

Understanding Distance Measures Among Elections

Niclas Boehmer¹, Piotr Faliszewski², Rolf Niedermeier¹, Stanisław Szufa² and Tomasz Was^{2,3}

¹Technische Universität Berlin, Algorithmics and Computational Complexity

²AGH University

³University of Warsaw

niclas.boehmer@tu-berlin.de, faliszew@agh.edu.pl, szufa@agh.edu.pl, t.was@mimuw.edu.pl

Abstract

Motivated by putting empirical work based on (synthetic) election data on a more solid mathematical basis, we analyze six distances among elections, including, e.g., the challenging-to-compute but very precise swap distance and the distance used to form the so-called *map of elections*. Among the six, the latter seems to strike the best balance between its computational complexity and expressiveness.

1 Introduction

We study the properties of several distances (metrics) among elections.¹ We focus on the ordinal model, where each voter ranks the candidates from the most to the least appealing one, and on metrics that are independent of renaming the candidates and voters. Such metrics were introduced by Faliszewski *et al.* [2019], who argued about their usefulness to compare two elections generated from some statistical models, or to evaluate which statistical model is most likely to generate elections similar to given real-life ones. Indeed, in such applications the names of the candidates and voters do not carry any information and should be disregarded. Unfortunately, while the metrics studied by Faliszewski *et al.* [2019] are naturally motivated—for example, one of them extends the widely accepted swap distance—many of them are not only NP-hard to compute (even approximately), but also difficult to compute in practice.

Yet, many of the motivating ideas of Faliszewski *et al.* [2019] were soon implemented in the *maps of elections*, introduced by Szufa *et al.* [2020] and extended by Boehmer *et al.* [2021b]. Briefly put, such a map is a collection of election instances, typically generated from some statistical models (but Boehmer *et al.* [2021b] also used real-life ones from PrefLib [Mattei and Walsh, 2013]), together with their distances. The elections are represented graphically as points in the plane, whose Euclidean distances are as similar to the distances between the respective elections as possible. Since maps of elections typically contain hundreds of elections, instead of using the appealing-but-hard-to-compute extension

¹Generally, we use the word *distance* when we refer to a value of a *metric*, but occasionally, reflecting the literature, we break this rule (e.g., to speak of the earth mover’s distance or the swap distance).

of the swap distance, Szufa *et al.* [2020] introduced and used a much simpler metric. Yet, the maps proved to be quite useful. For example, Szufa *et al.* [2020] used their map to find hard instances for some multiwinner voting rule, Boehmer *et al.* [2021b] and Boehmer and Schaar [2022] obtained insights about different types of real-life elections, and Boehmer *et al.* [2021a] used the maps to study the robustness of elections.

In this work, we take a step back and analyze the properties of several metrics, including those of Faliszewski *et al.* [2019] and Szufa *et al.* [2020]. Our goal is to help putting empirical work with election data on a more solid mathematical basis and, in particular, to understand if basing the election maps on the simpler metric was a good decision, what was its price, and if one should have used other metrics. We view the swap distance as a yardstick against which we measure the other ones. In particular, we consider the following issues:

1. As a basic test, we compare the metrics’ ability to distinguish nonisomorphic elections. Since some metrics act on aggregate representations of elections, they may sometimes fail at this task. We also study the complexity of computing an election with a given representation.
2. We analyze distances between four “compass” elections, which capture four types of (dis)agreement among the voters [Boehmer *et al.*, 2021b]. We find that two of them are the most distant elections under each of our metrics.
3. We compute the correlation between the values provided by the swap distance and the other metrics; we also compare the maps that they produce.
4. We note that the swap distance can be understood in terms of the shortest paths on a certain graph and we analyze to what extent this applies to the other metrics.

Some proofs and arguments are deferred to our full version [Boehmer *et al.*, 2022].

2 Preliminaries

For an integer n , let $[n] := \{1, \dots, n\}$. Let C be set of candidates. By $\mathcal{L}(C)$ we denote the set of all total orders over C , referred to either as preference orders or votes, depending on the context.

Elections. An election $E = (C, V)$ consists of a candidate set $C = \{c_1, \dots, c_m\}$ and a preference profile $V = (v_1, \dots, v_n)$, where each v_i is a preference order from $\mathcal{L}(C)$.

For example, a preference order $v_i: a \succ b \succ c$ indicates that the i -th voter ranks candidate a highest, followed by candidates b and c . Given a vote v and a candidate c , we write $\text{pos}_v(c)$ to denote the position of c in v (the top-ranked candidate has position 1, the next one has position 2, and so on).

Example 1. Consider elections $E = (C, V)$ and $E' = (C', V')$, where $C = \{a, b, c\}$, $C' = \{x, y, z\}$, $V = (v_1, v_2, v_3)$, $V' = (v'_1, v'_2, v'_3)$, and the votes are:

$$\begin{aligned} v_1: a \succ b \succ c, & \quad v_2: b \succ c \succ a, & \quad v_3: b \succ a \succ c, \\ v'_1: x \succ y \succ z, & \quad v'_2: x \succ y \succ z, & \quad v'_3: y \succ x \succ z. \end{aligned}$$

Aggregate Representations. Let $E = (C, V)$ be an election. We use the following aggregate representations of E :

1. For each two candidates $c, d \in C$, $\mathcal{M}_E(c, d)$ is the number of voters that prefer c to d in election E . We call it the weighted majority relation and represent it as an $m \times m$ matrix where rows and columns correspond to the candidates (the diagonal is undefined). A relative weighted majority relation is a weighted majority relation from whose entries we subtract $n/2$ (let n be even here).
2. For a candidate $c \in C$ and a position $i \in [m]$, $\mathcal{P}_E(c, i)$ is the number of voters from E that rank c on position i ; $\mathcal{P}_E(c) = (\mathcal{P}_E(c, 1), \dots, \mathcal{P}_E(c, m))$ is a (column) position vector of c . We view \mathcal{P}_E as a matrix with columns $\mathcal{P}_E(c_1), \dots, \mathcal{P}_E(c_m)$ and call it a position matrix.
3. For a candidate $c \in C$, $\mathcal{B}_E(c) := \sum_{i=1}^n (m - \text{pos}_{v_i}(c))$ is the Borda score of c in E . Then, \mathcal{B}_E is the Borda score vector, whose entries correspond to the candidates.

Example 2. Below we provide \mathcal{M}_E , \mathcal{P}_E , and \mathcal{B}_E , respectively, for election E from Example 1:

$$\begin{array}{c} \begin{array}{ccc} a & b & c \\ a & - & 1 & 2 \\ b & 2 & - & 3 \\ c & 1 & 0 & - \end{array} \end{array}, \quad \begin{array}{c} \begin{array}{ccc} a & b & c \\ 1 & 2 & 0 \\ 1 & 1 & 1 \\ 1 & 0 & 2 \end{array} \end{array}, \quad \begin{array}{c} \begin{array}{ccc} a & b & c \\ 3 & 5 & 1 \end{array} \end{array}.$$

(Pseudo)Metrics. Given a set X , a function $d: X \times X \rightarrow \mathbb{R}$ is a pseudometric over X if for each three elements $a, b, c \in X$ it holds that (i) $d(a, b) = d(b, a) \geq 0$, (ii) $d(a, a) = 0$, and (iii) $d(a, c) \leq d(a, b) + d(b, c)$. For brevity, we will refer to our pseudometrics as metrics (formally, a metric should assume value 0 only if both its arguments are identical).

Metrics Among Vectors. Let $x = (x_1, \dots, x_n)$ and $y = (y_1, \dots, y_n)$ be two real-valued vectors. Then, $\ell_1(x, y) := |x_1 - y_1| + \dots + |x_n - y_n|$ is the ℓ_1 -metric between x and y . Given a real-valued vector $z = (z_1, \dots, z_n)$, we write \hat{z} to denote its prefix-sum variant, i.e., an n -dimensional vector such that for each $i \in [n]$, its i -th entry is $\hat{z}_i = z_1 + z_2 + \dots + z_i$. If the entries of x and y sum up to the same value and contain only nonnegative entries, then their *earth mover's distance* is defined as [Rubner *et al.*, 2000]:

$$\text{emd}(x, y) := \ell_1(\hat{x}, \hat{y}).$$

Alternatively, $\text{emd}(x, y)$ is defined as the minimum total cost of a sequence of operations that transform vector x into vector y , where each operation is of the form “subtract value α from some x_i (where, at the time of performing the operation, we have $x_i \geq \alpha$) and add α to x_j ” and has cost $\alpha \cdot |j - i|$. Both definitions are equivalent [Rubner *et al.*, 2000].

Bijections. Given two equal-sized sets X and Y , let $\Pi(X, Y)$ denote the set of all one-to-one mappings from X to Y . For a positive integer n , let S_n be the set of all permutations of $[n]$ (i.e., $S_n = \Pi([n], [n])$). Given two equal-sized candidate sets C and D , a preference order $v \in \mathcal{L}(C)$, and a function $\sigma \in \Pi(C, D)$, we write $\sigma(v)$ to denote the preference order obtained from v by replacing each candidate $c \in C$ with the candidate $\sigma(c) \in D$. Given a preference profile $V = (v_1, \dots, v_n) \in (\mathcal{L}(C))^n$, by $\sigma(V)$ we mean $(\sigma(v_1), \dots, \sigma(v_n))$. We use analogous notation for other objects defined over candidate sets. For example, for an election $E = (C, V)$, we write $\sigma(\mathcal{M}_E)$ to denote the weighted majority relation of the election $(\sigma(C), \sigma(V)) = (D, \sigma(V))$.

3 Metrics Among Elections

In this section, we define the six metrics among elections that we study. All but the last one already appeared in the literature. We only consider distances between elections with the same numbers of candidates and the same numbers of voters.

Swap and Discrete Isomorphic Metrics

Let C be a candidate set and let $u, v \in \mathcal{L}(C)$ be two votes. Their discrete distance, $d_{\text{disc}}(u, v)$, is 0 if they are identical and it is 1 otherwise. Their swap distance, $d_{\text{swap}}(u, v)$, is the number of inversions between u and v , i.e., the number of pairs of candidates $c, d \in C$ such that u and v rank these candidates in opposite order.

Let d be either d_{disc} or d_{swap} . For two elections, $E = (C, V)$ and $E' = (C', V')$, where $|C| = |C'|$, $V = (v_1, \dots, v_n)$, and $V' = (v'_1, \dots, v'_n)$, Faliszewski *et al.* [2019] extended d to elections as follows:

$$d(E, E') := \min_{\sigma \in \Pi(C, C')} \min_{\rho \in S_n} \sum_{i=1}^n d(\sigma(v_i), v'_{\rho(i)}).$$

In other words, under the extended distance we match the candidates and the votes of the two input elections so that the sum of the distances between the matched votes is minimal.

Example 3. The discrete distance between elections E and E' from Example 1 is one, as witnessed by the candidate matching $\sigma(a) = x$, $\sigma(b) = y$, and $\sigma(c) = z$. For the same matching, the swap distance between E and E' is two.

We refer to the extensions of d_{disc} and d_{swap} as the discrete and swap *isomorphic metrics*. This stems from the fact that two elections are isomorphic (i.e., can be made identical by renaming candidates and reordering the votes) if and only if their discrete distance (equivalently, swap distance) is zero.

Faliszewski *et al.* [2019] have shown that while computing the discrete isomorphic distance can be done in polynomial time, the same task for the swap distance is NP-hard (in essence, this follows from the hardness proofs for the Kemeny voting rule [Bartholdi *et al.*, 1989; Dwork *et al.*, 2001]).

Positionwise and Pairwise Metrics

To circumvent the hardness of computing the isomorphic swap distance, Szufa *et al.* [2020] introduced two other metrics, based on analyzing aggregate representations of elections.

Let $E = (C, V)$ and $E' = (C', V')$ be two elections such that $|C| = |C'|$ and $|V| = |V'|$. The EMD-positionwise

distance between E and E' is:

$$d_{\text{pos}}^{\text{emd}}(E, E') := \min_{\sigma \in \Pi(C, C')} \sum_{c \in C} \text{emd}(\mathcal{P}_E(c), \mathcal{P}_{E'}(\sigma(c))).$$

Note that we view each candidate as her or his position vector and we seek a candidate matching that minimizes the earth mover's distances between the vectors of matched candidates.

Example 4. Consider the same elections and the same matching σ as in Example 3. We have that $\mathcal{P}_E(a) = (1, 1, 1)$, whereas $\mathcal{P}_{E'}(\sigma(a)) = (2, 1, 0)$ and their earth mover's distance is 2. Altogether, $d_{\text{pos}}^{\text{emd}}(E, E') = 2 + 1 + 1 = 4$.

Szufa *et al.* [2020] based their metric on EMD because they felt it was intuitively appropriate. We aim to verify this intuition and, thus, we also consider the ℓ_1 -positionwise metric, which uses the ℓ_1 distance instead of EMD.

Szufa *et al.* [2020] also introduced the pairwise metric, which works on top of the weighted majority relation:

$$d_{\text{pair}}(E, E') := \min_{\sigma \in \Pi(C, C')} \sum_{\substack{a, b \in C \\ a \neq b}} |\mathcal{M}_E(a, b) - \mathcal{M}_{E'}(\sigma(a), \sigma(b))|.$$

Example 5. Take the elections from Example 1 and matching $\sigma'(a) = y, \sigma'(b) = x$, and $\sigma'(c) = z$; \mathcal{M}_E and $\sigma'(\mathcal{M}_{E'})$ are:

$$\begin{matrix} & \begin{matrix} a & b & c \end{matrix} \\ \begin{matrix} a \\ b \\ c \end{matrix} & \begin{bmatrix} - & 1 & 2 \\ 2 & - & 3 \\ 1 & 0 & - \end{bmatrix}, \end{matrix} \quad \begin{matrix} & \begin{matrix} \sigma'(a) & \sigma'(b) & \sigma'(c) \end{matrix} \\ \begin{matrix} \sigma'(a) \\ \sigma'(b) \\ \sigma'(c) \end{matrix} & \begin{bmatrix} - & 1 & 3 \\ 2 & - & 3 \\ 0 & 0 & - \end{bmatrix} \end{matrix}$$

and we see that the pairwise distance of our elections (for this matching) is $0 + 1 + 0 + 0 + 1 + 0 = 2$.

The positionwise distances can be computed in polynomial time [Szufa *et al.*, 2020], whereas the pairwise distance is NP-hard to compute [Grohe *et al.*, 2018; Szufa *et al.*, 2020].

Bordawise Metric

We introduce one new metric, similar in spirit to the positionwise and pairwise ones, but defined on top of the election's Borda score vectors. Given two equal-sized elections E and E' , their Bordawise distance is defined as:

$$d_{\text{Borda}}(E, E') := \text{emd}(\text{sort}(\mathcal{B}_E), \text{sort}(\mathcal{B}_{E'})),$$

where for a vector x , $\text{sort}(x)$ means a vector obtained from x by sorting it in nonincreasing order. The Bordawise metric is defined to be as simple as possible, while trying to still be meaningful. For example, sorting the score vectors ensures that two isomorphic elections are at distance zero and removes the use of an explicit matching between the candidates.

Example 6. The distance between the elections from Example 1 is $\text{emd}((5, 3, 1), (5, 4, 0)) = 1$.

4 Aggregate Representations

First, we discuss how the aggregate representations that underlie our metrics affect their ability to distinguish nonisomorphic elections. Then, we study the complexity of deciding if a given representation indeed corresponds to some election.

$ C \times V $	ANECs	Positionwise	Pairwise	Bordawise
3×3	10	10	8	8
3×4	24	23	17	13
3×5	42	40	25	18
4×3	111	93	50	37
4×4	762	465	200	76
4×5	4095	1746	513	131

Table 1: Number of equivalence classes under our metrics.

Given a metric, two elections are in the same equivalence class if their distance is zero. An *anonymous, neutral equivalence class (ANEC)* consists of all isomorphic elections with a given number of candidates and voters [Eğecioğlu and Giritligil, 2013]. While ANECs are the equivalence classes of the swap and discrete metrics, the other metrics are less precise and their equivalence classes are unions of some ANECs.

To get a feeling as to how much precision is lost due to various aggregate representations, in Table 1 we compare the number of ANECs and the numbers of equivalence classes of the positionwise, pairwise, and Bordawise metrics, for small elections; we computed the table using exhaustive search² (note that EMD- and ℓ_1 -positionwise metrics have the same equivalence classes). Among these metrics, positionwise ones perform best and Bordawise performs worst. Next, we provide a partial theoretical explanation for this observation.

We say that a metric d is at least as fine as a metric d' if for each two elections A and B , $d(A, B) = 0$ implies that $d'(A, B) = 0$ (i.e., each equivalence class of d is a subset of some equivalence class of d'). Metric d is finer than d' if it is at least as fine as d' but d' is not at least as fine as d .

Proposition 1. *Swap and discrete isomorphic metrics are finer than EMD/ ℓ_1 -positionwise and pairwise, which both are finer than Bordawise. Neither EMD/ ℓ_1 -positionwise is finer than pairwise nor the other way round.*

Nonetheless, having many equivalence classes does not automatically make a metric desirable. For example, the discrete isomorphic metric is as fine as the swap one, but has many unappealing properties.

Let us now move to the problem of deciding if a given matrix/vector indeed represents some election. For position matrices, Boehmer *et al.* [2021b] have shown this to be easy: If a matrix M has nonnegative entries and its rows and columns sum up to the same value, then there is a polynomial-time computable election E with $M = \mathcal{P}_E$. Hence, one can work directly on the matrices and recover elections when needed. Unfortunately, for Borda score vectors and weighted majority relations analogous problems are NP-complete. Thus, by operating on them directly, we may leave the space of elections.

Theorem 1. Given a vector x of nonnegative integers, it is NP-complete to decide if there is an election E with $\mathcal{B}_E = x$.

The hardness of deciding whether there is an election with a given Borda score vector is not too surprising because it is closely related to strategic voting under the Borda rule, which is also NP-complete [Betzler *et al.*, 2011]. The case of weighted majority relation is more intriguing because

²There are exact formulas for some columns in Table 1, but not for all. See, e.g., the work of Eğecioğlu and Giritligil [2013].

the classic McGarvey’s theorem [McGarvey, 1953] gives a polynomial-time algorithm for recovering an election with a given *relative* weighted majority relation (but see also the work of Bachmeier *et al.* [2019]).

Theorem 2. Given an $m \times m$ matrix M , it is NP-complete to decide if there is an election E with $\mathcal{M}_E = M$.

Proof (First Part). We reduce from the NP-complete RESTRICTED X3C problem (RX3C) [Gonzalez, 1985]. Its instances consist of a universe $X = \{x_1, \dots, x_{3t}\}$ and a family $\mathcal{S} = \{S_1, \dots, S_{3t}\}$ of size-3 subsets of X , where each x_i appears in exactly three sets from \mathcal{S} . We ask if there is a family $\mathcal{S}' \subseteq \mathcal{S}$ of t sets such that $\bigcup_{S_i \in \mathcal{S}'} S_i = X$ (i.e., we ask if there is an exact cover of X). Let (X, \mathcal{S}) be an instance of RX3C. We form a candidate set:

$$C = \{d\} \cup X \cup \{s_1, \dots, s_{3t}\} \cup \{b_1, b_2, b_3\} \cup A \cup F,$$

where $A = \{a_1, \dots, a_{3t}\}$ and $F = \{f_1, \dots, f_{3t}\}$ are sets of “location” candidates. The candidates s_1, \dots, s_{3t} correspond to the sets from \mathcal{S} , b_1, b_2, b_3 will delineate blocks in the votes, and d will be used to encode the exact cover (if one exists). For each $i \in [3t]$, by $A(i)$ and $F(i)$ we mean the orderings:

$$\begin{aligned} A(i): & a_1 \succ \dots \succ a_i \succ s_i \succ a_{i+1} \succ \dots \succ a_{3t}, \\ F(i): & f_1 \succ s_1 \succ f_2 \succ s_2 \succ \dots \succ f_{i-1} \succ s_{i-1} \\ & \succ f_i \succ f_{i+1} \succ s_{i+1} \succ \dots \succ f_{3t} \succ s_{3t}. \end{aligned}$$

For a subset $C' \subseteq C$ of candidates, $[C']$ denotes an arbitrary, fixed ordering of the candidates from C' . For every $S_i = \{x_j, x_k, x_\ell\} \in \mathcal{S}$, let v_i denote the following vote:

$$b_1 \succ x_j \succ x_k \succ x_\ell \succ b_2 \succ A(i) \succ [X \setminus S_i] \succ b_3 \succ F(i).$$

Let $E' = (C \setminus \{d\}, (v_1, \dots, v_{3t}))$ be an election. We form a weighted majority relation M by taking $\mathcal{M}_{E'}$ and extending it to include d as follows: We require that d is placed behind b_1 in exactly t votes, d is placed behind each candidate from X in exactly one vote, and d is never placed behind any other candidate. We claim that there is an election E such that $M = \mathcal{M}_E$ if and only if \mathcal{S} contains an exact cover of X .

First, let $\mathcal{S}' \subseteq \mathcal{S}$ be an exact cover of X . To construct E , for each $S_i \in \mathcal{S}'$, we include vote v_i with d inserted right in front of b_2 , and for each $S_i \in \mathcal{S} \setminus \mathcal{S}'$, we include v_i with d inserted before b_1 . As \mathcal{S}' is an exact cover, $M = \mathcal{M}_E$.

The other direction is available in our full version. \square

5 Diameter and Compass Elections

In this section, we analyze the distances between four “compass” elections of Boehmer *et al.* [2021b]. These elections capture four different types of (dis)agreement among the voters and, thus, we expect good metrics to put them far apart. Since the smallest nonzero distances under all our metrics are either 1, 2, or 4, the larger are the distances between the compass elections, the more space there is between them for other elections. For technical reasons, we fix the number m of candidates to be even, and the number of voters to be $n = t \cdot m!$, where t is some positive integer (we will see ways to relax this assumption). The compass elections are defined as follows:

1. In the identity elections, denoted ID, all voters have the same, fixed preference order.

2. In the antagonism elections, denoted AN, half of the voters rank the candidates in one way and half of the voters rank them in the opposite way.
3. In the uniformity elections, denoted UN, each possible vote appears the same number of times.
4. In the stratification elections, denoted ST, the candidates are partitioned into two equal-sized sets A and B . Each possible preference order where all members of A are ranked ahead of B appears the same number of times.

The next proposition gives asymptotic distances between the compass elections.

Proposition 2. Let X and Y be two distinct compass elections. Then, $d_{\text{Borda}}(X, Y) = \Theta(nm^3)$, $d_{\text{swap}}(X, Y) = d_{\text{pos}}^{\text{emd}}(X, Y) = d_{\text{pair}}(X, Y) = \Theta(nm^2)$, $d_{\text{pos}}^{\ell_1}(X, Y) = \Theta(nm)$, and $d_{\text{disc}}(X, Y) = \Theta(n)$, except that $d_{\text{pair}}(\text{AN}, \text{UN}) = d_{\text{Borda}}(\text{AN}, \text{UN}) = 0$.

The distances between the compass elections are the largest under the swap and EMD-positionwise metrics, followed by those under ℓ_1 -positionwise and discrete isomorphic metrics. Pairwise and Bordawise metrics perform particularly badly because they cannot distinguish between UN and AN. Yet, except for this, the compass elections are asymptotically as distant under each of our metrics as possible. Indeed, ID and UN even form diameters of our election spaces (this also confirms a conjecture of Boehmer *et al.* [2021b]).

Theorem 3. Let d be one of our six metrics. For each two elections X and Y (with sizes as specified at the beginning of this section) it holds that $d(X, Y) \leq d(\text{ID}, \text{UN})$.

Proof (EMD-positionwise). The EMD-positionwise metric can be seen as working over matrices whose entries are nonnegative and whose rows and columns sum up to the same value. Let A and B be two such matrices, whose rows and columns sum up to some value n' . Let A/n' and B/n' be the same matrices, but with their entries divided by n' . Note that $d_{\text{pos}}^{\text{emd}}(A, B) = n' \cdot d_{\text{pos}}^{\text{emd}}(A/n', B/n')$. From now on, we focus on matrices whose entries are nonnegative and whose rows and columns sum up to 1 (they are called *bistochastic*).

Boehmer *et al.* [2021b] have shown that $d_{\text{pos}}^{\text{emd}}(\text{ID}, \text{UN}) = n(m^2 - 1)/3$. We claim that for each two $m \times m$ bistochastic matrices X and Y it holds that $d_{\text{pos}}^{\text{emd}}(X, Y) \leq (m^2 - 1)/3$. For the sake of contradiction, assume that the opposite holds. Let x_1, \dots, x_m be the columns of X and y_1, \dots, y_m be the columns of Y . Without loss of generality, we assume that $d_{\text{pos}}^{\text{emd}}(X, Y) = \sum_{i \in [m]} \text{emd}(x_i, y_i)$; otherwise we could reorder the columns of one of the matrices. By definition of the EMD metric, we have that $\sum_{i \in [m]} \text{emd}(x_i, y_i) = \sum_{i \in [m]} \ell_1(\hat{x}_i, \hat{y}_i)$. For each $i \in [m]$, we write $\hat{x}_{i,1}, \dots, \hat{x}_{i,m}$ to denote the entries of the cumulative vector \hat{x}_i ; we use analogous notation for \hat{y}_i . Note that in the matrices with columns $\hat{x}_1, \dots, \hat{x}_m$ and $\hat{y}_1, \dots, \hat{y}_m$, for each $j \in [m]$, the j -th row sums up to j (we refer to this as the *cumulative rows* property). Using these observations, we note that:

$$\begin{aligned} \sum_{i \in [m]} \text{emd}(x_i, y_i) &= \sum_{i \in [m]} \ell_1(\hat{x}_i, \hat{y}_i) \\ &= \sum_{i, j \in [m]} |\hat{x}_{i,j} - \hat{y}_{i,j}| \end{aligned}$$

$$\begin{aligned}
 &= \sum_{i,j \in [m]} (\max(\hat{x}_{i,j}, \hat{y}_{i,j}) - \min(\hat{x}_{i,j}, \hat{y}_{i,j})) \\
 &= \sum_{i,j \in [m]} (\hat{x}_{i,j} + \hat{y}_{i,j}) - 2 \sum_{i,j \in [m]} \min(\hat{x}_{i,j}, \hat{y}_{i,j}) \\
 &= 2 \left(\sum_{j \in [m]} j \right) - 2 \sum_{i,j \in [m]} \min(\hat{x}_{i,j}, \hat{y}_{i,j}) \\
 &= m(m+1) - 2 \sum_{i,j \in [m]} \min(\hat{x}_{i,j}, \hat{y}_{i,j}).
 \end{aligned}$$

Thus, if $d_{\text{pos}}^{\text{emd}}(X, Y) > (m^2 - 1)/3$, then it must hold that:

$$\begin{aligned}
 \sum_{i,j \in [m]} \min(\hat{x}_{i,j}, \hat{y}_{i,j}) &< \frac{1}{2} \left(m(m+1) - \frac{m^2-1}{3} \right) \\
 &= (2m^2 + 3m + 1)/6 = (2m+1)(m+1)/6.
 \end{aligned}$$

In the following, for each $i, k \in [m]$ with $i+k > m$, if $i+k$ is used as a column index, then we take it to be $i+k-m$ (i.e., column indices “cycle”). For each $k \in [m]$, we have $d_{\text{pos}}^{\text{emd}}(X, Y) \leq \sum_{i \in [m]} \text{emd}(x_i, y_{i+k})$; if this were not the case, then our assumption that $d_{\text{pos}}^{\text{emd}}(X, Y) = \sum_{i \in [m]} \text{emd}(x_i, y_i)$ would have been false. Consequently, for every $k \in [m]$, repeating the above reasoning, we get:

$$\sum_{i,j \in [m]} \min(\hat{x}_{i,j}, \hat{y}_{i+k,j}) < (2m+1)(m+1)/6,$$

If $a, b \in [0, 1]$ then $a \cdot b \leq \min(a, b)$. As for each $i, j \in [m]$ we have $\hat{x}_{i,j}, \hat{y}_{i,j} \in [0, 1]$, for each $k \in [m]$ we have:

$$\sum_{i,j \in [m]} \hat{x}_{i,j} \cdot \hat{y}_{i+k,j} < (2m+1)(m+1)/6.$$

By summing this inequality sidewise for all $k \in [m]$, we get:

$$\sum_{i,j \in [m]} \hat{x}_{i,j} \cdot \sum_{k \in [m]} \hat{y}_{k,j} < m(2m+1)(m+1)/6.$$

By applying the cumulative rows property, we obtain:

$$\sum_{j=1}^m j^2 < m(2m+1)(m+1)/6.$$

Since we know that $\sum_{j=1}^m j^2 = m(2m+1)(m+1)/6$, this is a contradiction. Hence, for all $m \times m$ bistochastic matrices X, Y we have $d_{\text{pos}}^{\text{emd}}(X, Y) \leq (m^2 - 1)/3$. Thus, for all elections with m candidates and n voters, their EMD-positionwise distance is at most $n(m^2 - 1)/3$. \square

In the above proof we do not work directly with elections, but, rather, with normalized position matrices. Viewed this way, ID is a unit diagonal matrix and UN is a matrix whose entries are all equal. Indeed, this is how Boehmer *et al.* [2021b] defined them. In this way, the proof works for any number of voters (this also applies to ℓ_1 -positionwise and, using normalized weighted majority relations, to pairwise).

For Bordawise, the proof of Theorem 3 shows that for each election the sum of its distances from ID and UN is the same. That is, under this metric every election lays on the diameter.

6 Maps and Correlations

While in the previous section we studied distances between hand-crafted elections, now we analyze automatically-generated ones. We test how our metrics correlate with the swap one, and we compare their maps of elections.

We use two datasets. The first one consists of all small elections, as in Section 4. The second one resembles those used in the maps of Szufa *et al.* [2020] and Boehmer *et al.* [2021b], but consists of elections with 10 candidates and 50 voters,³ generated according to the following statistical models (see the just-cited papers for more details):

$ C \times V $	EMD-Pos.	ℓ_1 -Pos.	ℓ_1 -Pair.	Bordawise	Discrete
3×3	0.942	0.748	0.860	0.587	0.614
3×4	0.900	0.697	0.860	0.659	0.636
3×5	0.920	0.759	0.843	0.606	0.680
4×3	0.850	0.577	0.735	0.442	0.402
4×4	0.782	0.561	0.689	0.415	0.434
4×5	0.772	0.567	0.672	0.439	0.432
10×50 (340 elections)	0.745	0.563	0.708	0.430	0.342

Table 2: Pearson correlation coefficients between swap distances and the other ones computed for our datasets.

IC, Urn, and Mallows. We generated 20 elections using the impartial culture model (IC), where each vote is selected uniformly at random, and 60 elections for each of the classic urn and Mallows models (we used the same sampling protocol as Boehmer *et al.* [2021b]).

SP, SC, and SPOC. We generated 20 single-peaked elections (SP elections) uniformly at random (this is known as the SP Walsh model [Walsh, 2015]), 20 such elections using the Conitzer model [Conitzer, 2009], and 20 single-peaked on a circle elections (SPOC elections), uniformly at random [Peters and Lackner, 2020]. We also generated 20 single-crossing elections (SC elections) using the sampling protocol of Szufa *et al.* [2020].

Euclidean. In these elections, each candidate and voter is a point from some Euclidean space and the voters rank the candidates with respect to their increasing distances from them. We have generated the points uniformly at random from (i) a 1D interval, (ii) a 2D sphere, (iii) a 2D disc, and (iv) a 3D cube; in each case we generated 20 elections.

Group-Separable. Group-separable elections were introduced by Inada [1964; 1969]. We use a definition based on trees (see, e.g., the works of Karpov [2019] and Faliszewski *et al.* [2022] for a discussion and motivation). Consider an ordered, rooted tree where each leaf is a unique candidate. To obtain a vote, for each of the nodes we can choose to reverse the order of its children and, then, rank the candidates by reading them off from left to right. Given such a tree, we sample a vote by reversing the order of each node’s children with probability $1/2$. We generated 20 elections using complete binary trees, and 20 elections using binary caterpillar trees (a binary caterpillar tree is defined recursively to either be a leaf, or a root whose one child is a leaf and whose other child is a root of a binary caterpillar tree.)

For each dataset we have computed the Pearson Correlation Coefficient (PCC) between the swap distances and those provided by the other metrics. PCC is a classic measure of correlation that takes values between -1 and 1 ; its absolute value gives the strength of the correlation and the sign indicates its positive or negative nature. Szufa *et al.* [2020]

³This is the largest size for which we could compute swap distances within a few weeks. Szufa *et al.* [2020] used 100×100 elections, and Boehmer *et al.* [2021b] used 10×100 ones.

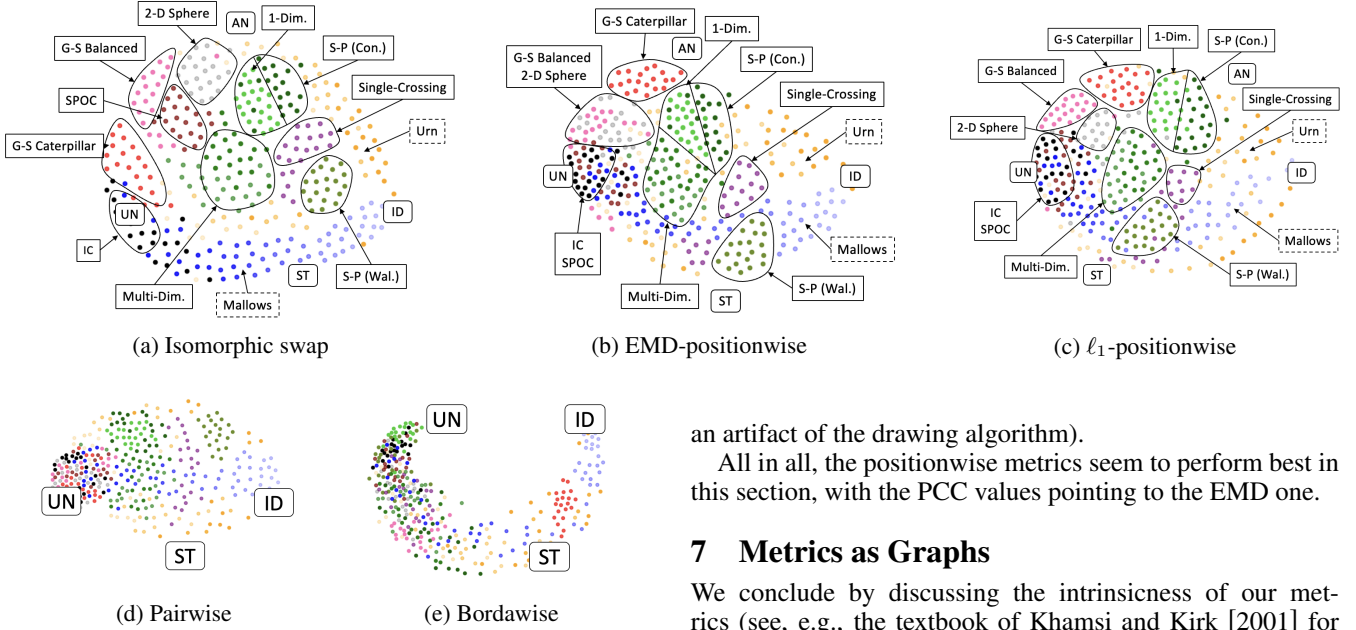


Figure 1: Maps of elections prepared using each of our metrics.

presented a similar experiment, but on a much smaller scale, and on a limited set of metrics. We present our results in Table 2. We see that EMD-positionwise is most strongly correlated with the swap metric, with a large advantage over all the metrics, except for pairwise (where the advantage is smaller).

Next, we used the techniques of Szufa *et al.* [2020] to draw the maps of elections from the 10×50 dataset. We show these maps in Figure 1 (we omit the discrete metric as its visualization is nearly meaningless); each dot is an election, its color corresponds to the statistical model it comes from, and the points are placed so that their Euclidean distances resemble those according to a given metric as much as possible (following Szufa *et al.* [2020], we used the algorithm of Fruchterman and Reingold [1991] to place the points). We also included the compass elections on the maps (for the swap metric, their location is approximate).

The maps provided by the isomorphic swap distance and both positionwise metrics are remarkably similar, but, nonetheless, there are some differences. For example, under the swap metric group-separable caterpillar elections are closer to the IC ones than the group-separable balanced elections (both on the map and in terms of actual distances), whereas according to the positionwise metrics this relation is reversed. Also, ℓ_1 -positionwise clearly distinguishes between 2D-Sphere and group-separable balanced elections (like the swap metric), but EMD-positionwise does not. Generally, the area between UN and AN is quite challenging for our metrics (fortunately, according to Boehmer *et al.* [2021b], only few real-life elections land there). The maps for the pairwise and Bordawise metrics illustrate their flaws identified in the previous sections (e.g., Bordawise and pairwise conflate UN and AN, and the former also puts all the elections on the diameter, which explains its elongated shape; the curvature is

an artifact of the drawing algorithm).

All in all, the positionwise metrics seem to perform best in this section, with the PCC values pointing to the EMD one.

7 Metrics as Graphs

We conclude by discussing the intrinsicness of our metrics (see, e.g., the textbook of Khamsi and Kirk [2001] for more details on this notion). Consider a graph whose vertices are equivalence classes of a given metric and where the edges connect those classes that are at minimum nonzero distance from each other. Under the swap distance, each edge corresponds to a swap of adjacent candidates and each shortest path corresponds to the distance between its endpoints. This means that the swap distance is intrinsic.

Definition 1. Let $\alpha \geq 1$ be a number. A metric d is α -intrinsic if for each pair of elections X and Y (with the same number of candidates and the same number of voters), there are elections $X = E_0, E_1, E_2, \dots, E_{k-1}, E_k = Y$ such that $\sum_{i=1}^k d(E_{i-1}, E_i) \leq \alpha \cdot d(X, Y)$, and for each $i \in [k]$, $d(E_{i-1}, E_i)$ is the smallest nonzero distance between elections under d . If $\alpha = 1$, then d is intrinsic.

An intrinsic metric can be viewed as performing a series of simple, unit operations. Among our metrics only swap and discrete are intrinsic, but the positionwise ones are 2-intrinsic (this is, perhaps, the most technically involved of our results).

Theorem 4. The swap and discrete metrics are intrinsic, but neither of the EMD/ ℓ_1 -positionwise, pairwise, and Bordawise metric is intrinsic, but the EMD- and ℓ_1 -positionwise metrics are 2-intrinsic yet not α -intrinsic for any $\alpha < 2$.

We conjecture that pairwise is not α -intrinsic for any $\alpha \geq 1$.

8 Summary

We found that the EMD- and ℓ_1 -positionwise metrics are quite similar to the swap one (which, computational issues aside, we view as ideal), but the EMD variant seems better. This justifies the choice of the EMD-positionwise metric for the maps of Szufa *et al.* [2020] and Boehmer *et al.* [2021b]. Yet, we ask for a metric that would perform even better, especially on elections between the uniform and antagonism ones. With our study of intrinsicness, we have initiated an axiomatic analysis of our metrics; it would be interesting to extend this analysis to better understand their properties.

Acknowledgements

NB was supported by the DFG project MaMu (NI 369/19) and by the DFG project ComSoc-MPMS (NI 369/22). TW was supported by the Polish National Science Center grant 2018/31/B/ST6/03201. 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

- [Bachmeier *et al.*, 2019] Georg Bachmeier, Felix Brandt, Christian Geist, Paul Harrenstein, Keyvan Kardel, Dominik Peters, and Hans Georg Seedig. k -majority digraphs and the hardness of voting with a constant number of voters. *Journal of Computer and System Sciences*, 105:130–157, 2019.
- [Bartholdi *et al.*, 1989] John Bartholdi, III, Craig Tovey, and Michael Trick. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare*, 6(2):157–165, 1989.
- [Betzler *et al.*, 2011] Nadja Betzler, Rolf Niedermeier, and Gerhard J. Woeginger. Unweighted coalitional manipulation under the Borda rule is NP-hard. In *Proceedings of IJCAI-2011*, pages 55–60, 2011.
- [Boehmer and Schaar, 2022] Niclas Boehmer and Nathan Schaar. Collecting, classifying, analyzing, and using real-world elections. Technical Report arXiv:2204.03589 [cs.GT], arXiv.org, 2022.
- [Boehmer *et al.*, 2021a] Niclas Boehmer, Robert Bredebeck, Piotr Faliszewski, and Rolf Niedermeier. Winner robustness via swap- and shift-bribery: Parameterized counting complexity and experiments. In *Proceedings of IJCAI-2021*, pages 52–58, 2021.
- [Boehmer *et al.*, 2021b] Niclas Boehmer, Robert Bredebeck, Piotr Faliszewski, Rolf Niedermeier, and Stanislaw Szufa. Putting a compass on the map of elections. In *Proceedings of IJCAI-2021*, pages 59–65, 2021.
- [Boehmer *et al.*, 2022] Niclas Boehmer, Piotr Faliszewski, Rolf Niedermeier, Stanislaw Szufa, and Tomasz Wąs. Understanding distance measures among elections. Technical Report arXiv:2205.00492 [cs.GT], arXiv.org, 2022.
- [Conitzer, 2009] Vincent Conitzer. Eliciting single-peaked preferences using comparison queries. *Journal of Artificial Intelligence Research*, 35:161–191, 2009.
- [Dwork *et al.*, 2001] Cynthia Dwork, Ravi Kumar, Moni Naor, and D. Sivakumar. Rank aggregation methods for the web. In *Proceedings of WWW-2001*, pages 613–622, March 2001.
- [Eğecioğlu and Giritligil, 2013] Ömer Eğecioğlu and Ayca Giritligil. The impartial, anonymous, and neutral culture model: A probability model for sampling public preference structures. *Journal of Mathematical Sociology*, 37(4):203–222, 2013.
- [Faliszewski *et al.*, 2019] Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, Stanislaw Szufa, and Nimrod Talmon. How similar are two elections? In *Proceedings of AAAI-2019*, pages 1909–1916, 2019.
- [Faliszewski *et al.*, 2022] Piotr Faliszewski, Alexander Karpov, and Svetlana Obraztsova. The complexity of election problems with group-separable preferences. *Autonomous Agents and Multi-Agent Systems*, 36(1):18, 2022.
- [Fruchterman and Reingold, 1991] Thomas Fruchterman and Edward Reingold. Graph drawing by force-directed placement. *Software: Practice and Experience*, 21(11):1129–1164, 1991.
- [Gonzalez, 1985] Teofilo Gonzalez. Clustering to minimize the maximum intercluster distance. *Theoretical Computer Science*, 38:293–306, 1985.
- [Grohe *et al.*, 2018] Martin Grohe, Gaurav Rattan, and Gerhard J. Woeginger. Graph similarity and approximate isomorphism. In *Proceedings of MFCS-2018*, pages 20:1–20:16, 2018.
- [Inada, 1964] Ken-ichi Inada. A note on the simple majority decision rule. *Econometrica*, 32(32):525–531, 1964.
- [Inada, 1969] Ken-ichi Inada. The simple majority decision rule. *Econometrica*, 37(3):490–506, 1969.
- [Karpov, 2019] Alexander Karpov. On the number of group-separable preference profiles. *Group Decision and Negotiation*, 28(3):501–517, 2019.
- [Khamsi and Kirk, 2001] Mohamed Khamsi and William Kirk. *An Introduction to Metric Spaces and Fixed Point Theory*. Wiley, 2001.
- [Mattei and Walsh, 2013] Nicholas Mattei and Toby Walsh. Preflib: A library for preferences. In *Proceedings of ADT-2013*, pages 259–270, 2013.
- [McGarvey, 1953] David McGarvey. A theorem on the construction of voting paradoxes. *Econometrica*, 21(4):608–610, 1953.
- [Peters and Lackner, 2020] Dominik Peters and Martin Lackner. Preferences single-peaked on a circle. *Journal of Artificial Intelligence Research*, 68:463–502, 2020.
- [Rubner *et al.*, 2000] Yossi Rubner, Carlo Tomasi, and Leonidas J. Guibas. The earth mover’s distance as a metric for image retrieval. *International Journal of Computer Vision*, 40(2):99–121, 2000.
- [Szufa *et al.*, 2020] Stanislaw Szufa, Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. Drawing a map of elections in the space of statistical cultures. In *Proceedings of AAMAS-2020*, pages 1341–1349, 2020.
- [Walsh, 2015] Toby Walsh. Generating single peaked votes. Technical Report arXiv:1503.02766 [cs.GT], arXiv.org, March 2015.