

How to Sample Approval Elections?

Stanisław Szufa¹, Piotr Faliszewski¹, Łukasz Janeczko¹, Martin Lackner², Arkadii Slinko³, Krzysztof Sornat¹ and Nimrod Talmon⁴

¹AGH University, Poland

²DBAI, TU Wien, Austria

³University of Auckland, New Zealand

⁴Ben-Gurion University of the Negev, Israel

Abstract

We extend the map-of-elections framework to the case of approval elections. While doing so, we study a number of statistical cultures, including some new ones, and we analyze their properties. We find that approval elections can be understood in terms of the average number of approvals in the votes, and the extent to which the votes are chaotic.

1 Introduction

In an approval election [Brams and Fishburn, 1983], each voter indicates which candidates he or she finds acceptable for a certain task (e.g., to be a president, to join the parliament, or to enter the final round of a competition) and a voting rule aggregates this data into the final outcome. In the single-winner setting (e.g., when choosing the president), the most popular rule is to pick the candidate with the highest number of approvals. In the multiwinner setting (e.g., in parliamentary elections or when choosing finalists in a competition) there is a rich spectrum of rules to choose from, each with different properties and advantages. Approval voting is particularly attractive due to its simplicity and low cognitive load imposed on the voters. Indeed, its practical applicability has already been tested in a number of field experiments, including those in France [Laslier and der Straeten, 2008; Baujard and Igersheim, 2011; Bouveret *et al.*, 2019] and Germany [Alós-Ferrer and Granić, 2012]. Over the recent years there was also tremendous progress regarding its theoretical properties (see, e.g., the overviews of Laslier and Sanver [2010] and Lackner and Skowron [2021]).

In spite of all these achievements, numerical experiments regarding approval voting are still challenging to design. One of the main difficulties is caused by the lack of consensus on which statistical election models to use. Below we list a few models (i.e., statistical cultures) that were recently used:

1. In the impartial culture setting, we assume that each vote is equally likely. Taken literally, this means that each voter approves each candidate with probability $1/2$ [Barrot *et al.*, 2017]. As this is quite unrealistic, several authors treat the approval probability as a parameter [Bredereck *et al.*, 2019; Faliszewski *et al.*, 2020] or require that all voters approve the same (small) number of can-

didates [Lackner and Skowron, 2020]. A further refinement is to choose an individual approval probability for each candidate [Lackner and Maly, 2021].

2. In Euclidean models, each candidate and voter is a point in \mathbb{R}^d , where d is a parameter, and a voter approves a candidate if they are sufficiently near. Such models are used, e.g., by Bredereck *et al.* [2019] and Godziszewski *et al.* [2021]. Naturally, the distribution of the candidate and voter points strongly affects the outcomes.
3. Some authors consider statistical cultures designed for the ordinal setting (where the voters rank the candidates from the most to the least desirable one) and let the voters approve some top-ranked candidates (e.g., a fixed number of them). This approach is taken, e.g., by Lackner and Skowron [2020] on top of the ordinal Mallows model (later on, Caragiannis *et al.* [2022] provided approval-based analogues of the Mallows model).

Furthermore, even if two papers use the same model, they often choose its parameters differently. Since it is not clear how the parameters affect the models, comparing the results from different papers is not easy. Our goal is to initiate a systematic study of approval-based statistical cultures and attempt to rectify at least some of these issues. We do so by extending the map-of-elections framework of Szufa *et al.* [2020] and Boehmer *et al.* [2021] to the approval setting.

Briefly put, a map-of-elections is a set of elections with a distance between each pair. We embed these elections in a plane, by representing each election as a point, so that the Euclidean distances between these points are similar to the original distances between the respective elections. Such maps help in obtaining insights about the elections they contain.

To create a map-of-elections for approval elections, we start by identifying two metrics between approval elections, the *isomorphic Hamming distance* and the *approvalwise distance*. The first one is accurate, but difficult to compute, whereas the second one is less precise, but easily computable. Fortunately, in our datasets the two metrics are strongly correlated, so we use the latter one.

Next, we analyze the space of approval elections with a given number of candidates and voters. For each $p \in [0, 1]$, by p -identity (p -ID) elections we mean those where all the votes are identical and approve the same p -fraction of candidates. By p -impartial culture (p -IC) elections we mean those

where each voter chooses to approve each candidate with probability p . We view p -ID and p -IC elections as two extremes on the spectrum of agreement between the voters and, intuitively, we expect that every election (where each voter approves on average a p fraction of candidates) is located somewhere between these two. In particular, for $p, \phi \in [0, 1]$, we introduce the (p, ϕ) -resampling model, which generates elections whose expected approvalwise distance from p -ID is exactly the ϕ fraction of the distance between p -ID and p -IC.

Armed with these tools, we proceed to draw maps of elections. First, we consider p -ID, p -IC, and (p, ϕ) -resampling elections, where the p and ϕ values are chosen to form a grid, and compute the approvalwise distances between them. We find that, for a fixed value of p , the (p, ϕ) -resampling elections indeed form lines between the p -ID and p -IC ones, whereas for fixed ϕ values they form lines between 0-ID and 1-ID ones (which we refer to as the *empty* and *full* elections). We obtain more maps by adding elections generated according to other statistical cultures; the presence of the (p, ϕ) -resampling grid helps in understanding the locations of these new elections. For each of our elections we compute several parameters, such as, e.g. the highest number of approvals that a candidate receives, the time required to compute the results of a certain multiwinner voting rule, or the cohesiveness level (see Section 2 for a definition). For each of the statistical cultures, we present maps where we color the elections according to these values. This gives further insight into the nature of the elections they generate. Finally, we compare the results for randomly generated elections with those appearing in real-life, in the context of participatory budgeting.

2 Preliminaries

For a given positive integer t , we write $[t]$ to denote the set $\{1, 2, \dots, t\}$, and $[t]_0$ as an abbreviation for $[t] \cup \{0\}$.

Elections. A (simple) approval election $E = (C, V)$ consists of a set of candidates $C = \{c_1, \dots, c_m\}$ and a collection of voters $V = \{v_1, \dots, v_n\}$. Each voter $v \in V$ casts an approval ballot, i.e., he or she selects a subset of candidates that he or she approves. Given a voter v , we denote this subset by $A(v)$. Occasionally, we refer to the voters or their approval ballots as votes; the exact meaning will always be clear from the context. An approval-based committee election (an ABC election) is a triple (C, V, k) , where (C, V) is a simple approval election and k is the size of the desired committee. We use simple elections when the goal is to choose a single individual, and ABC elections in the multiwinner setting. Given an approval election E (be it a simple election or an ABC election) and a candidate c , we write $score_{av}(c)$ to denote the number of voters that approve c . We refer to this value as the *approval score* of c . The single-winner approval rule (called AV) returns the candidate with the highest approval score (or the set of such candidates, in case of a tie).

Distances Between Votes. For two voters v and u , their Hamming distance is $ham(v, u) = |A(v) \Delta A(u)|$, i.e., the number of candidates approved by exactly one of them. Other distances include, e.g., the Jaccard one, defined as $jac(v, u) = \frac{ham(v, u)}{|A(v) \cup A(u)|}$. For other examples of such dis-

tances, we point to the work of Caragiannis *et al.* [2022].

Approval-Based Committee Voting Rules. An *approval-based committee voting rule* (an ABC rule) is a function that maps an ABC election (C, V, k) to a nonempty set of committees of size k . If an ABC rule returns more than one committee, then we consider them tied.

We introduce two prominent ABC rules. *Multiwinner Approval Voting (AV)* selects the k candidates with the highest approval scores. Given a committee W , its approval score is the sum of the scores of its members; $score_{av}(W) = \sum_{w \in W} score_{av}(w)$. If there is more than one committee that achieves a maximum score, AV returns all tied committees. The second rule is *Proportional Approval Voting (PAV)*. PAV outputs all committees with maximum PAV-score:

$$score_{pav}(W) = \sum_{v \in V} h(|A(v) \cap W|),$$

where $h(x) = \sum_{j=1}^x 1/j$ is the harmonic function. Intuitively, AV selects committees that contain the “best” candidates (in the sense of most approvals) and PAV selects committees that are in a strong sense proportional [Aziz *et al.*, 2017]. In contrast to AV, which is polynomial-time computable, PAV is NP-hard to compute [Aziz *et al.*, 2015; Skowron *et al.*, 2016]. In practice, PAV can be computed by solving an integer linear program [Peters and Lackner, 2020] or by an approximation algorithm [Dudycz *et al.*, 2020].

Cohesive Groups. Intuitively, a proportional committee should represent all groups of voters in a way that (roughly) corresponds to their size. To speak of proportional committees in ABC elections, Aziz *et al.* [2017] introduced the concept of *cohesive groups*.

Definition 1. Consider an ABC election (C, V, k) with n voters and some non-negative integer ℓ . A group of voters $V' \subseteq V$ is ℓ -cohesive if (i) $|V'| \geq \ell \cdot \frac{n}{k}$ and (ii) $|\bigcap_{v \in V'} A(v)| \geq \ell$.

An ℓ -cohesive group is large enough to deserve ℓ representatives in the committee and is cohesive in the sense that there are ℓ candidates that can represent it. A number of proportionality notions have been proposed based on cohesive groups, such as (extended) justified representation [Aziz *et al.*, 2017], proportional justified representation [Sánchez-Fernández *et al.*, 2017], proportionality degree [Skowron, 2021], and others. All these notions make certain proportionality guarantees for the cohesive groups (see also the survey of [Lackner and Skowron, 2021] for a comprehensive overview).

3 Statistical Cultures for Approval Elections

In the following, we present several statistical cultures (probabilistic models) for generating approval elections. Our input consists of the desired number of voters n and a set of candidates $C = \{c_1, \dots, c_m\}$. For models that already exist in the literature, we provide examples of papers that use them.

Resampling, IC, and ID Models. Let p and ϕ be two numbers in $[0, 1]$. In the (p, ϕ) -resampling model, we first draw a central ballot u , by choosing $\lfloor p \cdot m \rfloor$ approved candidates uniformly at random; then, we generate each new vote v by initially setting $A(v) = A(u)$ and executing the following procedure for every candidate $c_i \in C$: With probability $1 - \phi$, we

leave c_i 's approval intact and with probability ϕ we resample its value (i.e., we let c_i be approved with probability p). The resampling model is due to this paper and is one of our basic tools for analyzing approval elections. By fixing $\phi = 1$, we get the p -impartial culture model (p -IC) where each candidate in each vote is approved with probability p ; it was used, e.g., by Bredereck *et al.* [2019] and Faliszewski *et al.* [2020]. By fixing $\phi = 0$, we ensure that all votes in an election are identical (i.e., approve the same p fraction of the candidates). We refer to this model as p -identity (p -ID).

Disjoint Model. The (p, ϕ, g) -disjoint model, where p and ϕ are numbers in $[0, 1]$ and g is a non-negative integer, works as follows: We draw a random partition of C into g sets, C_1, \dots, C_g , and, to generate a vote, we choose $i \in [g]$ uniformly at random and sample the vote from a (p, ϕ) -resampling model with the central vote that approves exactly the candidates from C_i (while the central votes are independent of p , we still need this parameter for resampling).

Noise Models. Let p and ϕ be two numbers from $[0, 1]$ and let d be a distance between approval votes. We require that d is polynomial-time computable and, for each two approval votes u and v , $d(u, v)$ depends only on $|A(u)|$, $|A(v)|$, and $|A(u) \cap A(v)|$; both distances from Section 2 have this property. In the (p, ϕ, d) -Noise model we first generate a central vote u as in the resampling model and, then, each new vote v is generated with probability proportional to $\phi^{d(u, v)}$. Such noise models are analogous to the Mallows model for ordinal elections and were studied, e.g., by Caragiannis *et al.* [2022]. In particular, they gave a sampling procedure for the case of the Hamming distance. We extend it to arbitrary distances.

Proposition 1. *There is a polynomial-time sampling procedure for the (p, ϕ, d) -noise models (as defined above).*

In the reminder, we only use the noise model with the Hamming distance and we refer to it as the (p, ϕ) -noise model. Note that the roles of p and ϕ in this model are similar, but not the same, as in the resampling one (for example, for $\phi = 0$ we get the p -ID model, but for $\phi = 1$ get the 0.5-IC one).

Euclidean Models. In the t -dimensional Euclidean model, each candidate and each voter is a point from \mathbb{R}^t and a voter v approves candidate c if the distance between their points is at most r (this value is called the radius); such models were discussed, e.g., in the classical works of Enelow and Hinich [1984, 1990], and more recently by Elkind and Lackner [2015], Elkind *et al.* [2017], Bredereck *et al.* [2019], and Godziszewski *et al.* [2021]. We consider t -dimensional models for $t \in \{1, 2\}$, where the agents' points are distributed uniformly at random on $[0, 1]^t$. We refer to them as 1D-Uniform and 2D-Square models (note that to fully specify each of them, we also need to indicate the radius value).

Truncated Urn Models. Let p be a number in $[0, 1]$ and let α be a non-negative real number (the parameter of contagion). In the truncated Pólya-Eggenberger Urn Model [Berg, 1985] we start with an urn that contains all $m!$ possible linear orders over the candidate set. To generate a vote, we (1) draw a random order r from the urn, (2) produce an approval vote that consists of $\lceil p \cdot m \rceil$ top candidates according to r (this is the generated vote), and (3) return $\alpha m!$ copies of r to the

urn. For $\alpha = 0$, all votes with $\lceil p \cdot m \rceil$ approved candidates are equally likely, whereas for large values of α all votes are likely to be identical (so the model becomes similar to p -ID).

4 Metrics

Next we describe two (pseudo)metrics used to measure distances between approval elections. Since we are interested in distances between randomly generated elections, our metrics are independent of renaming the candidates and voters.

Consider two equally-sized candidate sets C and D , and a voter v with a ballot over C . For a bijection $\sigma: C \rightarrow D$, by $\sigma(v)$ we mean a voter with an approval ballot $A(\sigma(v)) = \{\sigma(c) \mid c \in C\}$. In other words, $\sigma(v)$ is the same as v , but with the candidates renamed by σ . We write $\Pi(C, D)$ to denote the set of all bijections from C to D . For a positive integer n , by S_n we mean the set of all permutations over $[n]$. Next, we define the isomorphic Hamming distance (inspired by the metrics of Faliszewski *et al.* [2019]).

Definition 2. *Let $E = (C, V)$ and $F = (D, U)$ be two elections, where $|C| = |D|$, $V = (v_1, \dots, v_n)$ and $U = (u_1, \dots, u_n)$. The isomorphic Hamming distance between E and F , denoted $d_{\text{ham}}(E, F)$, is defined as:*

$$\min_{\sigma \in \Pi(C, D)} \min_{\rho \in S_n} \left(\sum_{i=1}^n \text{ham}(\sigma(v_i), u_{\rho(i)}) \right).$$

Intuitively, under the isomorphic Hamming distance we unify the names of the candidates in both elections and match their voters to minimize the sum of the resulting Hamming distances. We call this distance *isomorphic* because its value is zero exactly if the two elections are identical, up to renaming the candidates and voters. Computing this distance is NP-hard (see also the related results for approximate graph isomorphism [Arvind *et al.*, 2012; Grohe *et al.*, 2018]).

Proposition 2. *Computing the isomorphic Hamming distance between two approval elections is NP-hard.*

Thus we compute this distance using a brute-force algorithm (which is faster than using, e.g., ILP formulations). Since this limits the size of elections we can deal with, we also introduce a simple, polynomial-time computable metric.

Definition 3. *Let E be an election with candidate set $\{c_1, \dots, c_m\}$ and n voters. Its approvalwise vector, denoted $av(E)$, is obtained by sorting the vector $(\text{score}_{av}(c_1)/n, \dots, \text{score}_{av}(c_m)/n)$ in non-increasing order.*

Definition 4. *The approvalwise distance between elections E and F with approvalwise vectors $av(E) = (x_1, \dots, x_m)$ and $av(F) = (y_1, \dots, y_m)$ is defined as:*

$$d_{\text{app}}(E, F) = |x_1 - y_1| + \dots + |x_m - y_m|.$$

In other words, the approvalwise vector of an election is a sorted vector of the normalized approval scores of its candidates, and an approvalwise distance between two elections is the ℓ_1 distance between their approvalwise vectors. We sort the vectors to avoid the explicit use of a candidate matching, as is needed in the Hamming distance. Occasionally we will speak of approvalwise distances between approvalwise vectors, without referring to the elections that provide them.

It is clear that the approvalwise distance is computable in polynomial time. Yet, it is not as clear that its values are




Figure 1: Distances between resampling elections.

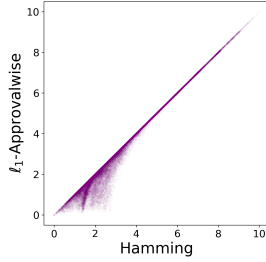


Figure 2: Correlation between Hamming and approvalwise metrics.

actually meaningful. In Section 6.3 we will show that in our datasets it is strongly correlated with the Hamming distance and, so, in the following, we focus on approvalwise distances.

5 A Grid of Approval Elections

To better understand the approvalwise metric space of elections, next we analyze expected distances between elections generated according to the (p, ϕ) -resampling model.

Fix some number m of candidates and parameters $p, \phi \in [0, 1]$, such that pm is an integer, and consider the process of generating votes from the (p, ϕ) -resampling model. In the limit, the approvalwise vector of the resulting election is:

$$\underbrace{((1-\phi) + (\phi \cdot p), \dots, (1-\phi) + (\phi \cdot p))}_{p \cdot m}, \underbrace{\phi \cdot p, \dots, \phi \cdot p}_{(1-p) \cdot m}.$$

Indeed, each of the pm candidates approved in the central ballot either stays approved (with probability $1 - \phi$) or is re-sampled (with probability ϕ , and then gets an approval with probability p). Analogous reasoning applies to the remaining $(1-p)m$ candidates. With a slight abuse of notation, we call the above vector $av(p, \phi)$. Furthermore, we refer to $av(p, 0)$ as the p -ID vector, to $av(p, 1)$ as the p -IC vector, and to 0-ID and 1-ID vectors as the *empty* and *full* ones, respectively (note that 0-ID = 0-IC and 1-ID = 1-IC).

Now, consider two additional numbers, $p', \phi' \in [0, 1]$, such that $p'm$ is an integer. Simple calculations show that:

$$\begin{aligned} d_{\text{app}}(\text{empty}, \text{full}) &= m, \\ d_{\text{app}}(p\text{-IC}, p\text{-ID}) &= 2mp(1-p), \\ d_{\text{app}}(av(p, \phi), av(p', \phi)) &= m \cdot |p - p'|, \\ d_{\text{app}}(av(p, \phi), av(p, \phi')) &= 2mp(1-p) \cdot |\phi - \phi'|. \end{aligned}$$

Thus $d_{\text{app}}(av(p, \phi), \text{empty}) = mp$ is a p fraction of the distance between *empty* and *full*, and $d_{\text{app}}(av(p, \phi), p\text{-ID}) = 2mp(1-p)\phi$ is a ϕ fraction of the distance between p -IC and p -ID (see also Figure 1). Furthermore, $d_{\text{app}}(\text{empty}, \text{full}) = m$ is the largest possible approvalwise distance.

Intuitively, (p, ϕ) -resampling elections form a grid that spans the space between extreme points of our election space; the larger the ϕ parameter, the more “chaotic” an election becomes (formally, the closer it is to the p -IC elections), and the larger the p parameter, the more approvals it contains (the closer it is to the *full* election). We use (p, ϕ) -resampling elections as a *background* dataset, which consists of 241 elections with 100 candidates and 1000 voters each, with the following p and ϕ parameters:

1. p is chosen from $\{0, 0.1, 0.2, \dots, 0.9, 1\}$ and ϕ is chosen from the interval $(0, 1)$,¹ or
2. ϕ is chosen from $\{0, 0.25, 0.5, 0.75, 1\}$ and p is chosen from the interval $(0, 1)$.

For each of these elections we compute a point in \mathbb{R}^2 , so that the Euclidean distances between these points are as similar to the approvalwise distances between the respective elections as possible. For this purpose, similarly to Szufa *et al.* [2020], we use the Fruchterman-Reingold force-directed algorithm [Fruchterman and Reingold, 1991]. For the resulting map, see the clear grid-like shape at the left side of Figure 3. Whenever we present maps of elections later in the paper, we compute them in the same way (but for datasets that include other elections in addition to the background ones).

6 Experiments

In this section, we use the map-of-elections approach to analyze quantitative properties of approval elections generated according to our models. In particular, we will see how an election’s position in the grid influences each of the properties, and what parameters to use to generate elections with the quantitative property in a desired range.

6.1 Experimental Design

Concretely, we consider the following four statistics:

Max. Approval Score. The highest approval score among all candidates in a given election, normalized by the maximum possible score, i.e., the number of voters.

Cohesiveness Level. The largest integer ℓ such that there exists an ℓ -cohesive group (for committee size 10).

Voters in Cohesive Groups. Fraction of voters that belong to at least one 1-cohesive group (for committee size 10).

PAV Runtime. Runtime (in seconds) required to compute a winning committee under the PAV rule, by solving an integer linear program provided by the `abcvoting` library [Lackner *et al.*, 2021], using the Gurobi ILP solver.

We use six datasets. Five of them are generated using our statistical cultures and consist of 100 candidates and 1000

¹By generating t elections with a parameter from interval (a, b) , we mean generating one election for each value $a + i \frac{b-a}{t+1}$, for $i \in [t]$.

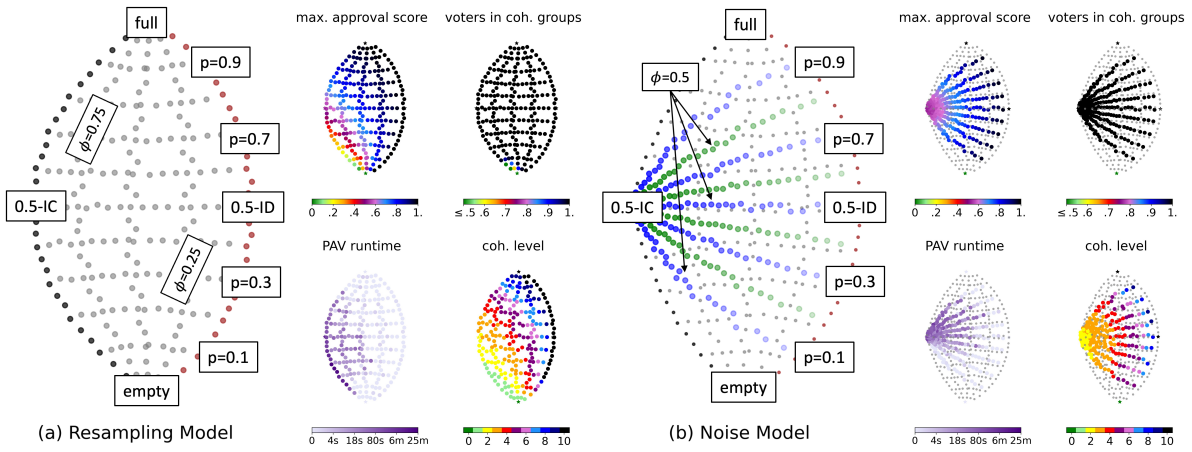


Figure 3: Maps for (a) the Resampling Model and (b) the Noise Model. For the latter, the darker a dot in the main plot, the higher is the ϕ value used to generate the election.

voters (except for the experiments related to the cohesiveness level, where we have 50 candidates and 100 voters, due to computation time). We have: 250 elections from the Disjoint Model (50 for each $g \in \{2, 3, 4, 5, 6\}$ with $\phi \in (0.05, 1/g)$); 225 elections from the Noise Model with Hamming distance (25 for each $p \in \{0.1, 0.2, \dots, 0.9\}$ with $\phi \in (0, 1)$); 225 elections from the Truncated Urn Model (25 for each $p \in \{0.1, 0.2, \dots, 0.9\}$ with $\alpha \in (0, 1)$); 200 elections from Euclidean Model (100 for 1D-Uniform, with radius in $(0.0025, 0.25)$, and 100 for 2D-Square, with radius in $(0.005, 0.5)$); these parameters are as used by Bredereck *et al.* [2019]. The sixth dataset uses real-life participatory budgeting data and contains 44 elections from Pabulib [Stolicki *et al.*, 2020], where for each (large enough) election we randomly selected a subset of 50 candidates and 1000 voters.

6.2 Experimental Results

Our visualizations are shown in Figures 3 and 4. We use the grid structure of the background dataset for comparison, with other datasets. Notably, some of them do not fill this grid. This is most striking for the real-world dataset, Pabulib (Figure 4d), which is placed very distinctly in the bottom left part.

To get an intuitive understanding of the four statistics, let us consider the background dataset in Figure 3a. We see that the highest approval score is lowest in the lower left side and increases towards up and right. This is sensible: If the average number of approved candidates increases, so does this statistic; also, if voters become more homogeneous, high-scoring candidates are likely to exist. Regarding voters in cohesive groups, it turns out that in most elections almost all voters belong to a 1-cohesive group, with the left lower part as an exception (where there are not enough approvals to form 1-cohesive groups). The time needed to find a winning committee under PAV is correlated with the distance from 0.5-IC. Similar to the highest approval score, the cohesiveness level increases when moving up or right in the diagram. Cohesive groups with levels close to the committee size only exist in very homogeneous elections (rightmost path) and elections with many approvals (top part).

We move on to the results for the five other datasets. Note that each figure also contains the background dataset (gray dots) for reference. These results help to understand the differences between our statistical cultures.

The maximum approval score statistic provides an insight into whether there is a candidate that is universally supported. Instances with a value close to 1 possess such a candidate. In a single-winner election, this candidate is likely to be a clear winner. This is undesirable when simulating contested elections or shortlisting. Also note that in the real-world data set (Pabulib) we do not observe the existence of such a candidate.

When looking at the PAV runtime, we find some statistical cultures that generate computationally difficult elections, e.g., the (p, ϕ) -resampling model with parameter values close to $p = 0.5$ and $\phi = 1$ (0.5-IC), the noise model with parameters $p \in [0.5, 0.9]$ and $\phi > 0.5$, and the disjoint model with $g = 2$. Yet, instances from the real-world dataset, as well as from the Euclidean and urn ones, can be computed very quickly.²

Concerning voters in cohesive groups, whenever this statistic is close to 1, it is easy to provide most voters with at least one approved candidate in the committee; such committees are easy to find [Aziz *et al.*, 2017]. Since many proportional rules take special care of voters that belong to cohesive groups, in such elections there are no voters that are at a systematic disadvantage. In many of our generated elections (almost) all voters belong to 1-cohesive groups, but this is not the case for the real-world, Pabulib data and it would be interesting to find statistical model that would behave similarly. For the cohesiveness level, we see that all models generate a full spectrum (i.e., $[0, 10]$) of cohesiveness levels. That said, we expect realistic elections to appear in the “lower left” part of our grid (with few approvals), and such elections tend to have low cohesiveness levels. Indeed, this is the case for Pabulib elections. Thus it is important how proportional rules treat ℓ -cohesive groups with small ℓ .

²Less than 1 second on a single core (Intel Xeon Platinum 8280 CPU @ 2.70GH) of a 224 core machine with 6TB RAM. In contrast, the worst-case instance (0.3-IC) required 25 minutes on 13 cores.

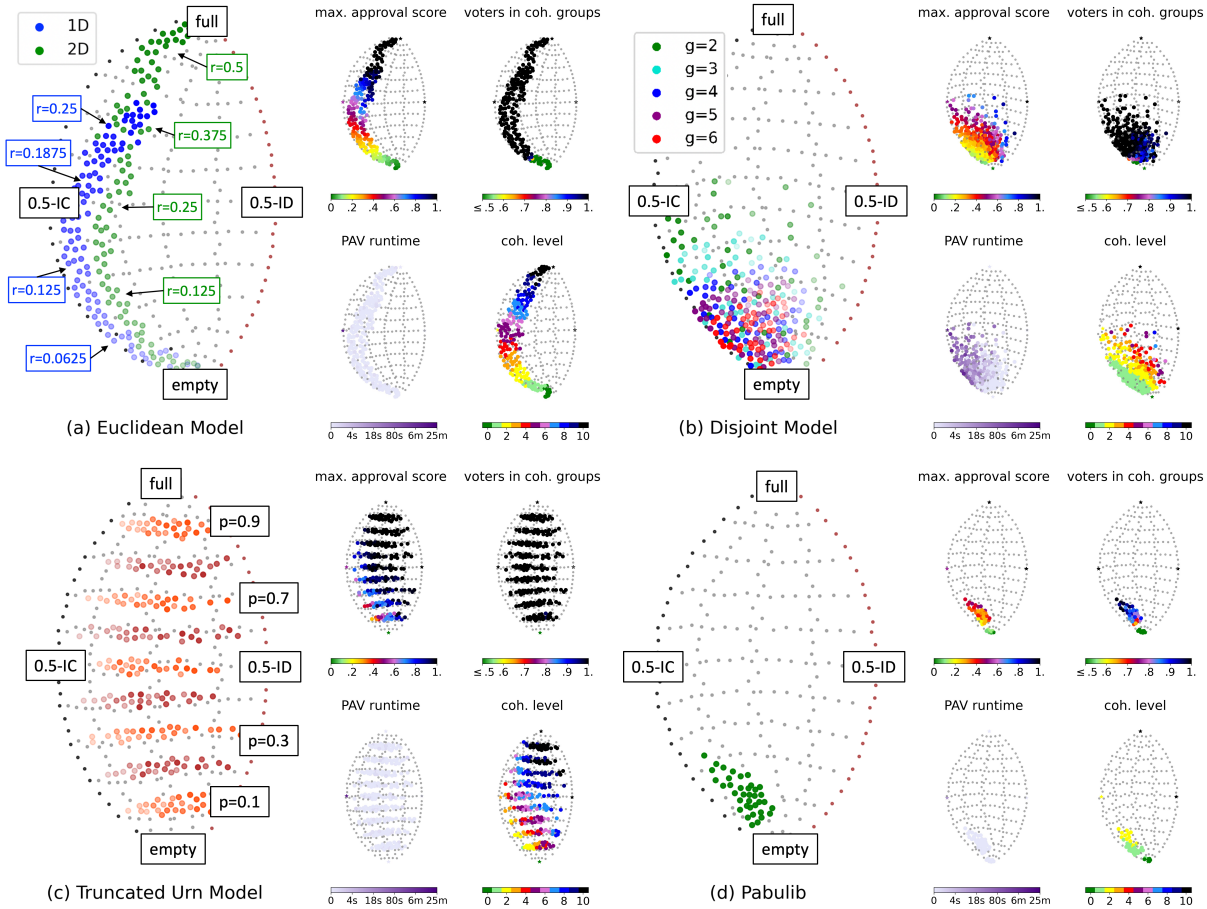


Figure 4: Maps for (a) the Euclidean Model, (b) the Disjoint Models, (c) the Truncated Urn Model, and (d) Pabulib. For (a), (b), and (c), the darker a dot in the main plot, the higher is the value of the radius, the ϕ parameter, or the α parameter, respectively.

6.3 Correlation

Figures 3 and 4 are based on the approvalwise distance, but they would not change much if we used the (computationally intractable) isomorphic Hamming distance. To verify this claim, we generated 363 elections with 10 candidates and 50 voters from the statistical cultures used in the previous experiment and compared Hamming and approvalwise distances between them: The results are in Figure 2, where each dot represents a pair of elections, and its coordinates are the respective distances between them. The Pearson Correlation Coefficient is 0.989, and for 67% pairs of elections the distances (after normalization) are identical.

This last observation requires comments. On the one hand, we observe that the two distances often coincide when the distances are large. This happens, e.g., if one election is dense (i.e., each voter approves many candidates) and the other one is sparse (i.e., each voter approves few candidates). In such cases, it is likely that every ballot $A(v_i)$ from the sparse election can be mapped to some ballot $B(u_j)$ in the dense election such that $A(i) \subseteq B(j)$. If this is the case, then both distances coincide. On the other hand, equality of the distances is unlikely, e.g., for elections generated from 0.5-IC.

7 Conclusions and Future Work

We propose to use the resampling model as a baseline in numerical experiments on approval voting. An important task for future work is to broadly study real-world datasets with the methods proposed in this paper. It would also be interesting to consider further statistics than the ones we used.

Acknowledgements

Martin Lackner was supported by the Austrian Science Fund (FWF), project P31890. Nimrod Talmon was supported by the Israel Science Foundation (ISF; GrantNo.630/19). 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

- [Alós-Ferrer and Granić, 2012] C. Alós-Ferrer and D. Granić. Two field experiments on approval voting in Germany. *Social Choice and Welfare*, 39(1):171–205, 2012.
- [Arvind *et al.*, 2012] V. Arvind, J. Köbler, S. Kuhnert, and Y. Vassudev. Approximate graph isomorphism. In *Proceedings of MFCS-2012*, pages 100–111, 2012.
- [Aziz *et al.*, 2015] H. Aziz, S. Gaspers, J. Gudmundsson, S. Mackenzie, N. Mattei, and T. Walsh. Computational aspects of multi-winner approval voting. In *Proceedings of AAMAS-2015*, pages 107–115, 2015.
- [Aziz *et al.*, 2017] H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 48(2):461–485, 2017.
- [Barrot *et al.*, 2017] N. Barrot, J. Lang, and M. Yokoo. Manipulation of Hamming-based approval voting for multiple referenda and committee elections. In *Proceedings of AAMAS-2017*, pages 597–605, 2017.
- [Baujard and Igersheim, 2011] A. Baujard and H. Igersheim. Framed-field experiment on approval voting and evaluation voting. Some teachings to reform the French presidential electoral system. In B. Dolez, B. Grofman, and A. Laurent, editors, *In Situ and Laboratory Experiments on Electoral Law Reform*, Studies in Public Choice, pages 69–89. Springer, 2011.
- [Berg, 1985] S. Berg. Paradox of voting under an urn model: The effect of homogeneity. *Public Choice*, 47(2):377–387, 1985.
- [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*, 2021.
- [Bouveret *et al.*, 2019] S. Bouveret, R. Blanch, A. Baujard, F. Durand, H. Igersheim, J. Lang, A. Laruelle, J. Laslier, I. Lebon, and V. Merlin. Voter autrement 2017 for the French presidential election - the data of the in situ experiment. Technical report, Zenodo, 2019. <http://doi.org/10.5281/zenodo.3548574>.
- [Brams and Fishburn, 1983] S. Brams and P. Fishburn. *Approval Voting*. Birkhäuser, Boston, 1983.
- [Bredereck *et al.*, 2019] R. Bredereck, P. Faliszewski, A. Kaczmarczyk, and R. Niedermeier. An experimental view on committees providing justified representation. In *Proceedings of IJCAI-2019*, pages 109–115, 2019.
- [Caragiannis *et al.*, 2022] I. Caragiannis, C. Kaklamanis, N. Karanikolas, and G. A. Krimpas. Evaluating approval-based multiwinner voting in terms of robustness to noise. *Autonomous Agents and Multiagent Systems*, 36(1):1–22, 2022.
- [Dudycz *et al.*, 2020] S. Dudycz, P. Manurangsi, J. Marcinkowski, and K. Sornat. Tight approximation for proportional approval voting. In *Proceedings of IJCAI-2020*, pages 276–282, 2020.
- [Elkind and Lackner, 2015] E. Elkind and M. Lackner. Structure in dichotomous preferences. In *Proceedings of IJCAI-2015*, pages 2019–2025, 2015.
- [Elkind *et al.*, 2017] E. Elkind, P. Faliszewski, J. Laslier, P. Skowron, A. Slinko, and N. Talmon. What do multiwinner voting rules do? An experiment over the two-dimensional Euclidean domain. In *Proceedings of AAAI-2017*, pages 494–501, 2017.
- [Enelow and Hinich, 1984] J. Enelow and M. Hinich. *The Spatial Theory of Voting: An Introduction*. Cambridge University Press, 1984.
- [Enelow and Hinich, 1990] J. Enelow and M. Hinich. *Advances in the Spatial Theory of Voting*. Cambridge University Press, 1990.
- [Faliszewski *et al.*, 2019] P. Faliszewski, P. Skowron, A. Slinko, S. Szufa, and N. Talmon. How similar are two elections? In *Proceedings of AAAI-2019*, pages 1909–1916, 2019.
- [Faliszewski *et al.*, 2020] P. Faliszewski, A. Slinko, and N. Talmon. Multiwinner rules with variable number of winners. In *Proceedings of ECAI-2020*, pages 67–74, 2020.
- [Fruchterman and Reingold, 1991] T. Fruchterman and E. Reingold. Graph drawing by force-directed placement. *Software: Practice and Experience*, 21(11):1129–1164, 1991.
- [Godziszewski *et al.*, 2021] M. Godziszewski, P. Batko, P. Skowron, and P. Faliszewski. An analysis of approval-based committee rules for 2D-Euclidean elections. In *Proceedings of AAAI-2021*, pages 5448–5455, 2021.
- [Grohe *et al.*, 2018] M. Grohe, G. Rattan, and G. Woeginger. Graph similarity and approximate isomorphism. In *Proceedings of MFCS-2018*, pages 20:1–20:16, 2018.
- [Lackner and Maly, 2021] M. Lackner and J. Maly. Approval-based shortlisting. In *Proceedings of AAMAS-2021*, pages 737–745, 2021.
- [Lackner and Skowron, 2020] M. Lackner and P. Skowron. Utilitarian welfare and representation guarantees of approval-based multiwinner rules. *Artificial Intelligence*, 288:103366, 2020.
- [Lackner and Skowron, 2021] M. Lackner and P. Skowron. Multiwinner voting with approval preferences. *arXiv preprint arXiv:2007.01795*, 2021.
- [Lackner *et al.*, 2021] M. Lackner, P. Regner, B. Krenn, and S. S. Forster. abc voting: A Python library of approval-based committee voting rules, 2021. Current version: <https://github.com/martinlackner/abcvoting>.
- [Laslier and der Straeten, 2008] J. Laslier and K. Van der Straeten. A live experiment on approval voting. *Experimental Economics*, 11(1):97–105, 2008.
- [Laslier and Sanver, 2010] J. Laslier and R. Sanver, editors. *Handbook on Approval Voting*. Springer, 2010.
- [Peters and Lackner, 2020] D. Peters and M. Lackner. Preferences single-peaked on a circle. *Journal of Artificial Intelligence Research*, 68:463–502, 2020.
- [Sánchez-Fernández *et al.*, 2017] L. Sánchez-Fernández, E. Elkind, M. Lackner, N. Fernández, J. A. Fisteus, P. Basanta Val, and P. Skowron. Proportional justified representation. In *Proceedings of AAAI-2017*, pages 670–676, 2017.
- [Skowron *et al.*, 2016] P. Skowron, P. Faliszewski, and J. Lang. Finding a collective set of items: From proportional multirepresentation to group recommendation. *Artificial Intelligence*, 241:191–216, 2016.
- [Skowron, 2021] P. Skowron. Proportionality degree of multiwinner rules. In *Proceedings of EC-2021*, pages 820–840, 2021.
- [Stolicki *et al.*, 2020] D. Stolicki, S. Szufa, and N. Talmon. Pabulib: A participatory budgeting library. *arXiv preprint arXiv:2012.06539*, 2020.
- [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.