

Diversity, Agreement, and Polarization in Elections

Piotr Faliszewski¹, Andrzej Kaczmarczyk¹, Krzysztof Sornat², Stanisław Szufa¹, Tomasz Wąs³

¹AGH University, Poland

²IDSIA, USI-SUPSI, Switzerland

³Pennsylvania State University, PA, USA

faliszew@agh.edu.pl, andrzej.kaczmarczyk@agh.edu.pl,
krzysztof.sornat@idsia.ch, szufa@agh.edu.pl, twas@psu.edu

Abstract

We consider the notions of agreement, diversity, and polarization in ordinal elections (that is, in elections where voters rank the candidates). While (computational) social choice offers good measures of agreement between the voters, such measures for the other two notions are lacking. We attempt to rectify this issue by designing appropriate measures, providing means of their (approximate) computation, and arguing that they, indeed, capture diversity and polarization well. In particular, we present “maps of preference orders” that highlight relations between the votes in a given election and which help in making arguments about their nature.

1 Introduction

The notions of *agreement*, *diversity*, and *polarization* of a society with respect to some issue are intuitively quite clear. In case of agreement, most members of the society have very similar views regarding the issue, in case of diversity there is a whole spectrum of opinions, and in case of polarization there are two opposing camps with conflicting views and with few people taking middle-ground positions (more generally, if there are several camps, with clearly separated views, then we speak of *fragmentation*; see, for example, the collection of Dynes and Tierney [1994]). We study these three notions for the case of ordinal elections—that is, for elections where each voter has a preference order (his or her vote) ranking the candidates from the most to the least appealing one—and analyze ways of quantifying them.¹

Interestingly, even though agreement, diversity, and polarization seem rather fundamental concepts for understanding the state of a given society (see, for example, the papers in a special issue edited by Levin *et al.* [2021]), so far (computational) social choice mostly focused on the agreement-disagreement spectrum. Let us consider the following notion:

Given an election, the voters’ agreement index for candidates a and b is the absolute value of the difference

¹More extensive analysis as well as all proofs are available in the full version of our paper [Faliszewski *et al.*, 2023] and the code of our experiments is available at <https://github.com/Project-PRAGMA/diversity-agreement-polarization-IJCAI23>.

between the fraction of the voters who prefer a to b and the fraction of those with the opposite view. Hence, if all voters rank a over b (or, all voters rank b over a) then the agreement index for these candidates is equal to 1. On the other hand, if half of the voters report $a \succ b$ and half of them report $b \succ a$, then the index is equal to 0. The agreement index of the whole election is the average over the agreement indices of all the candidate pairs.

For an election E , we denote its agreement index as $A(E)$. Alcalde-Unzu and Vorsatz [2013] viewed this index as measuring voter cohesiveness—which is simply a different term for voter agreement—and provided its axiomatic characterization. Hashemi and Endriss [2014] focused on measuring diversity and provided axiomatic and experimental analyses of a number of election indices, including $1 - A(E)$. $1 - A(E)$ was also characterized axiomatically by Can *et al.* [2015], who saw it as measuring polarization; their point of view was that for each pair of candidates one can measure polarization independently. (In Section 2 we briefly discuss other election indices from the literature; generally, they are strongly interrelated with the agreement one).

Our view is that $1 - A(E)$ is neither a measure of diversity nor of polarization, but of disagreement. Indeed, it has the same, highest possible, value on both the antagonism election (AN), where half of the voters report one preference order and the other half reports the opposite one, and on the uniformity election (UN), where each possible preference order occurs the same number of times. Indeed, both these elections arguably represent extreme cases of disagreement. Yet, the nature of this disagreement is very different. In the former, we see strong polarization, with the voters split into two opposing groups, and in the latter we see perfect diversity of opinion. The fundamental difference between these notions becomes clear in the text of Levin *et al.* [2021] which highlights “the loss of diversity that extreme polarization creates” as a central theme of the related special issue. Our main goal is to design election indices that distinguish these notions.

Our new indices are based on what we call the k -KEMENY problem. In the classic KEMENY RANKING problem (equivalent to 1-KEMENY), given an election we ask for a ranking whose sum of swap distances to the votes is the smallest (a swap distance between two rankings is the number of swaps of adjacent candidates needed to transform one ranking into the other). The k -KEMENY problem is defined anal-

ogously, but we ask for k rankings that minimize the sum of each vote’s distance to the closest one (readers familiar with multiwinner elections [Faliszewski *et al.*, 2017] may think of it as the Chamberlin–Courant rule [Chamberlin and Courant, 1983] for committees of rankings rather than candidates). We refer to this value as the k -Kemeny distance. Unfortunately, the k -KEMENY problem is intractable—just like KEMENY RANKING [Bartholdi *et al.*, 1989; Hemaspaandra *et al.*, 2005]—so we develop multiple ways (such as fast approximation algorithms) to circumvent this issue.

Our polarization index is a normalized difference between the 1-Kemeny and 2-Kemeny distances of an election, and our diversity index is a weighted sum of the k -Kemeny distances for $k = 1, 2, 3, \dots$. The intuition for the former is that if a society is completely polarized (that is, partitioned into two equal-sized groups with opposing preference orders), then 1-Kemeny distance is the largest possible, but 2-Kemeny distance is zero. The intuition for the latter is that if a society is fully diverse (consists of all possible votes) then each k -Kemeny distance is non-negligible (we use weights for technical reasons). Since our agreement index can also be seen as a variant of the KEMENY RANKING problem, where we measure the distance to the majority relation, all these indices are based on similar principles.

To evaluate our indices, we use the “map of elections” framework of Szufa *et al.* [2020], Boehmer *et al.* [2021], and Boehmer *et al.* [2022], applied to a dataset of randomly generated elections. In particular, we find that our indices are correlated with the distances from several characteristic elections and, hence, provide the map with a semantic meaning. Additionally, we develop a new form of a map that visualizes the relations between the votes of a single election (the original maps visualized relations between several elections from a given dataset). We use this approach to get an insight regarding the statistical cultures used to generate our dataset and to validate intuitions regarding the agreement, diversity, and polarization of its elections. In our experiments, we focused on elections with a relatively small number of candidates (8 candidates and 96 voters). While we believe that our main conclusions extend to all sizes of elections, it would be valuable to check this (however, this would require quite extensive computation that, currently, is beyond our reach).

2 Preliminaries

For every number $k \in \mathbb{N}$, by $[k]$ we understand the set $\{1, \dots, k\}$. For two sets A and B such that $|A| = |B|$, by $\Pi(A, B)$ we mean the set of all bijections from A to B .

2.1 Elections

An *election* $E = (C, V)$ is a pair, where C is a set of *candidates* and V is a collection of *voters* whose preferences (or, *votes*) are represented as linear orders over C (we use the terms *vote* and *voter* interchangeably, depending on the context). For a vote v , we write $a \succ_v b$ (or, equivalently, $v: a \succ b$) to indicate that v *prefers* candidate a over candidate b . We also extend this notation to more candidates. For example, for candidate set $C = \{a, b, c\}$ by $v: a \succ b \succ c$ we mean that v ranks a first, b second, and c third. For two

candidates a and b from election E , by $p_E(a, b)$ we denote the fraction of voters in E that prefer a over b .

We will often speak of the following three characteristic elections, introduced by Boehmer *et al.* [2021] as “compass elections” (we assume candidate set $C = \{c_1, \dots, c_m\}$ here; Boehmer *et al.* [2021] also considered the fourth election, i.e., *stratification*, but it will not play an important role for us):

Identity (ID). In an identity election all votes are identical. We view this election as being in perfect agreement.

Antagonism (AN). In an antagonism election, exactly half of the voters have one preference order (for example, $c_1 \succ c_2 \succ \dots \succ c_m$) and the other half has the reversed one ($c_m \succ c_{m-1} \succ \dots \succ c_1$). We view this election as being perfectly polarized.

Uniformity (UN). A uniformity election contains the same number of copies of every possible preference order. We view this election as being perfectly diverse.

2.2 Kemeny Rankings and Swap Distance

For two votes u and v over a candidate set C , by $\text{swap}(u, v)$ we mean their *swap distance*, that is, the minimal number of swaps of consecutive candidates required to transform u into v . This value is also known as *Kendall’s τ* distance and is equal to the number of candidate pairs $a, b \in C$ such that $a \succ_u b$ but $b \succ_v a$. A *Kemeny ranking* of an election $E = (C, V)$ is a linear order over C that minimizes the sum of its swap distances to the votes from V [Kemeny, 1959]. It is well known that computing a Kemeny ranking is NP-hard [Bartholdi *et al.*, 1989] and, more precisely, Θ_2^P -complete [Hemaspaandra *et al.*, 2005].

For two elections, $E = (C, V)$ and $F = (D, U)$, such that $|C| = |D|$, $V = (v_1, \dots, v_n)$, and $U = (u_1, \dots, u_n)$, by $d_{\text{swap}}(E, F)$ we denote their *isomorphic swap distance* [Faliszewski *et al.*, 2019], that is, the (minimal) sum of swap distances between the votes in both elections, given by optimal correspondences between their candidates and their voters. Formally:

$$d_{\text{swap}}(E, F) = \min_{\sigma \in \Pi([n], [n])} \min_{\pi \in \Pi(C, D)} \sum_{i=1}^n \text{swap}(\pi(v_i), u_{\sigma(i)}),$$

where by $\pi(v_i)$ we denote vote v_i with every candidate $c \in C$ replaced by candidate $\pi(c)$.

2.3 Maps of Elections

A *map of elections* is a collection of elections represented on a 2D plane as points, so that the Euclidean distances between the points reflect the similarity between the elections (the closer two points are, the more similar should their elections be). Maps of elections were introduced by Szufa *et al.* [2020] (together with an open-source Python library *mapel*, which we use and build on) and Boehmer *et al.* [2021], who used the distance based on position matrices of elections as a measure of similarity. We use the isomorphic swap distance instead. Indeed, Szufa *et al.* [2020] and Boehmer *et al.* [2021] admitted that isomorphic swap distance would be more accurate but avoided it because it is hard to compute (Boehmer *et al.* [2022] analyzed the consequences of using various distances). We are able to use the swap distance because we

focus on small candidate sets. To present a set of elections as a map, we compute the distance between each two elections and then run the multidimensional scaling algorithm (MDS) to find an embedding of points on a plane that reflects the computed distances. For an example of a map, see Fig. 2 at the end of the paper; we describe its elections in Section 5.

2.4 Agreement and Other Election Indices

Election index is a function that given an election outputs a real number. The next index is among the most studied ones and captures voter agreement.

Definition 1. *The agreement index of an election $E = (C, V)$ is: $A(E) = (\sum_{\{a,b\} \subseteq C} |p_E(a,b) - p_E(b,a)|) / \binom{|C|}{2}$.*

The agreement index takes values between 0 and 1, where 0 means perfect disagreement and 1 means perfect agreement. Indeed, we have $A(\text{ID}) = 1$ and $A(\text{UN}) = A(\text{AN}) = 0$.

There is also a number of other election indices in the literature. Somewhat disappointingly, they mostly fall into one or more of the following categories: (1) They are generalizations of the agreement index (or its linear transformation) [Alcalde-Unzu and Vorsatz, 2016; Can *et al.*, 2017]; (2) They are highly correlated with the agreement index (at least on our datasets) [Hashemi and Endriss, 2014; Karpov, 2017; Alcantud *et al.*, 2013]; (3) Their values come from a small set, limiting their expressiveness and robustness [Bosch, 2006; Hashemi and Endriss, 2014].

3 Diversity and Polarization Indices

In this section, we introduce our two new election indices, designed to measure the levels of diversity and polarization in elections. Both of them are defined on top of a generalization of the KEMENY RANKING problem (note that this generalization is quite different from that studied by Arrighi *et al.* [2021] under a related name).

Definition 2. *k -Kemeny rankings of election $E = (C, V)$ are the elements of a set $\Lambda = \{\lambda_1, \dots, \lambda_k\}$ of k linear orders over C that minimize $\sum_{v \in V} \min_{i \in [k]} \text{swap}(v, \lambda_i)$.*

The k -Kemeny distance, $\kappa_k(E)$, is equal to this minimum.

We can think of finding k -Kemeny rankings as finding an optimal split of votes into k groups and minimizing the sum of each group's distance to its Kemeny ranking. Hence, 1-Kemeny distance is simply the distance of the voters from the (standard) Kemeny ranking. We will later argue that $\kappa_1(E)$ is closely related to the agreement index.

We want our diversity index to be high for UN, but small for AN and ID. For identity, 1-Kemeny distance is equal to zero, but for both UN and AN, 1-Kemeny distance is equal to $|V| \cdot \binom{|C|}{2} / 2$, which is the maximal possible value (as shown, for example, by Boehmer *et al.* [2022]). However, for $k \geq 2$ we observe a sharp difference between k -Kemeny distances in these two elections. For AN, we get distance zero (it suffices to use the two opposing votes as the k -Kemeny rankings), and for UN we get non-negligible positive distances (as long as k is smaller than the number of possible votes). Motivated by this, we define the *diversity index* as a normalized sum of all k -Kemeny distances.

Definition 3. *The diversity index of an election $E = (C, V)$ is: $D(E) = (\sum_{k \in [|V|]} \kappa_k(E) / k) / (|V| \cdot \binom{|C|}{2})$.*

The sum in the definition is divided by the number of voters and the maximal possible distance $\binom{|C|}{2}$ between two votes. As a result, the values of the index are more consistent across elections with different number of voters and candidates (for example, diversity of AN is always equal to $1/2$). Apart from that, in the sum, each k -Kemeny distance is divided by k . This way, the values for large k have lesser impact on the total value, and it also improves scalability. However, we note that even with this division, diversity of UN seems to grow slightly faster than linearly with the growing number of candidates and there is a significant gap between the value for UN with all $m!$ possible votes and even the most diverse election with significantly smaller number of voters. The currently defined diversity index works well on our datasets (see Section 6), but finding a more robust normalization is desirable (the obvious idea of dividing by the highest possible value of the sum is challenging to implement and does not prevent the vulnerability to changes in the voters count).

To construct the polarization index, we look at AN and take advantage of the sudden drop from the maximal possible value of the 1-Kemeny distance to zero for the 2-Kemeny distance. We view this drop as characteristic for polarized elections because they include two opposing, but coherent, factions. Consequently, we have the following definition (we divide by $|V| \cdot \binom{|C|}{2} / 2$ for normalization; the index takes values between 0, for the lowest polarization, and 1, for the highest).

Definition 4. *The polarization index of an election $E = (C, V)$ is: $P(E) = 2(\kappa_1(E) - \kappa_2(E)) / (|V| \cdot \binom{|C|}{2})$.*

For AN polarization is one, while for ID it is zero. For UN with 8 candidates, it is 0.232. This is intuitive as in UN every vote also has its reverse. However, we have experimentally checked that with a growing number of candidates the polarization of UN seems to approach zero (e.g., it is 0.13, 0.054, and 0.024 for, respectively, 20, 100, and 500 candidates).

We note that there is extensive literature on polarization measures in different settings, such as regarding distributions over continuous intervals [Esteban and Ray, 1994] or regarding opinions on networks [Musco *et al.*, 2018; Huremović and Ozkes, 2022; Tu and Neumann, 2022; Zhu and Zhang, 2022] (we only mention a few example references). On the technical level, this literature is quite different from our setting, but finding meta connections could be very inspiring.

Concluding our discussion of the election indices, we note a connection between the agreement index and the 1-Kemeny distance. Let μ be the *majority relation* of an election $E = (C, V)$, that is, a relation such that for candidates $a, b \in C$, $a \succeq_\mu b$ if and only if $p_E(a,b) \geq p_E(b,a)$. If E does not have a *Condorcet cycle*, that is, there is no cycle within μ , then μ is identical to the Kemeny ranking. As noted by Can *et al.* [2015], the agreement index can be expressed as a linear transformation of the sum of the swap distances from all the votes to μ . Hence, if there is no Condorcet cycle, the agreement index is strictly linked to $\kappa_1(E)$ and all three of our indices are related.

4 Computation of k -Kemeny Distance

We define an optimization problem k -KEMENY in which the goal is to find the k -Kemeny distance of a given election (see Definition 2). In a decision variant of k -KEMENY, we check if the k -Kemeny distance is at most a given value. We note that k -KEMENY is NP-hard [Bartholdi *et al.*, 1989], even for $k = 1$ and $n = 4$ [Dwork *et al.*, 2001]. Hence, we seek polynomial-time approximation algorithms.

4.1 Approximation Algorithms

While there is a PTAS for 1-KEMENY [Kenyon-Mathieu and Schudy, 2007], it is not obvious how to approximate even 2-KEMENY. Yet, we observe that k -KEMENY is related to the classic facility location problem k -MEDIAN [Williamson and Shmoys, 2011]. In this problem we are given a set of clients X , a set of potential facility locations F , a natural number k , and a metric d defined over $X \cup F$. The goal is to find a subset $W = \{f_1, f_2, \dots, f_k\}$ of facilities which minimizes the *total connection cost* of the clients, that is, $\sum_{x \in X} \min_{f \in W} d(x, f)$. We see that k -KEMENY is equivalent to k -MEDIAN in which the clients are the votes from the input election, the set of facilities is the set of all possible votes, and the metric is the swap distance. Hence, to approximate k -KEMENY we can use approximation algorithms designed for k -MEDIAN. The issue is that there are $m!$ possible Kemeny rankings and the algorithms for k -MEDIAN run in polynomial time with respect to the number of facilities so they would need exponential time.

We tackle the above issue by reducing the search space from all possible rankings to those appearing in the input. We call this problem k -KEMENY AMONG VOTES and provide the following result.²

Theorem 1. *An α -approximate solution for k -KEMENY AMONG VOTES is a 2α -approximate solution for k -KEMENY.*

This allows us to use the rich literature on approximation algorithms for k -MEDIAN [Williamson and Shmoys, 2011]. For example, using the (currently best) 2.7-approximation algorithms for k -MEDIAN [Byrka *et al.*, 2017; Cohen-Addad *et al.*, 2023; Gowda *et al.*, 2023] we get the following.

Corollary 1. *There is a polynomial-time 5.4-approximation algorithm for k -KEMENY.*

The algorithms of Byrka *et al.* [2017], Cohen-Addad *et al.* [2023] and Gowda *et al.* [2023] are based on a complex procedure for rounding a solution of a linear program, which is difficult to implement. Moreover, there are large constants hidden in the running time. Fortunately, there is a simple local search algorithm for k -MEDIAN which achieves $(3 + 2/p)$ -approximation in time $|F|^p \cdot \text{poly}(|F|, |X|)$, where p is the swap size (as a basic building block, the algorithm uses a swap operation which replaces p centers with p other ones, to locally minimize the connection cost) [Arya *et al.*, 2001].

Corollary 2. *There is a local search $(6 + 4/p)$ -approximation algorithm for k -KEMENY, where p is the swap size.*

²We note that the special case of Theorem 1 for $k = 1$ and $\alpha = 1$ was proved by Endriss and Grandi [2014].

We implemented the local search algorithm for $p = 1$ and used it in our experiments (see Section 6). We note that there is a recent result [Cohen-Addad *et al.*, 2022] which shows that the same local search algorithm actually has an approximation ratio $2.83 + \epsilon$, but at the cost of an enormous swap size (hence also the running time)—for example, for approximation ratio below 3 one needs swap size larger than 10^{10000} .

In our experiments in Section 6, we also use a greedy algorithm, which solves k -KEMENY AMONG VOTES iteratively: Starting from an empty set of rankings, in each iteration, it adds a ranking (from those appearing among the votes) that decreases the k -Kemeny distance most. It is an open question if this algorithm achieves a bounded approximation ratio.

We also point out that using the PTAS for 1-KEMENY, we can obtain an approximation scheme in parameterized time for k -KEMENY (parameterized by the number of voters; note that an exact parameterized algorithm is unlikely as 1-KEMENY is already NP-hard for four voters [Dwork *et al.*, 2001]). The idea is to guess the partition of the voters and solve 1-KEMENY for each group.

Theorem 2. *For every $\epsilon > 0$, there is a $(1 + \epsilon)$ -approximation algorithm for k -KEMENY which runs in time FPT w.r.t. n .*

All algorithms in this section, besides solving the decision problem, also output the sought k -Kemeny rankings.

4.2 Hardness of k -Kemeny Among Votes

The reader may wonder why we use k -MEDIAN algorithms instead of solving k -KEMENY AMONG VOTES directly. Unfortunately, even this restricted variant is intractable.

Theorem 3. *k -KEMENY AMONG VOTES is NP-complete and W[2]-hard when parameterized by k .*

Proof (sketch). We give a reduction from the MAX K -COVER problem (which is equivalent to the well-known Approval Chamberlin-Courant voting rule [Procaccia *et al.*, 2008]). In MAX K -COVER we are given a set of elements $X = \{x_1, x_2, \dots, x_N\}$, a family $\mathcal{S} = \{S_1, S_2, \dots, S_M\}$ of nonempty, distinct subsets of X , and positive integers $K \leq M$ and T . The goal is to find K subsets from \mathcal{S} which together cover at least T elements from X .

We take an instance (X, \mathcal{S}, K, T) of MAX K -COVER and construct an instance of k -KEMENY AMONG VOTES as follows. We create three *pivot-candidates* p_1, p_2 , and p_3 . For every set $S \in \mathcal{S}$, we create two *set-candidates* c_S and d_S obtaining, in total, $m = 2M + 3$ candidates. Next, we create the votes, each with the following *vote structure*:

$$\{p_1, p_2, p_3\} \succ \{c_{S_1}, d_{S_1}\} \succ \{c_{S_2}, d_{S_2}\} \succ \dots \succ \{c_{S_M}, d_{S_M}\},$$

where $\{c, d\}$ means that the order of candidates c and d is not specified. Hence, when defining a vote we will only specify the voter's preference on the unspecified pairs of candidates.

For every set $S_j \in \mathcal{S}$, we create $L = N(M + 4)$ *set-voters* v_j (we do not need to distinguish between these copies, hence we call any of them v_j) with the following specification over the vote structure:

$$p_1 \succ_{v_j} p_2 \succ_{v_j} p_3; \quad d_{S_j} \succ_{v_j} c_{S_j}; \quad c_S \succ_{v_j} d_S, \text{ for } S \neq S_j.$$

For each two set-voters u and v , $\text{swap}(u, v) \in \{0, 2\}$ and it equals 0 if and only if u and v come from the same set.

For every element $x_i \in X$, we create an *element-voter* e_i with the following specification over the vote structure:

$$p_3 \succ_{e_i} p_2 \succ_{e_i} p_1; \quad d_S \succ_{e_i} c_S, \text{ for } e_i \in S; \\ c_S \succ_{e_i} d_S, \text{ for } e_i \notin S.$$

Note that for each element-voter e_i and set voter v_j , $\text{swap}(e_i, v_j) \geq 3$. In total we have $n = N(M^2 + 4M + 1)$ voters. We define $k = K$ and we ask if the k -Kemeny distance in k -KEMENY AMONG VOTES is at most $D = 2L(M - K) + \sum_{j \in [M]} |S_j| + 4N - 2T$.

One direction follows by taking k set-voters corresponding to a solution for MAX K -COVER. The other one follows by observing that a solution to k -KEMENY AMONG VOTES may contain only set-voters (because there are $N(M + 4)$ copies of each) and, hence, we can derive a corresponding solution for MAX K -COVER.

In order to achieve the theorem statement we notice that MAX K -COVER is W[2]-hard w.r.t. K [Cygan *et al.*, 2015],³ $k = K$, and the reduction runs in polynomial time. \square

In the full version of the paper, we show that using the same reduction as in the proof of Theorem 3, we can provide more fine-grained hardness results.

5 Statistical Cultures of Our Dataset

Before we move on to our main experiments, we describe and analyze our dataset. It consists of 292 elections with 8 candidates and 96 voters each, generated from several *statistical cultures*, that is, models of generating random elections. For example, under *impartial culture (IC)* each vote is drawn uniformly at random from all possible votes (thus, it closely resembles UN). We present our dataset as a map of elections in Fig. 2. In the full version of the paper we consider two additional datasets.

Below, we discuss each statistical culture used in our dataset and build an intuition on how our indices should evaluate elections generated from them. To this end, we form a new type of a map, which we call a *map of preferences*, where we look at relations between votes within a single election. In other words, a map of elections gives a bird’s eye view of the space of elections, and a map of preferences is a microscope view of a single election.

5.1 Maps of Preferences

To generate a map of preferences for a given election, we first compute the (standard) swap distance between each pair of its votes. Then, based on these distances, we create a map in the same way as for maps of elections (that is, we use the multidimensional scaling algorithm). We obtain a collection of points in 2D, where each point corresponds to a vote in the election, and Euclidean distances between the points resemble the swap distances between the votes they represent.

For each model, we generated a *single* election and created its map of preferences. The results are shown in Fig. 1. The elections have 1000 voters instead of 96, so that the pictures look similar each time we draw an election from the model.

³Actually, the result comes from W[2]-hardness of the SET COVER problem and a folklore reduction to MAX K -COVER by setting $T = N$.

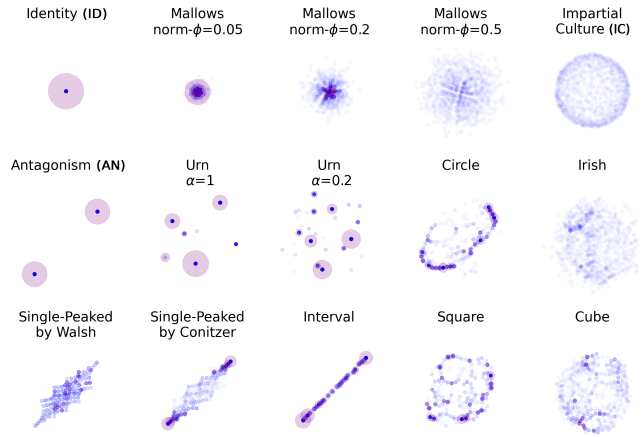


Figure 1: Maps of Preferences (8 candidates, 1000 voters). If there are more than 10 copies of the same vote, we add a purple disc with a radius proportional to the number of voters.

5.2 Model Definitions and Analysis

ID, AN, and IC

We first consider ID, AN, and IC elections (which, for the time being, covers for UN). ID and AN are shown as the first two entries of the first two rows in Fig. 1. The former, with 1000 copies of the same vote, presented as a single point with a large purple disc, embodies perfect agreement. The latter, with 500 votes of one type and 500 its reverses, represents a very polarized society, which is well captured by the two faraway points with large discs on its map. Under IC, whose map is the last one in the first row, we see no clear structure except that, of course, there are many pairs of votes at high swap distance (they form the higher-density rim). Yet, for each such pair there are also many votes in between. Hence, it is close to being perfectly diverse.

We do not present UN in our maps because it requires at least $m!$ votes. Indeed, from now on, we will consider an approximate variant of UN, i.e., UN*, generated by sampling votes from the scaled position matrix of UN.

Mallows Model

The Mallows model is parameterized by the central vote u and the dispersion parameter $\phi \in [0, 1]$. Votes are generated independently and the probability of generating a vote v is proportional to $\phi^{\text{swap}(u,v)}$. Instead of using the parameter ϕ directly, we follow Boehmer *et al.* [2021] and use its normalized variant, $\text{norm-}\phi \in [0, 1]$, which is internally converted to ϕ (see their work for details; with 8 candidates the conversion is nearly linear). For $\text{norm-}\phi = 1$, the Mallows model is equivalent to IC, for $\text{norm-}\phi = 0$ it is equivalent to ID, and for values in between we get a smooth transition between these extremes (or, between agreement and diversity, to use our high-level notions). We see this in the first row of Fig. 1.

Urn Model

In the Pólya-Eggenberger urn model [Berg, 1985; McCabe-Dansted and Slinko, 2006], we have a parameter of contagion $\alpha \in [0, \infty)$. We start with an urn containing one copy of each possible vote and we repeat the following process n times:

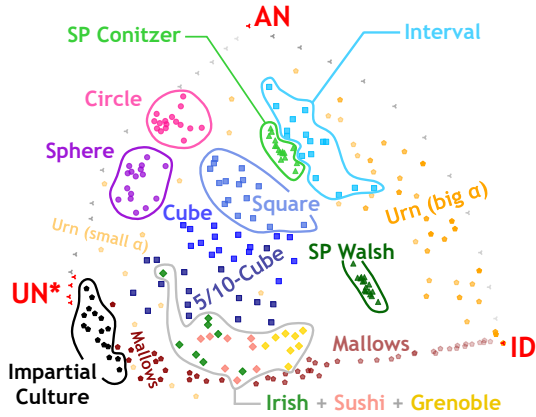


Figure 2: A map of elections in our dataset obtained using isomorphic swap distance and MDS.

We draw a vote from the urn, its copy is included in the election, and the vote, together with $\alpha \cdot m!$ copies, is returned to the urn. For $\alpha = 0$ the model is equivalent to IC. The larger is the α value, the stronger is the correlation between the votes.

In Fig. 1, urn elections (shown in the middle of the second row) consist of very few distinct votes. For example, for $\alpha = 1$ we only have seven votes, thus this election’s map looks similarly to that for AN—few points with discs. Such elections, with several popular views but without a spectrum of opinions in between, are known as *fragmented* [Dynes and Tierney, 1994]. Hence, we expect their diversity to be small. As α decreases, urn elections become less fragmented.

We upper-bound the expected number of different votes in an urn election with m candidates, n voters (where n is significantly smaller than $m!$), and parameter α by $\sum_{i=1}^n 1/(1+(i-1)\alpha)$ (the first vote is always unique, the second one is drawn from the original $m!$ votes from the urn with probability $1/(1+\alpha)$, and so on; if we draw one of the original votes from the urn it still might be the same as one of the previous ones, but this happens with a small probability when n is significantly smaller than $m!$). For example, for $n = 1000$ and α equal to 1, our formula gives 7.48. In the literature, authors often use $\alpha = 1$ [Erdélyi *et al.*, 2015; Keller *et al.*, 2019; Walsh, 2011], sometimes explicitly noting the strong correlations and modifying the model [Erdélyi *et al.*, 2015]. However, smaller values of α are also used [Skowron *et al.*, 2015; McCabe-Dansted and Slinko, 2006]. Since $\alpha = 1$ gives very particular elections, it should be used consciously.

Single-Peaked Elections

Single-peaked elections [Black, 1958] capture scenarios where voters have a spectrum of opinions between two extremes (like choosing a preferred temperature in a room).

Definition 5 (Black [1958]). *Let C be a set of candidates and let \triangleright be an order over C , called the societal axis. A vote is single-peaked with respect to \triangleright if for each $t \in [|C|]$, its top t candidates form an interval w.r.t. \triangleright . An election is single-peaked (w.r.t. \triangleright) if its votes are.*

We use the *Walsh* [Walsh, 2015] and the *Conitzer (random peak)* models [Conitzer, 2009] of generating single-peaked

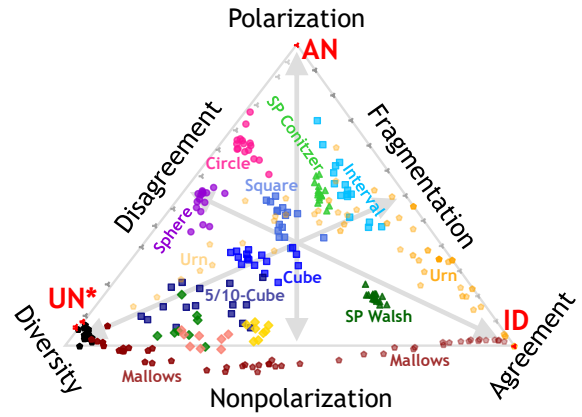


Figure 3: An affine transformation of a plot where x/y coordinates of the elections are their agreement/diversity indices.

elections. In the former, we fix the societal axis and choose votes single-peaked with respect to it uniformly at random (so we can look at it as IC over the single-peaked domain). In the Conitzer model we also first fix the axis, and then generate each vote as follows: We choose the top-ranked candidate uniformly at random and fill-in the following positions by choosing either the candidate directly to the left or directly to the right of the already selected ones on the axis, with probability $1/2$ (at some point we run out of the candidates on one side and then only use the other one).

In Fig. 1, Conitzer and Walsh elections are similar, but the former one has more votes at large swap distance. Indeed, under the Conitzer model, we generate a vote equal to the axis (or its reverse) with probability $2/m$, which for $m = 8$ is 25%. Under the Walsh model, this happens with probability 1.5% (it is known there are 2^{m-1} different single-peaked votes and Walsh model chooses each of them with equal probability). Hence, our Conitzer elections are more polarized (see the purple discs at the farthest points) than the Walsh ones, and Walsh ones appear to be more in agreement (in other words, the map for the Conitzer election is more similar to that for AN, and the map for Walsh election is more similar to ID).

Euclidean Models

In d -dimensional Euclidean elections every candidate and every voter is a point in \mathbb{R}^d , and a voter prefers candidate a to candidate b if his or her point is closer to that of a than to that of b . To generate such elections, we sample the candidate and voter points as follows: (a) In the d -Cube model, we sample the points uniformly at random from a d -dimensional hypercube $[0, 1]^d$, and (b) in the Circle and Sphere models we sample them uniformly at random from a circle (embedded in 2D space) and a sphere (embedded in 3D space). We refer to the 1-Cube, 2-Cube, and 3-Cube models as, respectively, the Interval, Square, and Cube models. In Fig. 1, we see that as the dimension increases, the elections become more similar to the IC one. The Interval election is very similar to those of Conitzer and Walsh, because 1-Euclidean elections are single-peaked. It is also worth noting that the Circle election is quite polarized (we see an increased density of votes on two opposite sides of its map).

Irish and Other Elections Based on Real-Life Data

We also consider elections generated based on real-life data from a 2002 political election in Dublin [Mattei and Walsh, 2013]. We treat the full Irish data as a distribution and sample votes from it as from a statistical culture. The Irish election in Fig. 1 is, in some sense, between the Cube and Mallows ones for $\text{norm-}\phi = 0.5$. Intuitively, we would say that it is quite diverse. In the dataset, we also include Sushi and Grenoble elections, similarly generated using different real-life data [Mattei and Walsh, 2013].

6 Final Experiments and Conclusion

In this section we present the results of computing the agreement, diversity, and polarization indices on our dataset.

6.1 Computing the Indices in Practice

First, we compared three ways of computing k -Kemeny distances: the greedy approach, the local search with swap size equal to 1, and a combined heuristic where we first calculate the greedy solution and then try to improve it using the local search. We ran all three algorithms for all $k \in [96]$ and for every election in our dataset. The conclusion is that the local search and the combined heuristic gave very similar outcomes and both outperformed the greedy approach. Hence, in further computations, we used the former two algorithm and took the smaller of their outputs.

6.2 Understanding the Map via Agreement, Diversity, and Polarization

Using the $\kappa_k(E)$ values computed in the preceding experiment, we calculated diversity and polarization indices of all the elections from our dataset, along with their agreement indices (which are straightforward to compute). We illustrate the results in several ways.

First, we consider Fig. 4. In the left plot, each election from our dataset is represented as a dot whose x/y coordinates are the values of the diversity index and the distance from UN*, and whose color corresponds to the statistical culture from which it comes (it is the same as in Fig. 2). The plot on the right is analogous, except that it regards polarization and distance from AN. An analogous plot for agreement and distance from ID is almost perfectly linear. The Pearson correlation coefficient between each of the three indices and the distance from the respective compass election is below -0.9 , which means that the correlation is very strong. This is our first indication that the locations on the map of elections, in particular, the one from Fig. 2, can be understood in terms of agreement, diversity, and polarization.

Next, for all three pairs of our indices we plotted our dataset in such a way that each election’s x/y coordinates are the values of the respective indices (these plots can be found in the full version of the paper). We observed that each of these plots resembles the original map from Fig. 2. Hence, for the sake of clearer comparison, we took the plot for agreement and diversity indices and, by an affine transformation, converted it to a roughly equilateral triangle spanned between ID, AN, and UN*. Fig. 3 presents the result of this operation.




Figure 4: Correlation between our indices and the distance from the respective compass election.

The similarity between Figs. 2 and 3 is striking as most elections can be found in analogous locations. Even the positions of the outliers in the groups are approximately preserved. Yet, there are also differences. For example, in Fig. 3 elections from most of the statistical cultures are closer to each other than in Fig. 2. Nonetheless, the similarity between these two figures is our second argument for understanding the map in terms of agreement, diversity, and polarization. Specifically, the closer an election is to ID, AN, or UN*, the more agreement, polarization, or diversity it exhibits.

6.3 Validation Against Intuition

Finally, let us check our intuitions from Section 5 against the actually computed values of the indices, as presented on the plot from Fig. 3. We make the following observations:

1. We see that Mallows elections indeed progress from ID (for which we use $\text{norm-}\phi = 0$) to IC (for which we use $\text{norm-}\phi = 1$), with intermediate values of $\text{norm-}\phi$ in between. The model indeed generates elections on the agreement-diversity spectrum.
2. Elections generated using the urn model with large value of α appear on the agreement-polarization line. Indeed, for very large values of α nearly all the votes are identical, but for smaller values we see polarization effects. Finally, as the values of α go toward 0, the votes become more and more diverse.
3. Walsh elections are closer to agreement (ID) and Conitzer elections are closer to polarization (AN).
4. High-dimensional Cube elections have fairly high diversity. Circle and Sphere elections are between diversity and polarization.
5. Irish elections are between Mallows and high-dimensional Cube elections.

All in all, this confirms our intuitions and expectations.

7 Summary

The starting point of our work was an observation that some of the measures of diversity and polarization used in computational social choice literature should, rather, be seen as measures of disagreement. We have proposed two new measures and we have argued that they do capture diversity and polarization. On the negative side, our measures are computationally intractable. Hence, finding a measure that would be easy to compute but that would maintain the intuitive appeal of our ones is an interesting research topic.

Acknowledgements

Krzysztof Sornat was supported by the SNSF Grant 200021_200731/1. 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).



References

- [Alcalde-Unzu and Vorsatz, 2013] J. Alcalde-Unzu and M. Vorsatz. Measuring the cohesiveness of preferences: An axiomatic analysis. *Social Choice and Welfare*, 41(4):965–988, 2013.
- [Alcalde-Unzu and Vorsatz, 2016] J. Alcalde-Unzu and M. Vorsatz. Do we agree? Measuring the cohesiveness of preferences. *Theory and Decision*, 80(2):313–339, 2016.
- [Alcantud *et al.*, 2013] J. C. R. Alcantud, R. de Andrés Calle, and J. M. Cascón. A unifying model to measure consensus solutions in a society. *Mathematical and Computer Modelling*, 57(7-8):1876–1883, 2013.
- [Arrighi *et al.*, 2021] E. Arrighi, H. Fernau, D. Lokshtanov, M. de Oliveira Oliveira, and P. Wolf. Diversity in Kemeny rank aggregation: A parameterized approach. In *Proceedings of IJCAI-2021*, pages 10–16, 2021.
- [Arya *et al.*, 2001] V. Arya, N. Garg, R. Khandekar, A. Meyerson, K. Munagala, and V. Pandit. Local search heuristic for k -median and facility location problems. In *Proceedings of STOC-2001*, pages 21–29, 2001.
- [Bartholdi *et al.*, 1989] J. Bartholdi, III, C. Tovey, and M. Trick. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare*, 6(2):157–165, 1989.
- [Berg, 1985] S. Berg. Paradox of voting under an urn model: The effect of homogeneity. *Public Choice*, 47(2):377–387, 1985.
- [Black, 1958] D. Black. *The Theory of Committees and Elections*. Cambridge University Press, 1958.
- [Boehmer *et al.*, 2021] N. Boehmer, R. Bredereck, P. Faliszewski, R. Niedermeier, and S. Szufa. Putting a compass on the map of elections. In *Proceedings of IJCAI-2021*, pages 59–65, 2021.
- [Boehmer *et al.*, 2022] N. Boehmer, P. Faliszewski, R. Niedermeier, S. Szufa, and T. Wąs. Understanding distance measures among elections. In *Proceedings of IJCAI-2022*, pages 102–108, 2022.
- [Bosch, 2006] R. Bosch. *Characterizations on Voting Rules and Consensus Measures*. PhD thesis, Tilburg University, 2006.
- [Byrka *et al.*, 2017] J. Byrka, T. W. Pensyl, B. Rybicki, A. Srinivasan, and K. Trinh. An improved approximation for k -median and positive correlation in budgeted optimization. *ACM Transactions on Algorithms*, 13(2):23:1–23:31, 2017.
- [Can *et al.*, 2015] B. Can, A. I. Ozkes, and T. Storcken. Measuring polarization in preferences. *Mathematical Social Sciences*, 78:76–79, 2015.
- [Can *et al.*, 2017] B. Can, A. Ozkes, and T. Storcken. Generalized measures of polarization in preferences. Technical report, HAL, 2017.
- [Chamberlin and Courant, 1983] B. Chamberlin and P. Courant. Representative deliberations and representative decisions: Proportional representation and the Borda rule. *American Political Science Review*, 77(3):718–733, 1983.
- [Cohen-Addad *et al.*, 2022] V. Cohen-Addad, A. Gupta, L. Hu, H. Oh, and D. Saulpic. An improved local search algorithm for k -median. In *Proceedings of SODA-2022*, pages 1556–1612, 2022.
- [Cohen-Addad *et al.*, 2023] V. Cohen-Addad, F. Grandoni, E. Lee, and C. Schwiegelshohn. Breaching the 2 LMP approximation barrier for facility location with applications to k -median. In *Proceedings of SODA-2023*, pages 940–986, 2023.
- [Conitzer, 2009] V. Conitzer. Eliciting single-peaked preferences using comparison queries. *Journal of Artificial Intelligence Research*, 35:161–191, 2009.
- [Cygan *et al.*, 2015] M. Cygan, F. V. Fomin, Ł. Kowalik, D. Lokshtanov, D. Marx, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. *Parameterized Algorithms*. Springer, 2015.
- [Dwork *et al.*, 2001] C. Dwork, R. Kumar, M. Naor, and D. Sivakumar. Rank aggregation methods for the Web. In *Proceedings of WWW-2001*, pages 613–622, 2001.
- [Dynes and Tierney, 1994] R. Dynes and K. Tierney, editors. *Disasters, Collective Behavior, and Social Organization*. University of Delaware Press, 1994.
- [Endriss and Grandi, 2014] U. Endriss and U. Grandi. Binary aggregation by selection of the most representative voters. In *Proceedings of AAI-2014*, pages 668–674, 2014.
- [Erdélyi *et al.*, 2015] G. Erdélyi, M. Fellows, J. Rothe, and L. Schend. Control complexity in Bucklin and fallback voting: An experimental analysis. *Journal of Computer and System Sciences*, 81(4):661–670, 2015.
- [Esteban and Ray, 1994] J.-M. Esteban and D. Ray. On the measurement of polarization. *Econometrica*, 62(4):819–851, 1994.
- [Faliszewski *et al.*, 2017] P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Multiwinner voting: A new challenge for social choice theory. In U. Endriss, editor, *Trends in Computational Social Choice*. AI Access Foundation, 2017.
- [Faliszewski *et al.*, 2019] P. Faliszewski, P. Skowron, A. Slinko, S. Szufa, and N. Talmon. How similar are two elections? In *Proceedings of AAI-2019*, pages 1909–1916, 2019.

- [Faliszewski *et al.*, 2023] P. Faliszewski, A. Kaczmarczyk, K. Sornat, S. Szufa, and T. Wąs. Diversity, agreement, and polarization in elections. Technical Report arXiv:2305.09780 [cs.GT], arXiv.org, May 2023.
- [Gowda *et al.*, 2023] K. N. Gowda, T. W. Pensyl, A. Srinivasan, and K. Trinh. Improved bi-point rounding algorithms and a golden barrier for k -median. In *Proceedings of SODA-2023*, pages 987–1011, 2023.
- [Hashemi and Endriss, 2014] V. Hashemi and U. Endriss. Measuring diversity of preferences in a group. In *Proceedings of ECAI-2014*, pages 423–428, 2014.
- [Hemaspaandra *et al.*, 2005] E. Hemaspaandra, H. Spakowski, and J. Vogel. The complexity of Kemeny elections. *Theoretical Computer Science*, 349(3):382–391, 2005.
- [Huremović and Ozkes, 2022] K. Huremović and A. I. Ozkes. Polarization in networks: Identification–alienation framework. *Journal of Mathematical Economics*, 102:102732, 2022.
- [Karpov, 2017] A. Karpov. Preference diversity orderings. *Group Decision and Negotiation*, 26(4):753–774, 2017.
- [Keller *et al.*, 2019] O. Keller, A. Hassidim, and N. Hazon. New approximations for coalitional manipulation in scoring rules. *Journal of Artificial Intelligence Research*, 64:109–145, 2019.
- [Kemeny, 1959] J. Kemeny. Mathematics without numbers. *Daedalus*, 88:577–591, 1959.
- [Kenyon-Mathieu and Schudy, 2007] C. Kenyon-Mathieu and W. Schudy. How to rank with few errors. In *Proceedings of STOC-2007*, pages 95–103, 2007.
- [Levin *et al.*, 2021] S. A. Levin, H. V. Milner, and C. Perrings. The dynamics of political polarization. *Proceedings of the National Academy of Sciences*, 118(50):e21116950118, 2021.
- [Mattei and Walsh, 2013] N. Mattei and T. Walsh. Preflib: A library for preferences. In *Proceedings of ADT-2013*, pages 259–270, 2013.
- [McCabe-Dansted and Slinko, 2006] J. McCabe-Dansted and A. Slinko. Exploratory analysis of similarities between social choice rules. *Group Decision and Negotiation*, 15:77–107, 2006.
- [Musco *et al.*, 2018] C. Musco, C. Musco, and C. E. Tsourakakis. Minimizing polarization and disagreement in social networks. In *Proceedings of WWW-2018*, pages 369–378, 2018.
- [Procaccia *et al.*, 2008] A. D. Procaccia, J. S. Rosenschein, and A. Zohar. On the complexity of achieving proportional representation. *Social Choice and Welfare*, 30(3):353–362, 2008.
- [Skowron *et al.*, 2015] P. Skowron, P. Faliszewski, and A. Slinko. Achieving fully proportional representation: Approximability result. *Artificial Intelligence*, 222:67–103, 2015.
- [Szufa *et al.*, 2020] S. Szufa, P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Drawing a map of elections in the space of statistical cultures. In *Proceedings of AAMAS-2020*, pages 1341–1349, 2020.
- [Tu and Neumann, 2022] S. Tu and S. Neumann. A viral marketing-based model for opinion dynamics in online social networks. In *Proceedings of WWW-2022*, pages 1570–1578, 2022.
- [Walsh, 2011] T. Walsh. Where are the hard manipulation problems. *Journal of Artificial Intelligence Research*, 42(1):1–29, 2011.
- [Walsh, 2015] T. Walsh. Generating single peaked votes. Technical Report arXiv:1503.02766 [cs.GT], arXiv.org, March 2015.
- [Williamson and Shmoys, 2011] D. P. Williamson and D. B. Shmoys. *The Design of Approximation Algorithms*. Cambridge University Press, 2011.
- [Zhu and Zhang, 2022] L. Zhu and Z. Zhang. A nearly-linear time algorithm for minimizing risk of conflict in social networks. In *Proceedings of KDD-2022*, pages 2648–2656, 2022.