathcal{D}_{\mathrm{Mal}}^{v^*,\phi}(v) = \frac{1}{Z}\phi^{\kappa(v,v^*)}$, where $Z = 1 \cdot (1 + \phi) \cdot (1 + \phi + \phi^2) \cdot \cdots \cdot (1 + \cdots + \phi^{m-1})$ is a normalizing constant and $\kappa(v, v^*)$ is the swap distance between v and v^* (i.e., the number of swaps of adjacent candidates needed to transform v into v^*). In our experiments, we consider a new parameterization, introduced by Boehmer et al. (2021b). It uses a *normalized dispersion parameter* $\mathrm{norm}\text{-}\phi$, which is converted to a value of ϕ so that the expected swap distance between the central vote v^* and a sampled vote v is $\frac{\mathrm{norm}\text{-}\phi}{2}$ times the maximum swap distance between two votes (so, $\mathrm{norm}\text{-}\phi = 1$ is equivalent to IC and for $\mathrm{norm}\text{-}\phi = 0.5$ we get elections that lie close to the middle of the UN–ID path).

Our goal is now to compute the frequency matrix of $\mathcal{D}_{\mathrm{Mal}}^{v^*,\phi}$. That is, given the candidate ranked in position j in the central vote, we want to compute the probability that he or she appears in a

given position $i \in [m]$ in the sampled vote. Given a positive integer m, consider the candidate set $C(m) = \{c_1, \dots, c_m\}$ and the central vote $v_m^*\colon c_1 \succ \cdots \succ c_m$. Fix a candidate $c_j \in C(m)$, and a position $i \in [m]$. For every integer k between 0 and $m(m-1)/2$, let $S(m,k)$ be the number of votes in $\mathcal{L}(C(m))$ that are at swap distance k from v_m^*, and define $T(m,k,j,i)$ to be the number of such votes that have c_j in position i. One can compute $S(m,k)$ in time polynomial in m (OEIS Foundation Inc., 2020); using $S(m,k)$, we show that the same holds for $T(m,k,j,i)$.

Lemma 3.6. *There is an algorithm that computes $T(m,k,j,i)$ in polynomial time with respect to m.*

We can now express the probability of sampling a vote v, where the candidate ranked in position j in the central vote v^* ends up in position i under $\mathcal{D}_{\mathrm{Mal}}^{v^*,\phi}$, as:

$$f_m(\phi, j, i) = \frac{1}{Z} \sum_{k=0}^{m(m-1)/2} T(m,k,j,i)\phi^k. \tag{1}$$

The correctness follows from the definitions of T and $\mathcal{D}_{\mathrm{Mal}}^{v^*,\phi}$. By Lemma 3.6, we have the following.

Theorem 3.7. *There exists an algorithm that, given a number m of candidates, a vote v^*, and a parameter ϕ, computes the frequency matrix of $\mathcal{D}_{\mathrm{Mal}}^{v^*,\phi}$ using polynomially many operations in m.*

Note that Equation (1) only depends on ϕ, j and i (and, naturally, on m). Using this fact, we can also compute frequency matrices for several variants of the Mallows distribution.

Remark 3.8. Given a vote v, two dispersion parameters ϕ and ψ, and a probability $p \in [0,1]$, we define the distribution p-$\mathcal{D}_{\mathrm{Mal}}^{v,\phi,\psi}$ as $p \cdot \mathcal{D}_{\mathrm{Mal}}^{v,\phi} + (1-p) \cdot \mathcal{D}_{\mathrm{Mal}}^{\mathrm{rev}(v),\psi}$, i.e., with probability p we sample a vote from $\mathcal{D}_{\mathrm{Mal}}^{v,\phi}$ and with probability $1-p$ we sample a vote from $\mathcal{D}_{\mathrm{Mal}}^{\mathrm{rev}(v),\psi}$. The probability that candidate c_j appears in position i in the resulting vote is $p \cdot f_m(\phi, j, i) + (1-p) \cdot f_m(\psi, m-j+1, i)$.

Remark 3.9. Consider a candidate set $C = \{c_1, \dots, c_m\}$. Given a vote distribution $\mathcal{D}$ over $\mathcal{L}(C)$ and a parameter ϕ, define a new distribution $\mathcal{D}'$ as follows: Draw a vote $\hat{v}$ according to $\mathcal{D}$ and then output a vote v sampled from $\mathcal{D}_{\mathrm{Mal}}^{\hat{v},\phi}$; indeed, such models are quite natural, see, e.g., the work of Kenig & Kimelfeld (2019). For each $t \in [m]$, let $g(j,t)$ be the probability that c_j appears in position t in a vote sampled from $\mathcal{D}$. The probability that c_j appears in position $i \in [m]$ in a vote sampled from $\mathcal{D}'$ is $\sum_{t=1}^{m} g(j,t) \cdot f(\phi, t, i)$. In terms of matrix multiplication, this means that $\#\mathrm{freq}(\mathcal{D}') = \#\mathrm{freq}(\mathcal{D}_{\mathrm{Mal}}^{v^*,\phi}) \cdot \#\mathrm{freq}(\mathcal{D})$, where v^* is $c_1 \succ \cdots \succ c_m$. We write ϕ-Conitzer (ϕ-Walsh) to refer to this model where we use the Conitzer (Walsh) distribution as the underlying one and normalized dispersion parameter ϕ.

4 Skeleton Map

Our goal in this section is to form what we call a *skeleton map of vote distributions* (skeleton map, for short), evaluate its quality and robustness, and compare it to the map of Boehmer et al. (2021b). Throughout this section, whenever we speak of a distance between elections or matrices, we mean the positionwise distance (occasionally we will also refer to the Euclidean distances on our maps, but we will always make this explicit). Let $\Phi = \{0, 0.05, 0.1, \dots, 1\}$ be a set of normalized dispersion parameters that we will be using for Mallows-based distributions in this section.

We form the skeleton map following the general approach of Szufa et al. (2020) and Boehmer et al. (2021b). For a given number of candidates, we consider the four compass matrices (UN, ID, AN, ST) and paths between each matrix pair consisting of their convex combinations (gray dots), the frequency matrices of the Mallows distribution with normalized dispersion parameters from Φ (blue triangles), and the frequency matrices of the Conitzer (CON), Walsh (WAL), and GS/caterpillar distribution (CAT). Moreover, we add the frequency matrices of the following vote distributions (we again use the dispersion parameters from Φ): (i) The distribution $1/2$-$\mathcal{D}_{\mathrm{Mal}}^{v,\phi,\phi}$ as defined in Remark 3.8 (red triangles), (ii) the distribution where with equal probability we mix the standard Mallows distribution and $1/2$-$\mathcal{D}_{\mathrm{Mal}}^{v,\phi,\phi}$ (green triangles), and (iii) the ϕ-Conitzer and ϕ-Walsh distributions as defined in Remark 3.9 (magenta and orange crosses). For each pair of these matrices we compute their positionwise distance. Then we find an embedding of the matrices into a 2D plane, so that each matrix is a point and the Euclidean distances between these points are as similar to the positionwise distances as possible (we use the MDS algorithm, as implemented in the Python sklearn.manifold.MDS package). In Figure 3 we show our map for the case of 10 candidates (the

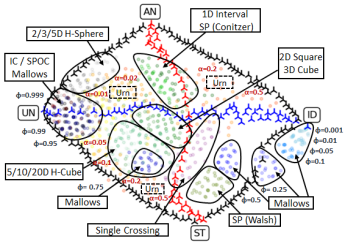

Figure 2: Map of elections, together with the compass matrices, as presented by Boehmer et al. (2021b).

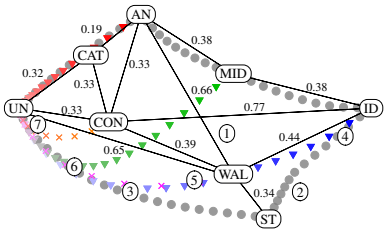

Figure 3: The skeleton map with 10 candidates. We have MID $= {}^1\!/_2$AN $+ {}^1\!/_2$ID. Each point labeled with a number is a real-world election as described in Section 5.

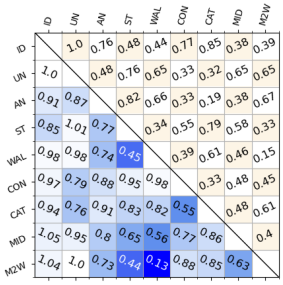

	ID	UN	AN	ST	WAL	CON	CAT	MID	M2W
ID		1.0	0.76	0.48	0.44	0.71	0.85	0.38	0.39
UN	1.0		0.48	0.76	0.65	0.33	0.32	0.65	0.65
AN	0.91	0.87		0.82	0.66	0.33	0.19	0.38	0.67
ST	0.85	1.01	0.77		0.34	0.55	0.79	0.58	0.33
WAL	0.98	0.98	0.74	0.45		0.39	0.61	0.46	0.15
CON	0.97	0.79	0.88	0.95	0.98		0.33	0.48	0.45
CAT	0.94	0.76	0.91	0.83	0.82	0.55		0.48	0.61
MID	1.05	0.95	0.8	0.65	0.56	0.77	0.86		0.4
M2W	1.04	1.0	0.73	0.44	0.13	0.88	0.85	0.63	

Figure 4: In the top-right part, we show the normalized positionwise distance. In the bottom-left one, we show the embedding misrepresentation ratios.

lines between some points/matrices show their positionwise distances; to maintain clarity, we only provide some of them).

We now verify the credibility of the skeleton map. As the map does not have many points, we expect its embedding to truly reflect the positionwise distances between the matrices. This, indeed, seems to be the case, although some distances are represented (much) more accurately than the others. In Figure 4 we provide the following data for a number of matrices (for $m = 10$; matrix M2W is the Mallows matrix in our data set that is closest to the Walsh matrix). In the top-right part (the white-orange area), we give positionwise distances between the matrices, and in the bottom-left part (the blue area), for each pair of matrices X and Y we report the *misrepresentation ratio* $\frac{\mathrm{Euc}(X,Y)}{\mathrm{nPOS}(X,Y)}$, where $\mathrm{Euc}(X, Y)$ is the Euclidean distance between X and Y in the embedding, normalized by the Euclidean distance between ID and UN. The closer they are to 1, the more accurate is the embedding. The misrepresentation ratios are typically between 0.8 and 1.15, with many of them between 0.9 and 1.05. Thus, in most cases, the map is quite accurate and offers good intuition about the relations between the matrices. Yet, some distances are represented particularly badly. As an extreme example, the Euclidean distance between the Walsh matrix and the closest Mallows matrix, M2W, is off by almost a factor of 8 (these matrices are close, but not as close as the map suggests). Thus, while one always has to verify claims suggested by the skeleton map, we view it as quite credible. This conclusion is particularly valuable when we compare the skeleton map and the map of Boehmer et al. (2021b), shown in Figure 2. The two maps are similar, and analogous points (mostly) appear in analogous positions. Perhaps the biggest difference is the location of the Conitzer matrix on the skeleton map and Conitzer elections in the map of Boehmer et al., but even this difference is not huge. We remark that the Conitzer matrix is closer to UN and AN than to ID and ST, whereas for the Walsh matrix the opposite is true. Boehmer et al. (2021b) make a similar observation; our results allow us to make this claim formal. In Appendix E, we analyze the robustness of the skeleton map with respect to varying the number of candidates. We find that except for pairs including the Walsh or GS/caterpillar matrices, which "travel" on the map as the number of candidates increases, the distance between each pair of matrices in the skeleton map stays nearly constant.

5 Learning Vote Distributions

We demonstrate how the positionwise distance and frequency matrices can be used to fit vote distributions to given real-world elections. Specifically, we consider the Mallows model ($\mathcal{D}^{v,\phi}_{\mathrm{Mal}}$) and the ϕ-Conitzer and ϕ-Walsh models. Naturally, we could use more distributions, but we focus on showcasing the technique and the general unified approach. Concerning our results, among others, we verify that for Mallows model our approach is strongly correlated with existing maximum-likelihood approaches. Moreover, unlike in previous works, we compare the capabilities of different distributions to fit the given elections. We remark that if we do not have an algorithm for computing a frequency matrix of a given vote distribution, we can obtain an approximate matrix by sampling sufficiently many votes from this distribution. In principle, it is also possible to deal with distributions

Table 1: Closest distributions for seven illustrative real-world elections. For each, we include the parameterization that produces the closest frequency matrix, and give the normalized positionwise distance of the elections from this matrix.

id	source	$\mathcal{D}_{\mathrm{Mal}}^{v,\phi}$		ϕ-Walsh		ϕ-Conitzer	
		ϕ	dist	ϕ	dist	ϕ	dist
1	f. skate	0.44	0.32	0.14	0.38	0.11	0.44
2	f. skate	0.23	0.24	0	0.36	0.02	0.56
3	aspen	0.68	0.18	0.3	0.16	0.23	0.27
4	f. skate	0.05	0.09	0	0.37	0	0.69
5	s. skate	0.46	0.11	0.15	0.18	0.16	0.35
6	irish	0.75	0.1	0.42	0.12	0.36	0.16
7	cities	0.93	0.06	0.69	0.06	0.63	0.06

over elections that do not correspond to vote distributions and hence are not captured by expected frequency matrices (as is the case, e.g., for the Euclidean models where candidates do not have fixed positions; see the work of Szufa et al. (2020) for examples of such models in the context of the map of elections): If we want to compute the distance of such a distribution, we sample sufficiently many elections and compute their average distance from the input one. However, it remains unclear how robust this approach is.

Approach. To fit our vote distributions to a given election, we compute the election's distance to the frequency matrices of $\mathcal{D}_{\mathrm{Mal}}^{v,\phi}$, ϕ-Conitzer, and ϕ-Walsh, for $\phi \in \{0, 0.001, \ldots, 1\}$. We select the distribution corresponding to the closest matrix.

Data. We consider elections from the real-world datasets used by Boehmer et al. (2021b). They generated 15 elections with 10 candidates and 100 voters (with strict preferences) from each of the eleven different real-world election datasets (so, altogether, they generated 165 elections, most of them from Preflib (Mattei & Walsh, 2013)). They used four datasets of political elections (from North Dublin (Irish), various non-profit and professional organizations (ERS), and city council elections from Glasgow and Aspen), four datasets of sport-based elections (from Tour de France (TDF), Giro d'Italia (GDI), speed skating, and figure skating) and three datasets with survey-based elections (from preferences over T-shirt designs, sushi, and cities). We present the results of our analysis for seven illustrative and particularly interesting elections in Table 1 and also include them in our skeleton map from Figure 3.

Basic Test. There is a standard maximum-likelihood estimator (MLE; based on Kemeny voting (Mandhani & Meilă, 2009)) that given an election provides the most likely dispersion parameter of the Mallows distribution that might have generated this election. To test our approach, we compared the parameters provided by our approach and by the MLE for our 165 elections and found them to be highly correlated (with Pearson correlation coefficient around 0.97). In particular, the average absolute difference between the dispersion parameter calculated by our approach and the MLE is only 0.02. See Appendix F for details.

Fitting Real-World Elections. Next, we consider the capabilities of $\mathcal{D}_{\mathrm{Mal}}^{v,\phi}$, ϕ-Conitzer, and ϕ-Walsh to fit the real-world elections of Boehmer et al. (2021b). Overall, we find that these three vote distributions have some ability to capture the considered elections, but it certainly is not perfect. Indeed, the average normalized distance of these elections to the frequency matrix of the closest distribution is 0.14. To illustrate that some distance is to be expected here, we mention that the average distance of an election sampled from impartial culture ($\mathcal{D}_{\mathrm{IC}}$, with 10 candidates and 100 voters) to the distribution's expected frequency matrix is 0.09 (see Appendix E.4 for a discussion of this and how it may serve as an estimator for the "variance of a distribution"). There are also some elections that are not captured by any of the considered distributions to an acceptable degree; examples of this are elections nr. 1 and nr. 2, which are at distance at least 0.32 and 0.25 from all our distributions, respectively. Remarkably, while coming from the same dataset, elections nr. 1 and nr. 2 are still quite different from each other and, accordingly, the computed dispersion parameter is also quite different. It remains a challenge to find distributions capturing such elections.

Comparing the power of the three considered models, nearly all of our elections are best captured by the Mallows model rather than ϕ-Conitzer or ϕ-Walsh. There are only twenty elections that are closer to ϕ-Walsh or ϕ-Conitzer than to a Mallows model (election nr. 3 is the most extreme example), and, unsurprisingly, both ϕ-Walsh and ϕ-Conitzer perform particularly badly at capturing elections close to ID (see election nr. 4). That is, ϕ-Conitzer and ϕ-Walsh are not needed to ensure good coverage of the space of elections; the average normalized distance of our elections to the closest Mallows model is only 0.0007 higher than their distance to the closest distribution (elections nr. 3-6 are three examples of elections which are well captured by the Mallows model and distributed over the entire map).[2] Nevertheless, ϕ-Walsh is also surprisingly powerful, as the average normalized distance of our elections to the closest ϕ-Walsh distribution is only 0.03 higher than their distance to the closest distribution (however, this might be also due to the fact that most of the considered real-world elections fall into the same area of the map, which ϕ-Walsh happens to capture particularly well (Boehmer et al., 2021b)). ϕ-Conitzer performs considerably worse: there are only three elections for which it produces a (slightly) better result than ϕ-Walsh.

Moreover, our results also emphasize the complex nature of the space of elections: Election nr. 7 is very close to $\mathcal{D}_{\mathrm{Mal}}^{v,0.95}$, hinting that its votes are quite chaotic. At the same time, this election is very close to 0.63-Conitzer and 0.69-Walsh distributions, which suggests at least a certain level of structure among its votes (because votes from Conitzer and Walsh distributions are very structured, and the Mallows filter with dispersion between 0.63 and 0.69 does not destroy this structure fully). However, as witnessed by the fact that the frequency matrix of GS/balanced (which is highly structured) is UN, such phenomena can happen. Lastly, note that most of our datasets are quite "homogenous", in that the closest distributions for elections from the dataset are similar and also at a similar distance. However, there are also clear exceptions, for instance, elections nr. 1 and nr. 4 from the figure skating dataset. Moreover, there are two elections from the speed skating dataset where one election is captured best by $\mathcal{D}_{\mathrm{Mal}}^{v,0.76}$ and the other by $\mathcal{D}_{\mathrm{Mal}}^{v,0.32}$.

6 Summary

We have computed the frequency matrices (Szufa et al., 2020; Boehmer et al., 2021b) of several well-known distributions of votes. Using them, we have drawn a "skeleton map", which shows how these distributions relate to each other, and we have analyzed its properties. Moreover, we have demonstrated how our results can be used to fit vote distributions to capture real-world elections.

For future work, it would be interesting to compute the frequency matrices of further popular vote distributions, such as the Plackett–Luce model (we conjecture that its frequency matrix is computable in polynomial time). It would also be interesting to use our approach to fit more complex models, such as mixtures of Mallows models, to real-world elections. Further, it may be interesting to use expected frequency matrices to reason about the asymptotic behavior of our models. For example, it might be possible to formally show where, in the limit, do the matrices of our models end up on the map as we increase the number of candidates.

Acknowledgments

NB was supported by the DFG project MaMu (NI 369/19) and by the DFG project ComSoc-MPMS (NI 369/22). This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 101002854).

[2] For each election, we also computed the closest frequency matrix of two mixtures of Mallows models with reversed central votes p-$\mathcal{D}_{\mathrm{Mal}}^{v,\phi,\psi}$ using our approach. However, this only decreased the average minimum distance by around 0.02, with the probability p of flipping the central vote being (close to) zero for most elections.

References

Awasthi, P., Blum, A., Sheffet, O., and Vijayaraghavan, A. Learning mixtures of ranking models. In *Proceedings of NIPS-2014*, pp. 2609–2617, 2014.

Betzler, N., Bredereck, R., and Niedermeier, R. Theoretical and empirical evaluation of data for exact Kemeny rank aggregation. *Autonomous Agents and Multiagent Systems*, 28(5):721–748, 2014.

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

Boehmer, N. and Schaar, N. Collecting, classifying, analyzing, and using real-world elections. Technical Report arXiv:2204.03589 [cs.GT], arXiv.org, 2022.

Boehmer, N., Bredereck, R., Faliszewski, P., and Niedermeier, R. Winner robustness via swap- and shift-bribery: Parameterized counting complexity and experiments. In *Proceedings of IJCAI-2021*, pp. 52–58, 2021a.

Boehmer, N., Bredereck, R., Faliszewski, P., Niedermeier, R., and Szufa, S. Putting a compass on the map of elections. In *Proceedings of IJCAI-2021*, pp. 59–65, 2021b.

Boehmer, N., Faliszewski, P., Niedermeier, R., Szufa, S., and Was, T. Understanding distance measures among elections. In *Proceedings of IJCAI-2022*, pp. 102–108. ijcai.org, 2022.

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

Brandt, F., Conitzer, V., and Endriss, U. Computational social choice. In Weiß, G. (ed.), *Multiagent Systems*. MIT Press, 2013.

Conitzer, V. Computing Slater rankings using similarities among candidates. In *Proceedings of AAAI-2006*, pp. 613–619, July 2006.

Conitzer, V. Eliciting single-peaked preferences using comparison queries. *Journal of Artificial Intelligence Research*, 35:161–191, 2009.

Conitzer, V., Sandholm, T., and Lang, J. When are elections with few candidates hard to manipulate? *Journal of the ACM*, 54(3):Article 14, 2007.

Dwork, C., Kumar, R., Naor, M., and Sivakumar, D. Rank aggregation methods for the web. In *Proceedings of WWW-2001*, pp. 613–622, 2001.

Elkind, E., Faliszewski, P., and Slinko, A. Clone structures in voters' preferences. In *Proceedings of EC-2012*, pp. 496–513, 2012.

Elkind, E., Lackner, M., and Peters, D. Structured preferences. In Endriss, U. (ed.), *Trends in Computational Social Choice*. AI Access Foundation, 2017.

Elkind, E., Lackner, M., and Peters, D. Preference restrictions in computational social choice: A survey, May 2022.

Faliszewski, P., Lackner, M., Peters, D., and Talmon, N. Effective heuristics for committee scoring rules. In *Proceedings of AAAI-2018*, pp. 1023–1030, 2018.

Faliszewski, P., Karpov, A., and Obraztsova, S. The complexity of election problems with group-separable preferences. In *Proceedings of IJCAI-2020*, pp. 203–209, 2020.

Gormley, I. C. and Murphy, T. B. Exploring voting blocs within the irish electorate: A mixture modeling approach. *Journal of the American Statistical Association*, 103(483):1014–1027, 2008.

Guiver, J. and Snelson, E. L. Bayesian inference for plackett-luce ranking models. In *Proceedings of ICML-2009*, pp. 377–384. ACM, 2009.

Hemaspaandra, E., Hemaspaandra, L., and Rothe, J. Exact analysis of Dodgson elections: Lewis Carroll's 1876 voting system is complete for parallel access to NP. *Journal of the ACM*, 44(6): 806–825, 1997.

Hunter, D. R. MM algorithms for generalized bradley-terry models. *The Annals of Statistics*, 32(1): 384–406, 2004.

Inada, K. A note on the simple majority decision rule. *Econometrica*, 32(32):525–531, 1964.

Karpov, A. On the number of group-separable preference profiles. *Group Decision and Negotiation*, 28(3):501–517, 2019.

Kenig, B. and Kimelfeld, B. Approximate inference of outcomes in probabilistic elections. In *Proceedings of AAAI-2021*, pp. 2061–2068, 2019.

Lu, T. and Boutilier, C. Effective sampling and learning for mallows models with pairwise-preference data. *Journal of Machine Learning Research*, 15(1):3783–3829, 2014.

Mallows, C. Non-null ranking models. *Biometrica*, 44:114–130, 1957.

Mandhani, B. and Meilă, M. Tractable search for learning exponential models of rankings. *Journal of Machine Learning Research—Proceedings Track*, 5:392–399, 2009.

Mattei, N. and Walsh, T. Preflib: A library for preferences. In *Proceedings of ADT-2013*, pp. 259–270, 2013.

Meila, M. and Chen, H. Dirichlet process mixtures of generalized mallows models. In *Proceedings of UAI-2010*, pp. 358–367. AUAI Press, 2010.

Minka, T. Power EP. Technical report, Microsoft Research, 2004.

Murphy, T. B. and Martin, D. Mixtures of distance-based models for ranking data. *Computational Statistics & Data Analysis*, 41(3-4):645–655, 2003.

OEIS Foundation Inc. The on-line encyclopedia of integer sequences, 2020. URL http://oeis.org/A008302.

Szufa, S., Faliszewski, P., Skowron, P., Slinko, A., and Talmon, N. Drawing a map of elections in the space of statistical cultures. In *Proceedings of AAMAS-2020*, pp. 1341–1349, 2020.

Vitelli, V., Sørensen, Ø., Crispino, M., Frigessi, A., and Arjas, E. Probabilistic preference learning with the Mallows rank model. *Journal of Machine Learning Research*, 18:158:1–158:49, 2017.

Walsh, T. Generating single peaked votes. Technical Report arXiv:1503.02766 [cs.GT], arXiv.org, 2015.

Wang, J., Sikdar, S., Shepherd, T., Zhao, Z., Jiang, C., and Xia, L. Practical algorithms for multi-stage voting rules with parallel universes tiebreaking. In *Proceedings of AAAI-2021*, pp. 2189–2196, 2019.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [No]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [No]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Missing Material From Section 3.2

Theorem 3.2. *There is an algorithm that given a tree $\mathcal{T}$ computes $\#\mathrm{freq}(D^{\mathcal{T}}_{\mathrm{GS}})$ using polynomially many arithmetic operations with respect to the number of nodes in $\mathcal{T}$.*

Proof. Let $\mathcal{T}$ be the input tree, and let $C = \{c_1, \ldots, c_m\}$ be its candidate set. We will give an algorithm for computing the probability that a given candidate c_j appears in position $i \in [m]$ in a vote sampled from $\mathcal{D}^{\mathcal{T}}_{\mathrm{GS}}$.

Let x be some node of $\mathcal{T}$ (either internal or a leaf). Let $\mathcal{T}_x$ be the tree obtained from $\mathcal{T}$ by deleting all descendants of x, so that x becomes a leaf, and for each subset S of internal nodes of $\mathcal{T}_x$, let $\mathcal{T}^S_x$ be the ordered tree obtained by starting with $\mathcal{T}_x$ and reversing the nodes in the set S. For each $t \in \{0\} \cup [m-1]$ we define $f(x,t)$ to be the probability that if we reversed each internal node of $\mathcal{T}_x$ with probability $1/2$ then x would be preceded by exactly t candidates in the frontier of the resulting ordered tree. We compute $f(x,t)$ using dynamic programming.

Let $root$ be the root of $\mathcal{T}$. Then $f(root, 0) = 1$, and $f(root, t) = 0$ for $t \in [1, m-1]$. Next, let x be some node of $\mathcal{T}$ other than the root, let p be the parent of x, and let ℓ and r be the number of leaves that are descendants of x's left siblings and x's right siblings in $\mathcal{T}$, respectively. We claim that for each $t \in \{0\} \cup [m-1]$ we have:

$$f(x,t) = \tfrac{1}{2} f(p, t-\ell) + \tfrac{1}{2} f(p, t-r).$$

To see why this formula is correct, observe that if $p \notin S$ then x appears in position t in the frontier of $\mathcal{T}^S_x$ if and only if p appears in position $t-\ell$ in the frontier of $\mathcal{T}^S_p$: indeed, in the frontier of $\mathcal{T}^S_x$ the node x appears after all predecessors of p in the frontier of $\mathcal{T}^S_p$ as well as after the ℓ leaves that are the descendants of x's left siblings in $\mathcal{T}$. Similarly, if $p \in S$ then x appears in position t in the frontier of $\mathcal{T}^S_x$ if and only if p appears in position $t-r$ in the frontier of $\mathcal{T}^{S\setminus\{p\}}_p$: indeed, in the frontier of $\mathcal{T}^S_x$ the node x appears after all predecessors of p in the frontier of $\mathcal{T}^{S\setminus\{p\}}_p$ as well as after the r leaves that are the descendants of x's right siblings in $\mathcal{T}$. Since p is reversed with probability $\frac{1}{2}$, the recurrence follows.

The above formula and standard dynamic programming allow us to compute all the values of f using $O(m^2)$ arithmetic operations (note that there are at most $O(m)$ internal nodes). To complete the proof, observe that the probability that candidate c_j ends up in position i is $f(c_j, i-1)$. □

B Missing Material From Section 3.3

We use the following lemma, which is implicit in the work of Faliszewski et al. (2020).

Lemma B.1. *Let $\mathcal{T} = \mathrm{CP}(c_1, \ldots, c_m)$. A ranking v over $\{c_1, \ldots, c_m\}$ belongs to the support of $\mathcal{D}^{\mathcal{T}}_{\mathrm{GS}}$ if and only if there exists a subset $C' \subseteq C \setminus \{c_m\}$ such that in v:*

1. *(1) all alternatives in C' are ranked above c_m and all alternatives in $(C \setminus \{c_m\}) \setminus C'$ are ranked below c_m;*
2. *(2) for all $c_i, c_j \in C'$ with $i < j$ the alternative c_i is ranked above c_j;*
3. *(3) for all $c_i, c_j \notin C' \cup \{c_m\}$ with $i < j$ the alternative c_i is ranked below c_j.*

That is, in v the alternatives in C' appear in the increasing order of indices, followed by c_m, followed by the remaining alternatives in the decreasing order of indices, i.e., the sequence of indices in v is "single-peaked". Using this observation, we establish a bijection between the votes in the support of $\mathcal{D}^{\mathcal{T}}_{\mathrm{GS}}$ and single-peaked votes.

Theorem 3.3. *Given a ranking v over $C = \{c_1, \ldots, c_m\}$, let $\widehat{v}$ be another ranking over C such that, for each $j \in [m]$, if c_j is ranked in position i in v then c_i is ranked in position $m-j+1$ in $\widehat{v}$. Suppose that v is in the support of $\mathcal{D}^{\mathcal{T}}_{\mathrm{GS}}$, where $\mathcal{T} = \mathrm{CP}(c_1, \ldots, c_m)$. Then $\widehat{v}$ is single-peaked with respect to $c_1 \lhd \cdots \lhd c_m$.*

Proof. Suppose that c_m is ranked in position z in v. Then c_z is ranked first in $\widehat{v}$.

Consider two candidates c_x, c_y with $x < y < z$. We will prove that in $\widehat{v}$ candidate c_y is ranked above c_x, i.e., $\mathrm{pos}_{\widehat{v}}(c_x) > \mathrm{pos}_{\widehat{v}}(c_y)$. Let $k = \mathrm{pos}_{\widehat{v}}(c_x)$, $\ell = \mathrm{pos}_{\widehat{v}}(c_y)$. Then, in v, alternative c_{m-k+1} is ranked in position x and alternative $c_{m-\ell+1}$ is ranked in position y. As we have $x < y < z$ and v is sampled from $\mathcal{D}^{\mathcal{T}}_{\mathrm{GS}}$, by Lemma B.1 we have $m - k + 1 < m - \ell + 1$, and hence $k > \ell$.

Similarly, if we have two alternatives c_x, c_y with $z < y < x$, we can show that $\mathrm{pos}_{\widehat{v}}(c_x) > \mathrm{pos}_{\widehat{v}}(c_y)$. Thus, $\widehat{v}$ is single-peaked with respect to $c_1 \lhd \cdots \lhd c_m$, as claimed. □

C Missing Material From Section 3.4

Recall that our candidate set is $C = \{c_1, \dots, c_m\}$ and the societal axis is $c_1 \lhd \cdots \lhd c_m$. We consider the Conitzer distribution. Let $f(i, j)$ be the probability that in a sampled vote candidates $c_i, \dots, c_j$ appear in the top $j - i + 1$ positions. Next we find the values of $f(i, j)$ for all $i, j \in [m]$, and using them we establish the frequency matrix for the Conitzer distribution.

Proposition C.1. *Let i, j be two integers with $1 < i \le j < m$. Then $f(\ell, \ell) = 1/m$ for all $\ell \in [m]$, $f(1, m) = 1$, $f(i, j) = 1/m$, $f(1, i) = (i+1)/2m$, and $f(j, m) = (m-j+2)/2m$.*

Proof. The quantity $f(\ell, \ell)$ is simply the probability that c_ℓ is ranked first, so we have $f(\ell, \ell) = 1/m$ by the definition of the Conitzer distribution.

The equality $f(1, m) = 1$ is immediate from the definition of $f(i, j)$.

To show that $f(i, j) = 1/m$, we give a proof by induction on $j - i$. The base step holds because for each $\ell \in [m]$ we have $f(\ell, \ell) = 1/m$. Assume that for all integers x, y such that $1 < x \le y < m$ and $y - x < j - i$ we have $f(x, y) = 1/m$. The only way for candidates $c_i, \dots, c_j$ to be ranked in top $j - i + 1$ positions under the Conitzer model is that, while generating the vote, we placed candidates $c_i, \dots, c_{j-1}$ in top $j - i$ positions and then extended the vote with c_j (the probability of this latter step is $1/2$), or we placed candidates $c_{i+1}, \dots, c_j$ in top $j - i$ positions and then extended the vote with c_i (again, the probability of the latter step is $1/2$). Thus:

$$f(i, j) = \tfrac{1}{2} f(i, j-1) + \tfrac{1}{2} f(i+1, j) = \tfrac{1}{2m} + \tfrac{1}{2m} = \tfrac{1}{m}.$$

This proves the claim for $f(i, j)$.

Next we show that $f(1, i) = (i+1)/2m$, using induction over $i < m$. For $i = 1$ we have $f(1, 1) = 1/m = 2/2m$, so our claim holds. Assume that it holds up to $i - 1$. There are only two ways in which candidates $c_1, \dots, c_i$ can be ranked in the top i positions: Either we first place $c_1, \dots c_{i-1}$ in top $i - 1$ positions and then extend the vote with c_i (the latter step has probability 1, since there is no candidate to the left of c_1), or we first place candidates $c_2, \dots, c_i$ in the top $i - 1$ positions and then extend the vote with c_1 (the latter step has probability $1/2$, since the vote could also be extended with c_{i+1}). Thus, we have:

$$f(1, i) = f(1, i-1) + \tfrac{1}{2} f(2, i) = \tfrac{i}{2m} + \tfrac{1}{2m} = \tfrac{i+1}{2m}.$$

This completes the proof for $f(1, i)$. The expression for $f(j, m)$ can be derived by symmetry: we have $f(j, m) = f(1, m - j + 1)$. □

Theorem 3.5. *Let $c_1 \lhd \cdots \lhd c_m$ be the societal axis, where m is an even number, and let v be a random vote sampled from $\mathcal{D}^{\mathrm{Con}}_{\mathrm{SP}}$ for this axis. For $j \in [\frac{m}{2}]$ and $i \in [m]$ we have:*

$$\mathbb{P}[\mathrm{pos}_v(c_j) = i] = \begin{cases} 2/2m & \text{if } i < j, \\ (j+1)/2m & \text{if } i = j, \\ 1/2m & \text{if } j < i < m - j + 1, \\ (m-j+1)/2m & \text{if } i = m - j + 1, \\ 0 & \text{if } i + j > m. \end{cases}$$

Further, for each candidate $c_j \in C$ and each position $i \in [m]$ we have $\mathbb{P}[\mathrm{pos}_v(c_j) = i] = \mathbb{P}[\mathrm{pos}_v(c_{m-j+1}) = i]$.

Proof. Consider a candidate c_j, $j \in [\frac{m}{2}]$, and a position $i \in [m]$. We proceed by case analysis:

1. If $i < j$, then there are two ways to generate a vote with c_j in position i: Either candidates $c_{j+1}, \dots, c_{i+j-1}$ are ranked in the top $i-1$ positions or candidates $c_{j-i+1}, \dots, c_{j-1}$ are ranked in the top $i-1$ positions. In both cases, c_j is ranked in the i-th position with probability $1/2$ (indeed, we have $i + j < m/2 + m/2 = m$, so in the former case both c_j and c_{i+j} could have been placed in position i, and also $j - i \geq 1$, so in the latter case both c_j and c_{j-i} could have been placed in position i). Thus we have:

$$\begin{aligned}\mathbb{P}[\mathrm{pos}_v(c_j) = i] &= \tfrac{1}{2} \cdot f(j+1, i+j-1) \\ &\quad + \tfrac{1}{2} \cdot f(j-i+1, j-1) = \tfrac{1}{2m} + \tfrac{1}{2m} = \tfrac{1}{m}.\end{aligned}$$

2. If $i = j$, then either candidates $c_1, \dots, c_{j-1}$ are ranked in the top $j-1$ positions and the vote is extended with c_j (with probability 1), or candidates $c_{j+1}, \dots, c_{2j-1}$ are ranked in the top $j-1$ positions and the vote is extended with c_j (with probability $1/2$). Thus, we have:

$$\begin{aligned}\mathbb{P}[\mathrm{pos}_v(c_j) = j] &= f(1, j-1) + \tfrac{1}{2} \cdot f(j+1, 2j-1) \\ &= \tfrac{j}{2m} + \tfrac{1}{2m} = \tfrac{j+1}{2m}.\end{aligned}$$

3. If $j < i < m - j + 1$ then there is only one possibility for c_j to be ranked i-th: It must be that candidates $c_{j+1}, \dots, c_{j+i-1}$ are ranked in the top $i-1$ positions and the vote is extended with c_j (which happens with probability $\frac{1}{2}$ because $j + i \leq m$).[3] Thus, we have:

$$\mathbb{P}[\mathrm{pos}_v(c_j) = i] = \tfrac{1}{2} \cdot f(j+1, i+j-1) = \tfrac{1}{2m}.$$

4. If $i = m - j + 1$, the analysis is similar to the case $i = j$: candidates $c_{j+1}, \dots, c_m$ must be ranked in the top $m - j$ positions, in which case c_j gets ranked in the $(m-j+1)$-st position (with probability 1). Thus, we have:

$$\mathbb{P}[\mathrm{pos}_v(c_j) = m - j + 1] = f(j+1, m) = \frac{m-j+1}{2m}.$$

5. If $i > m - j + 1$ then $\mathbb{P}[\mathrm{pos}_v(c_j) = i] = 0$, because both to the left of c_j and to the right of c_j there are fewer than $i - 1$ candidates.

The fact that for each candidate $c_j \in C$ and each position $i \in [m]$ we have $\mathbb{P}[\mathrm{pos}_v(c_j) = i] = \mathbb{P}[\mathrm{pos}_v(c_{m-j+1}) = i]$ follows directly from the symmetry of the Conitzer distribution and the fact that m is even. □

D Missing Material From Section 3.5

Proposition D.1 (OEIS Foundation Inc. (2020))**.** *There is an algorithm that computes $S(m, k)$ using at most polynomially many operations.*

Proof. First, we note that for each $m' \in [m]$ we have $S(m', 0) = 1$. Further, for each $m' \in [m]$ and $k' \in [k]_0$ the following recursion holds:

$$S(m', k') = S(m', k'-1) + S(m'-1, k') - S(m'-1, k'-m').$$

Using these two facts and standard dynamic programming, we can compute $S(m, k)$ using $O(mk)$ arithmetic operations. Since k is at most $O(m^2)$, the running time is at most $O(m^3)$. □

Lemma D.2. *There is an algorithm that computes $T(m, k, j, i)$ in polynomial time with respect to m.*

Proof. Our algorithm is based on dynamic programming. Fix some $m > 0$, $k \in [m(m-1)/2]$, and $j, i \in [m]$. We claim that:

$$T(m, k, m, i) = S(m-1, k-(m-i)).$$

[3] It is impossible for the candidates from the left side of c_j to take the top $i-1$ positions because there are fewer than $i - 1$ of them.

Indeed, let v be a vote over $C(m)$ that ranks c_m in position i, and let v' be its restriction to $\{c_1, \dots, c_{m-1}\}$. Then v can be obtained from v' by inserting c_m right behind the candidate in position $i-1$. If v' is at swap distance k' from v^*_{m-1}, then the resulting vote is at swap distance $k' + (m-i)$ from v^*_m, since c_m contributes $m-i$ additional swaps. Thus, $T(m,k,m,i)$ is equal to the number of votes in $\mathcal{L}(C(m-1))$ at swap distance $k-(m-i)$ from v^*_{m-1}.

Next, we claim that for each $j < m$, we have:

$$\begin{aligned} T(m,k,j,i) = &\textstyle\sum_{\ell=i+1}^{m} T[m-1, k-(m-\ell), j, i] \\ &+ \textstyle\sum_{\ell=1}^{i-1} \quad T[m-1, k-(m-\ell), j, i-1]. \end{aligned}$$

To see why this holds, again consider inserting c_m at some position in a vote v' over $\{c_1, \dots c_{m-1}\}$. Candidate c_j will end up in position i in the resulting vote if (1) c_j was in position i in v' and c_m was inserted after c_j, or if (2) c_j was in position $i-1$ in v' and c_m was inserted ahead of c_j. Considering all positions in which c_m can be inserted, we obtain the above equality.

Using the above equalities and the fact that $T(1,0,1,1) = 1$ (as there is just a single vote over $C(1)$), we can compute $T(m,k,j,i)$ by dynamic programming; our algorithm runs in polynomial time with respect to m. □

E Missing Material From Section 4

E.1 Distance to Compass Matrices

We analyze the distances between our matrices for different numbers of candidates. In Figure 5 we show these distances for the Conitzer and Walsh matrices and the compass matrices: For Conitzer, they are nearly constant, and for Walsh they vary significantly. Indeed, the more candidates we have, the closer the Walsh matrix is to ID (e.g., for 10 candidates their distance is 0.44, and for 300 candidates it is 0.09).

Figure 6 depicts the distance between the frequency matrix for GS/caterpillar and the four compass matrices, for varying number of candidates. As for the matrix for the Walsh model, its distance to the compass matrices changes as the number of candidates increases: The matrix moves closer and closer to AN.

In Figure 7, we display the distance between the frequency matrix for the Mallows model for different values of the dispersion parameter ϕ and the compass matrices (in contrast to the previous figures, we only consider up to 100 candidates, as for more than 100 candidates computing the matrix for the Mallows model becomes very memory-consuming). Independent of the chosen value of the dispersion parameter, the distance of the respective matrix to the four compass matrices changes significantly when we increase the number of candidates. In fact, for any fixed dispersion parameter ϕ, the resulting matrix will always move closer and closer to ID as the number of candidates increases.

In contrast, if we use the normalized version of the Mallows model, the matrices remain more or less at a constant distance from the compass matrices. Figure 8 shows the distance of the frequency matrix of the normalized version of the Mallows model for different values of $\mathrm{norm}\text{-}\phi$, as the number of candidates increases.

E.2 Distances of Pairs of Vote Distributions on the Skeleton Map

As we have observed above, some vote distributions stay more or less at constant normalized positionwise distance from the four compass matrices. This raises the question whether these matrices also stay at a constant normalized positionwise distance from each other. This would imply that the data on which the skeleton map is based is independent of the number of candidates, and thus that the map is likely to look very similar for different numbers of candidates. To check this, we conducted the following experiment. We put together a set of vote distributions/matrices that do not structurally change when increasing the number of candidates (like the change happening for the Walsh model). First, we add the four compass matrices and the frequency matrix of the Conitzer model. Second, we add the frequency matrices of different variants of the Mallows model (similar as on the skeleton map as described in Section 4): the normalized Mallows model, the normalized Mallows model where the central vote is reversed with probability $1/2$, the normalized Mallows model

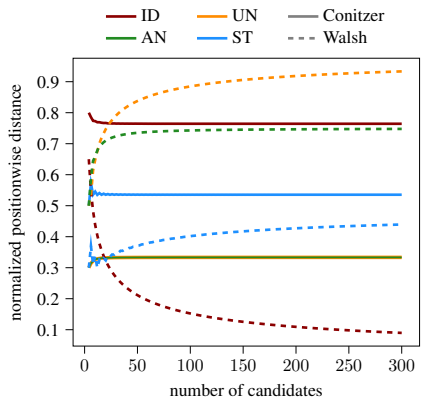

Figure 5: Normalized positionwise distance between the Conitzer [Walsh] matrix and the compass matrices in solid [dashed] lines, for varying number of candidates.

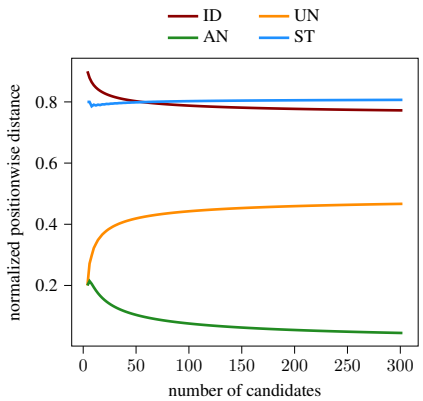

Figure 6: Normalized positionwise distance between the frequency matrix of the GS/caterpillar distribution and the compass matrices, for varying number of candidates.

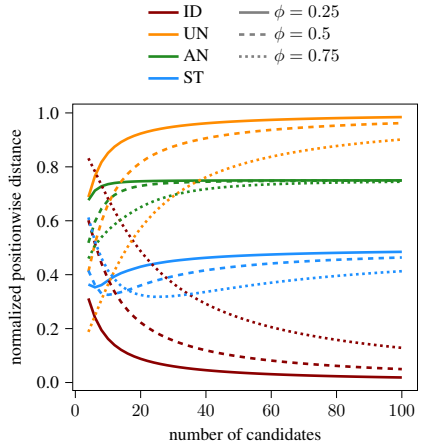

Figure 7: For different values of ϕ, normalized positionwise distance between the frequency matrix of the Mallows distribution and the compass matrices, for varying number of candidates.

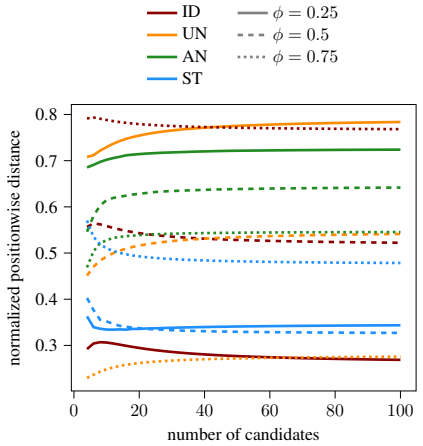

Figure 8: For different values of $\text{norm-}\phi$, normalized positionwise distance between the frequency matrix of the normalized Mallows distribution and the compass matrices, for varying number of candidates.

where the central vote is reversed with probability $1/4$, and the distribution where we first sample a vote v from the Conitzer distribution and then sample the final vote from the normalized Mallows model with v as the central vote. For each of these variants, we include their frequency matrix for $\text{norm-}\phi \in \{0.2, 0.4, 0.6, 0.8\}$.

For each pair of matrices from the created set, we compute their normalized positionwise distance for 100 candidates. Subsequently, for $m \in \{4, 6, \ldots, 98, 100\}$ candidates, we compute the normalized positionwise distance of the frequency matrices of the two considered models for this number of candidates as well as the absolute and relative difference between their normalized distance for m and 100 candidates (where we normalize by their normalized distance for 100 candidates). Finally, for each $m \in \{4, 6, \ldots, 98, 100\}$, we take the maximum over the computed absolute/relative differences for all pairs of matrices. In Figure 9, we present these maxima for all considered values of m. Examining the maximum absolute difference (the blue line in Figure 9), what stands out is that for 20 or more candidates the normalized positionwise distance of any pair of considered vote distributions/matrices differs only by at most 0.0287 from the pair's normalized positionwise distance for 100 candidates (for 10 or more candidates the error is at most 0.0547). As the diameter of our

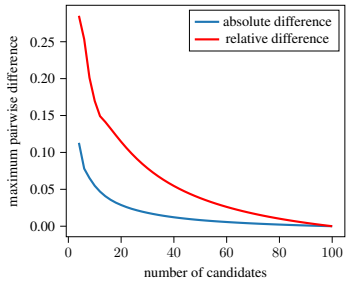

Figure 9: Results from our experiments described in Appendix E.2

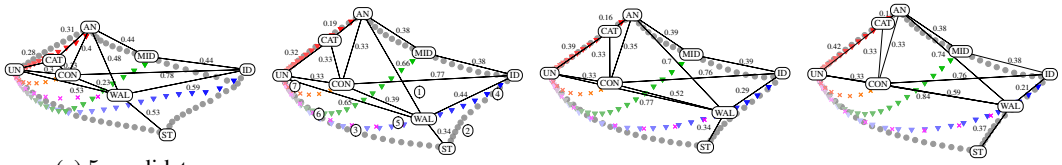

(a) 5 candidates (b) 10 candidates (c) 25 candidates (d) 50 candidates

Figure 10: Skeleton map for different number of candidates.

space has at least length 1, this change is quite small, and the considered vote distributions indeed remain at nearly the same distance for more than 20 candidates. Considering the relative difference (the red line in Figure 9), the picture appears to be a bit worse: for 20 or more candidates, the normalized positionwise distance of any pair of considered vote distributions/matrices differs at most by 11.44% from their normalized positionwise distance for 100 candidates. Nevertheless, this value is still relatively low, indicating an overall high robustness of the normalized positionwise distances of each pair of considered distributions with respect to the number of candidates.

E.3 Skeleton Map for Different Number of Candidates

After we have provided various arguments for why large parts of the skeleton map are presumably quite robust with respect to changing the number of candidates in the previous two subsections, in Figure 10, we present the skeleton map for $5/10/25/50$ candidates. While the map for 5 candidates looks a bit different from the other maps, the maps for 10, 25, and 50 candidates differ only in that the frequency matrix for GS/caterpillar moves closer to AN and that the frequency matrix for the Walsh model moves closer to ID (both phenomena that we have already observed earlier).

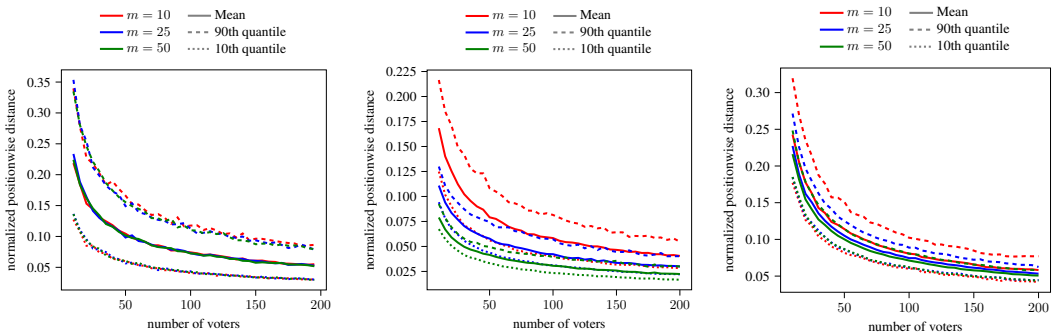

(a) Conitzer model (b) Walsh model (c) GS/caterpillar model

Figure 11: For different vote distributions, behavior of the normalized positionwise distance between elections sampled from this distributions and the distribution's frequency matrix, for 10/25/50 candidates and between 10 and 200 voters.

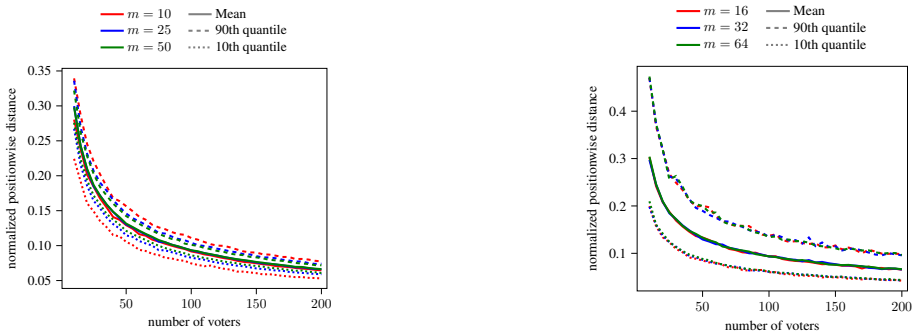

(a) Impartial Culture

(b) GS/balanced model

Figure 12: For different vote distributions, behavior of the normalized positionwise distance between elections sampled from this distributions and the distribution's frequency matrix, for 10/25/50 candidates and between 10 and 200 voters.

E.4 Variance of a Vote Distribution

As argued in Section 3.1, a frequency matrix of a distribution can be interpreted as a matrix of an "ideal" election sampled from this distribution. In this section we ask how far, on average, the elections sampled from our distributions land from the "ideal" ones. This distance may also serve as a measure of "diversity" for elections sampled from a given distribution.

For each of our distributions, we consider elections with 10, 25, and 50 candidates, and vary the number of voters from 10 to 200 (with a step of 5). For each combination of these parameters, we sample 600 elections and, for each election, compute the positionwise distance between its frequency matrix and the matrix of the respective distribution.[4] We show the results in Figures 11 and 12. As expected, for all vote distributions, increasing the number of votes decreases the average distance of an election from the distribution's matrix (indeed, in the limit this distance is zero). What is more surprising, this distance does not seem to depend on the number of candidates for the Conitzer, IC, and GS/balanced distribution. For the Walsh distribution (and, to a lesser extent, for GS/caterpillar), the sampled elections get slightly closer to the respective matrix as we increase the number of candidates. Moreover, if we fix the number of candidates and voters, then for all our distributions the elections sampled from them are, roughly, at the same distance from the distribution's matrix. For an illustration of this effect consider Figure 11a for the Conitzer distribution; there, the 10th quantile (dotted) and the 90th quantile (dashed) only differ by a factor of four. Lastly, comparing the plots for IC and GS/balanced (Figure 12), who both have UN as their frequency matrix, the average distance of elections sampled form one of these two models to UN (which is their frequency matrix) is the same for both distributions. However, for IC, the 10th and 90th quantile of the distances of elections to UN are closer to the average than for GS/balanced, which indicates that IC produces in some sense less varied elections than GS/balanced.

In Figure 13 we compare the average distances between elections sampled from various vote distributions and the distribution's matrices (we fix the number of candidates to 50 and vary the number of voters). While, on average, IC elections and GS/balanced elections end up at nearly the same distance from UN (which is their frequency matrix), Conitzer elections and GS/caterpillar elections end up closer to their distribution's matrices, and for Walsh this effect is considerably stronger. Overall, it is remarkable that even for 200 voters, for the Conitzer, IC, GS/balanced, and GS/caterpillar, the average distance of a sampled election from the respective matrix is still above 0.05 (so at least 5% of the diameter of the whole space). We also performed the same experiment for the Mallows model with different values of the normalized dispersion parameter (see Figure 14): For a varying number of voters, we depict the average distance of 600 elections with 50 candidates sampled from Mallows model for different values of norm-ϕ to the distribution's frequency matrix. Quite intuitively, the more swaps we make to the central vote (i.e., the higher norm-ϕ is), the higher is the average distance of a sampled elections from the distribution's frequency matrix.

[4] For the GS/balanced distribution we consider 16, 32, and 64 candidates, as this model requires the number of candidates to be a power of two and we do not consider GS/flat trees, as this distribution is too simple.

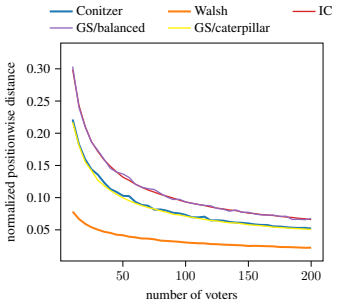

Figure 13: Average normalized positionwise distance between elections sampled from various voter distributions and the frequency matrices of the respective models, for 50 candidates (64 for GS) and between 10 and 200 voters.

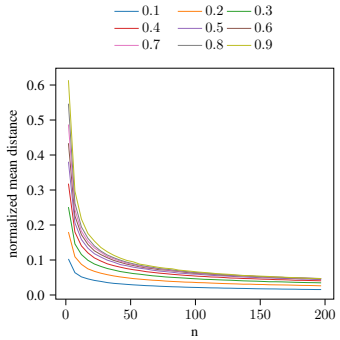

Figure 14: For different values of norm-ϕ, average normalized positionwise distance of elections with 50 candidates and between 10 and 200 voters sampled from the normalized Mallows model from the frequency matrix of the respective distribution.

It is interesting to contrast the data from Figure 11a with the map of Boehmer et al. (2021b) in Figure 2. Boehmer et al. considered 10 candidates and 100 voters. For these parameters, in Figure 11a we see that 10% of the elections are still farther from the Conitzer matrix than about 12% of the distance from UN to ID. This is roughly reflected by the size of the area taken by Conitzer elections in Figure 2. Similar observations hold for the other distributions too. While this might be a coincidence, it confirms the value of their map.

F Validation of Location Framework

Again, in this section, distances and dispersion parameters are always normalized. To validate whether our approach from Section 5 correctly identifies the "nature" of an election, we test its capabilities to find for a given election the dispersion parameter of the closest Mallows distribution with a single central vote. That is, given an election, we computed for all $\phi \in \{0, 0.001, \dots, 1\}$ its distance to $\mathcal{D}_{\text{Mal}}^{v,\phi}$ and returned the minimizing ϕ value. We compare the computed value to the maximum-likelihood estimator for the dispersion parameter of the underlying Mallows distribution computed from the Kemeny consensus ranking (Mandhani & Meilă, 2009) in two experiments.

We start by sampling for $\phi \in \{0, 0.05, \dots, 1\}$ an election with 10 candidates and 100 voters from Mallows model with dispersion parameter ϕ and computed estimates for the dispersion parameter based on our and the Kemeny approach. For all elections, the returned estimates differ by at most 0.01. Thus, the dispersion parameter returned by our approach is always very close to the maximum-likelihood estimate. However, the estimated dispersion parameter might deviate a bit from the originally used dispersion parameter: On average, the absolute difference between the dispersion parameter returned by our approach and the underlying dispersion parameter is 0.0179 with the maximum difference being 0.069; for Kemeny, the average is 0.016 and the maximum is 0.082. While it might seem surprising that the estimated dispersion parameter is different from the underlying one, recall from Appendix E.4 that elections sampled from a vote distribution typically have non-zero distance from the distribution's frequency matrix. To illustrate this idea, we can think of a Mallows distribution as a normal distribution placed in the space of elections with the dispersion parameter being its mean (in particular as soon as the dispersion parameter is greater zero, all elections have a non-zero probability of being sampled). So multiple Mallows distributions for different dispersion parameters translate to multiple partly overlapping normal distributions and it might as well happen that an election sampled from a Mallows distribution with one dispersion parameter is in fact closer to the mean of the Mallows distribution with a different dispersion parameter.

We also repeated the above experiment to measure the capabilities of our approach to estimate the parameters of a mixed Mallows distribution. For each pair of $p \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$ and $\phi \in \{0.05, 0.1, \dots, 0.95\}$ we sampled an election with 10 candidates and 100 voters from p-$\mathcal{D}_{\text{Mal}}^{v,\phi,\phi}$ (i.e., we sample a vote from the Mallows distribution with dispersion parameter ϕ and

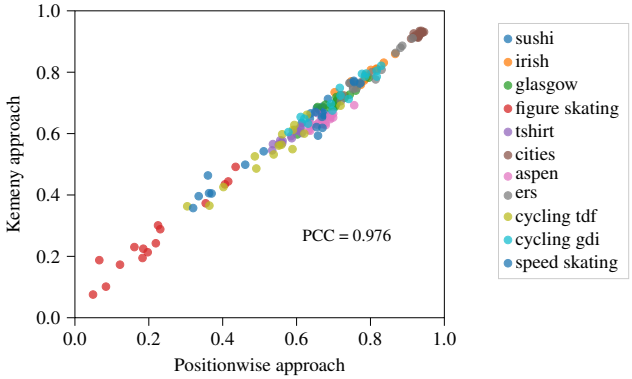

Figure 15: Correlation between the predicted dispersion parameter of our real-world data by our (positionwise) approach and by the maximum-likelihood Kemeny approach.

subsequently flip the sampled vote with probability p). Subsequently we computed for which value of $p \in \{0, 0.05, \ldots, 0.5\}$ and $\phi \in \{0, 0.05, \ldots, 1\}$ the distance of the sampled election is closest to the frequency matrix of the induced distribution p-$\mathcal{D}_{\text{Mal}}^{v,\phi,\phi}$ (the Kemeny consensus ranking can no longer be used here to provide a maximum-likelihood estimate). The average difference between the estimated and underlying dispersion parameter is 0.028 and the average difference between the estimated and underlying flipping probability is 0.068. For 63 out of the 95 elections is the difference between the estimated and underlying dispersion parameter and weight smaller equal 0.05 (which is the smallest non-zero difference). The error of the estimated dispersion parameter here is around three times larger than for Mallows elections with a single central vote. This can be explained by the fact that for mixtures of Mallows distributions, the "overlap" between different distributions is even larger; in fact, some paramterizations even result in the same distributions (e.g., for $\phi = 1$ all flipping probabilities result in the same distribution).

Second, while producing good estimates for elections that have been sampled from a Mallows distribution is a good sanity check, we are ultimately interested in computing to which distribution (unknown) real-world elections are closest. To do so, we again compare the estimated dispersion parameter for a Mallows distribution with a single central vote computed by our approach with the one estimated via the Kemeny consensus ranking (as described in the beginning of this section); however, this time instead of considering elections sampled from Mallows model, we examine 165 real-world elections used by Boehmer et al. (2021b) (see the data part of Section 5 for details on the dataset). The estimated dispersion parameters returned by both methods are highly correlated with a Pearson correlation coefficient of 0.976 and an average difference of 0.017, median difference of 0.0105, and maximum difference of 0.197, indicating the power of our approach. Interestingly, the correlation is particularly strong for larger dispersion parameters (see Figure Figure 15 for a plot showing the correlation between the two approaches). Together with the estimated normalized dispersion parameter, both approaches also return the central order v of the closest Mallows model, which are typically again quite similar: the average swap distance between the two estimators is 2.81 out of 45 possible swaps.