

Optimization-Based Voting Rule Design: The Closer to Utopia the Better

Piotr Faliszewski
AGH University
Krakow, Poland
faliszew@agh.edu.pl

Stanisław Szufa
Jagiellonian University
Krakow, Poland
s.szufa@gmail.com

Nimrod Talmon
Ben-Gurion University of the Negev
Be'er Sheva, Israel
talmonn@bgu.ac.il

ABSTRACT

In certain situations, such as elections in the Euclidean domain, it is possible to specify clear requirements for the operation of a multiwinner voting rule, for it to provide committees that correspond to some desirable intuitive notions (such as individual excellence of the committee members or their diversity). We formally describe several such requirements, which we refer to as “utopias”. Supplied with such utopias, we develop an optimization-based mechanism for constructing committee scoring rules that provide results as close to these utopias as possible; we test our mechanism on weakly separable and OWA-based rules. In particular, using our method we recover some connections between known multiwinner voting rules and certain applications.

KEYWORDS

multiwinner voting; committee scoring rules; optimization

ACM Reference Format:

Piotr Faliszewski, Stanisław Szufa, and Nimrod Talmon. 2018. Optimization-Based Voting Rule Design: The Closer to Utopia the Better. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018)*, Stockholm, Sweden, July 10–15, 2018, IFAAMAS, 9 pages.

1 INTRODUCTION

Multiwinner voting is a formalism for selecting a set of items (a committee), based on the preferences of a group of agents (the voters) [12, 20, 21]. For example, a group of judges may need to select a set of finalists of a competition, a hiring committee may need to select a set of people to invite for on-site interviews, and an Internet store may need to decide which items to present on its homepage (depending on how the preferences of its customers are perceived). In each of these examples, we need committees with different properties; the judges should select individually best candidates, the Internet store should select a diverse set of items that covers interests of as many of its customers as possible, and the hiring committee should, in a certain sense, balance these two requirements (on the one hand, we certainly wish to invite as good candidates as possible, but we also wish to maintain some diversity; e.g., a computer science department is likely to interview candidates working in several areas, and not only in the “hottest” one).

More generally, following the recent overview of Faliszewski et al. [20], multiwinner elections might be categorized into three classes, based on what their goals are: *Individual Excellence*, for selecting individually best candidates (as in the judges example

above); *Proportional Representation*, for accurately and proportionally representing the electorate views; and *Diversity*, for reflecting a wide spectrum of voters’ views (as in the Internet store above).

So far, to address these varied goals and needs researchers typically analyzed existing multiwinner voting rules, studied their computational complexity [4, 5, 8, 24, 30, 32], analyzed their axiomatic properties [1, 12, 17, 31, 34], evaluated them experimentally [6, 10, 11, 19, 35], and—based on this evidence—argued which rules are best for which application (e.g., the k -Borda rule [9] is seen as appropriate for choosing individually excellent candidates, whereas the Chamberlin–Courant rule [7] is appropriate for identifying diverse committees that cover a wide spectrum of opinions; the references provided here are meant to present the wide range of results obtained, and are certainly not complete). In other words, typically, researchers analyzed existing rules and checked which ones behave appropriately for a given setting. There are also cases where researchers hand-designed multiwinner rules to achieve their goals (e.g., for situations such as the hiring committee above, Faliszewski et al. [19] designed a spectrum of rules achieving various levels of compromise between the goals of excellence and diversity; Elkind et al. [14], Aziz et al. [2] and Sekar et al. [32] proposed multiwinner variants of the Condorcet rule).

In this paper we take a radically different approach from the previous ones: Given a specification of the kind of committees one is interested in, we use an optimization algorithm to automatically design—in a principled way—rules that match this specification. (Related to our approach, we mention the position paper of Xia [37], which takes a normative approach and suggests the use of machine learning to automatically design voting rules.)

Motivation. Our work is driven by two main motivations. The first one, suggested above, is that we wish to develop a methodology for designing voting rules which would satisfy certain desired properties. There are many multiwinner rules (such as k -Borda [9], Bloc, Chamberlin–Courant [7], Proportional Approval Voting [36], Monroe [26], and many others) that seem to have good properties for *some* idealized goals (recall the three types of multiwinner elections discussed above), but here we wish to develop a general mechanism that, when supplied with *any* arbitrary goal (specified in an appropriate way), can output a multiwinner rule appropriate for this goal. Indeed, designing voting rules which satisfy certain properties is in the heart of social choice. As a proof of concept, in this paper we focus on (two subclasses of) the class of committee scoring rules [12] and design an algorithm that searches for an appropriate rule among them. We focus on the rules from these classes because they are parameterized through sets of numeric values that we can tweak to manipulate their properties; this aspect is important for our optimization-based approach.

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden. © 2018 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

Our second motivation relates to the richness of the class of committee scoring rules. So far, researchers have analyzed the general structure of committee scoring rules [34], considered a few of their subclasses [12, 17, 18], and studied several specific rules [1, 12, 33] or spectra of rules [19]. However, as there are so many committee scoring rules, it might be that, in spite of the effort outlined above, some important rules might have been missed. Our mechanism of designing rules tailored for particular goals explores specified subclasses of committee scoring rules and, for each setting, either finds one of the already-known rules (thus confirming that it is highly appropriate for a given setting) or discovers a new rule.

Methodology. Providing intuitive descriptions of what kinds of committees are appropriate for certain types of multiwinner elections is quite easy and, indeed, above we only needed a few words to provide several examples. Providing formal specifications is far less trivial because there are many ways to understand intuitions and many ways to implement them. Here we take an approach inspired by a recent work of Elkind et al. [11] and consider Euclidean elections, in which it is relatively easy to interpret election results; this allows us to formally specify the goals of our rules.

In the Euclidean election model we assume that each candidate and each voter is represented by an ideal point in a given Euclidean space. The closer a candidate’s point is to a given voter’s point, the more this voter likes this candidate [15, 16] (e.g., in the context of politics, positions on a 2-dimensional plane may indicate the level of a person’s belief in personal and economic freedom).

To design a new voting rule, we proceed as follows. First, we assume some distribution of ideal points of the candidates and the voters (we use the uniform distribution on a 1D interval). Second, we specify the desired distribution of committee members; we refer to this distribution as the *utopic distribution* (e.g., as suggested by the experiments of Elkind et al. [11], for committees with individually excellent candidates we use the distribution that puts all the weight on the single center point, whereas for the case of diverse committees, we use the uniform distribution over the whole interval). Third, for a given multiwinner voting rule we evaluate how closely its results follow the utopic distribution (using a measure inspired by the Earth mover’s distance). We perform this evaluation on rules from a given class, and seek one whose distance to the utopic distribution is the smallest. Specifically, to identify the rule which is “the closest to Utopia”, we use an optimization algorithm. (Our work is close in spirit to that on distance rationalizations of voting rules [13, 23, 25, 27], but completely different technically.)

Search Space of Voting Rules. We focus on two subclasses of committee scoring rules, the class of weakly separable rules and the class of OWA-based rules (within which we consider Borda-based rules only). Briefly put, a weakly separable rule is defined by a vector of points that the voters assign to the candidates, depending on how highly they rank them. We add up the points that the candidates receive from the voters and the candidates with the highest scores form the winning committee. This class includes, for example, the well-known k -Borda, SNTV, and Bloc rules.

OWA-based rules are defined by vectors of points associated with how highly a voter ranks a given candidate (as in the previous case) and the ordered-weighted average operators (the OWA vectors), which specify the importance of the committee members,

depending on their ranking by a voter. OWA-based rules include, e.g., the Chamberlin–Courant rule and the PAV rule (see, e.g., the works of Skowron et al [33] and Aziz et al. [3]). Indeed, together the classes of weakly separable rules and OWA-based rules cover a very large fraction of the committee scoring rules analyzed to date.

Our Contribution. We believe that our most important contribution is a proof of concept: We show that it is possible to use optimization algorithms that minimize the distance from utopic distributions as a tool for designing rules with given properties. This is witnessed by the following results, which confirm previous studies and intuitions:

- (1) When we look for a weakly separable rule that selects candidates as close to the center of our Euclidean distribution as possible, our algorithm, in essence, finds the k -Borda rule. Similarly, when we look for a rule closest to a utopic distribution that models the behavior of the Bloc rule, we indeed find a very similar rule.
- (2) When we look for an OWA-based rule whose winners are spread as uniformly over our Euclidean distribution as possible, we recover the Chamberlin–Courant rule (as expected). As we change the utopic distribution to be more and more focused on the center of the interval (corresponding to rules achieving compromises between individual excellence of committee members and their diversity), we discover rules similar to those analyzed for this scenario by Faliszewski et al. [19]; interestingly, our algorithm also discovers some new, intuitively appealing, rules for this scenario.

Besides confirming good performance of already observed rules, our algorithm demonstrates that, among weakly separable rules, SNTV is the best diversity-oriented rule. Since SNTV is believed not to perform very well for this task (because of its focus on top preferences), the fact that it is the best our algorithm can find suggests that there are simply no truly good diversity-oriented rules among weakly separable ones. While this is somewhat disappointing—we hoped to discover a diversity-oriented weakly separable rule (and, hence, polynomial-time computable) that could be used instead of the Chamberlin–Courant rule (which is NP-hard [5, 24, 30])—our results indicate that we should look for polynomial-time diversity-oriented rules someplace else.

2 PRELIMINARIES

An election $E = (C, V)$ consists of a set of candidates $C = \{c_1, \dots, c_m\}$ and a collection of voters $V = (v_1, \dots, v_n)$, where each voter v_i has a linear order \succ_{v_i} , ranking the candidates from the one that v_i appreciates most to the one that v_i appreciates least. We refer to \succ_{v_i} as the preference order of voter v_i (and, sometimes, as the vote of v_i). For a voter v and a candidate c , we write $\text{pos}_v(c)$ to denote the position of c in v ’s preference order (the top-ranked candidate has position 1, the next one has position 2, and so on). A multiwinner voting rule is a function \mathcal{R} that, given an election $E = (C, V)$ and an integer k , $1 \leq k \leq |C|$, outputs a family of size- k subsets of C (i.e., a family of committees) that tie as winners of this election.

For each integer t , we write $[t]$ to denote the set $\{1, \dots, t\}$. In particular, if m is the number of candidates, we often interpret the set $[m]$ as the set of positions that candidates may take in a

preference order. A single-winner scoring function (for an election with m candidates) is a non-increasing function $\gamma_m : [m] \rightarrow \mathbb{R}$ that associates each position in a vote with a score value. We define the γ_m -score of a candidate c in an election $E = (C, V)$ to be γ -score $_E(c) = \sum_{v \in V} \gamma_m(\text{pos}_v(c))$. We use normalized scoring functions, so that $\gamma_m(1) = 1$ and $\gamma_m(0) = 0$. For example, the Borda scoring function (denoted β_m), is defined as $\beta_m(i) = m-i/m-1$, and the t -Approval scoring function (denoted α_t , where $t \in [m]$ is a parameter) is a function that associates score 1 with the first t positions, and score 0 with the remaining ones (note that the subscript t does not refer to the number of candidates in the election).

Committee scoring functions are defined analogously to the single-winner ones, but for a generalized notion of a position. Let us fix committee size k . Then, given a committee S and a vote v , we define the position of S in v , denoted $\text{pos}_v(S)$, to be the sequence of positions of the members of S in v , sorted in the increasing order (i.e., we obtain $\text{pos}_v(S)$ by sorting the set $\{\text{pos}_v(s) \mid s \in S\}$ in the increasing order). We write $[m]_k$ to denote the set of all length- k increasing sequences of elements from $[m]$ (and we interpret elements of $[m]_k$ as committee positions). We say that committee position $I = (i_1, \dots, i_k)$ weakly dominates committee position $J = (j_1, \dots, j_k)$, denoted $I \geq J$, if for each $t \in [k]$ it holds that $i_t \leq j_t$. A committee scoring function (for m candidates and committee size k) is a function $f_{m,k} : [m]_k \rightarrow \mathbb{R}$, such that for each two committee positions $I, J \in [m]_k$, if $I \geq J$ then $f(I) \geq f(J)$. The $f_{m,k}$ -score of committee S in election $E = (C, V)$ is defined as $\sum_{v \in V} f_{m,k}(\text{pos}_v(S))$. For a family $f = (f_{m,k})_{k \leq m}$ of committee scoring functions (one for each number of candidates and committee size), we define the committee scoring rule \mathcal{R}_f as follows: Given an election $E = (C, V)$ with m candidates and committee size k , it outputs all size- k committees S with the highest $f_{m,k}$ -score.

Example 2.1. Let us fix an election E with m candidates and committee size k . The SNTV rule is defined by committee scoring functions of the form $f_{m,k}^{\text{sntv}}(i_1, \dots, i_k) = \alpha_1(i_1)$. This means that the rule selects a committee of k candidates that are ranked first most frequently (or several such committees, in case of ties). The Bloc rule uses functions of the form $f_{m,k}^{\text{bloc}}(i_1, \dots, i_k) = \alpha_k(i_1) + \dots + \alpha_k(i_k)$, which can be interpreted as saying that each voter gives one point to each of his or her k most favorite candidates, and the k candidates with the highest score form the winning committee. The k -Borda rule chooses k candidates with the highest Borda scores and is defined through the functions $f_{m,k}^{\text{kb}}(i_1, \dots, i_k) = \beta_m(i_1) + \dots + \beta_m(i_k)$. The Chamberlin–Courant rule (the CC rule) uses scoring functions of the form $f_{m,k}^{\text{cc}}(i_1, \dots, i_k) = \beta_m(i_1)$. This means that given a committee S , each voter associates it with the Borda score of this member of S that he or she ranks highest (this candidate is called the representative of the voter). Finally, the Harmonic-Borda rule [19] (the HB rule) uses the scoring function $f_{m,k}^{\text{hb}}(i_1, \dots, i_k) = \beta_m(i_1) + 1/2\beta_m(i_2) + \dots + 1/k\beta_m(i_k)$.

Consider a setting with m candidates, where the desired committee size is k , and where \mathcal{R} is a committee scoring rule:

- (1) We say that \mathcal{R} is *weakly separable* if its committee scoring function is of the form $f(i_1, \dots, i_k) = \gamma(i_1) + \dots + \gamma(i_k)$, where γ is a single-winner scoring function.

- (2) We say that \mathcal{R} is *OWA-based* if its committee scoring function is of the form $f(i_1, \dots, i_k) = \lambda_1\gamma(i_1) + \dots + \lambda_k\gamma(i_k)$, where $\Lambda = (\lambda_1, \dots, \lambda_k)$ is a sequence of non-negative real numbers and γ is a single-winner scoring function (we refer to the vector Λ as the OWA vector).

Note that every weakly separable rule is OWA-based (with the all 1s OWA vector). For the purpose of this paper, we normalize OWA vectors, so that λ_1 is always 1. We say that a committee scoring rule is *OWA/Borda-based* if it is OWA-based, uses the Borda scoring function and a non-increasing OWA-vector.

Example 2.2. The k -Borda and Bloc rules are both weakly separable (and k -Borda is also OWA/Borda-based). The CC and HB rules are OWA/Borda-based (using OWA vectors $(1, 0, \dots, 0)$ and $(1, 1/2, \dots, 1/k)$, respectively). The SNTV rule is both OWA-based and weakly separable, but it is not OWA/Borda-based.

3 METHODOLOGY

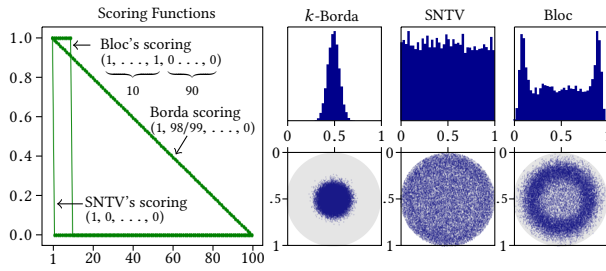
In this section we describe our technique of designing voting rules that match given utopic distributions. First we describe the model of Euclidean elections and the corresponding way of presenting election results, then we present our utopic distributions and our measures of closeness to these distributions, and finally, we describe our algorithm(s) for finding rules that match given utopias. In Section 4 we apply this methodology to design specific rules.

3.1 Euclidean Elections

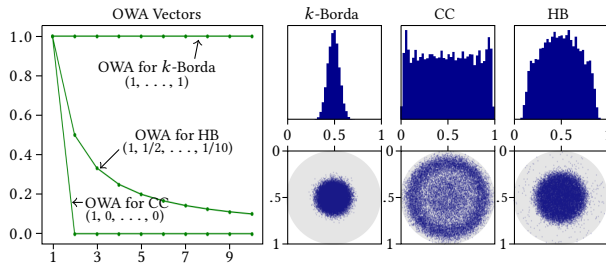
In the t -dimensional Euclidean model of elections, each individual u (i.e., each candidate and each voter) is represented by a point $p(u) \in \mathbb{R}^t$ in the t -dimensional space. Intuitively, the coordinates of this point may correspond to u 's position regarding some t issues. E.g., in the 1-dimensional model this single issue may be the acceptable level of taxation, while in the 2-dimensional model the two issues may indicate the levels of belief in personal and economic freedom [15, 16]. Each voter forms his or her preference order by sorting the candidates in increasing order of the distances of the candidates' ideal points from the voter's ideal point (i.e., the closer a candidate is to a voter, the higher the voter ranks the candidate).

In our computations, we use either 1-dimensional Euclidean elections, where we generate the ideal points of candidates and voters by drawing them uniformly at random from the $[0, 1]$ interval, or 2-dimensional elections, where we draw the ideal points uniformly at random from a disc centered at point $(0.5, 0.5)$ with radius 0.5. We refer to the former as the *interval model* and to the latter as the *disc model*. We always generate elections with 100 candidates and 100 voters, and we seek committees of size 10. We chose these parameters to ensure that our results are comparable to those already present in the literature [11, 19].

We use the interval model throughout the whole process of designing voting rules, and we use the disc model to check whether the rules that we produce maintain their features after changing (and, in a sense, generalizing) the setting. Arguably, the disc model is very similar to the interval one, and the reason for choosing it was that, on the one hand, we wanted to test our rules in a (somewhat) different setting, but, on the other hand, we wanted this setting to be sufficiently similar, so we could maintain the same intuitions while observing how the behavior of our rules changes.



(a) Weakly separable rules and their scoring functions.



(b) OWA/Borda-based rules and their OWA vectors.

Figure 1: Visualization of election results under various rules. Histograms for the interval model are on top and scatter plots for the disc model are on the bottom. To the left, we show scoring-functions/OWA-vectors defining the rules.

Following Elkind et al. [11], we present the results of our elections visually. For a given voting rule \mathcal{R} and a given election model (interval or disc), we generate a number of elections according to the model, compute the \mathcal{R} winning committee for each election (if there are ties, we break them arbitrarily), and—depending on the model—present them as follows:

- (1) For the interval model, we partition the $[0, 1]$ interval into subintervals, count how many times a candidate from a given subinterval was in a winning committee, and present these numbers as a histogram. Note that we do not normalize the histograms; different ones have different scales (the point of these histograms is to show the “shape” of the election results).
- (2) For the disc model, we show a scatter plot, where each member of a winning committee is indicated as a blue dot (thus, as opposed to the work of Elkind et al. [11], our plots for the disc model are *not* histograms; this way fewer election results suffice to obtain a meaningful figure). In addition to the blue dots, we also show the gray disc from which the candidates’ and voters’ points are drawn.

As an example—and for future comparisons—in In Figure 1 we show visualizations of the results of k -Borda, SNTV, Bloc, CC, and Harmonic-Borda (HB). All 1D plots are based on 1000 interval elections and all 2D plots are based on 2000 disc elections; for 1D histograms we use 40 subintervals. For weakly separable rules, we show plots of their scoring functions (so on the x -axis we have the

100 possible positions in a vote), and for OWA/Borda-based rules, we show their OWA vectors (so on the x -axis we have 10 entries); k -Borda is shown in both plots as it belongs to both classes of rules.

To compute results of weakly separable rules, we use their direct polynomial-time algorithms. For OWA/Borda-based rules, we compute winning committees by solving integer linear programs (ILPs) provided for this task by Peters [29] (we use the CPLEX ILP solver). Peters showed that using his formulations gives a polynomial-time algorithm for the case of single-peaked elections; since elections generated for the interval model are single-peaked, we enjoy this guaranteed efficiency (but this no longer holds for the disc model).

3.2 Utopic Distributions and Distance Measures

We model the requirements regarding our rules as *utopic distributions*, that is, as probability distributions over the $[0, 1]$ interval that represent how, ideally, we would like the winners of our interval elections to be distributed (or, roughly speaking, how we would like their visualization to look like). For example, the utopic distribution that models the goal of individual excellence associates the whole probability mass with the center of the interval, whereas the distribution associated with covering the whole spectrum is, simply, the uniform distribution over the interval. (We also use several other utopic distributions, which we describe later.)

Let \mathcal{U} be some utopic distribution. Given a committee $W = \{w_1, \dots, w_k\}$ for some interval election, we define d_W , the distribution associated with W , so that for each $x \in [0, 1]$:

$$d_W(x) = \frac{\|\{w_i \mid p(w_i) = x\}\|}{k}.$$

To measure how closely W fits utopia \mathcal{U} , we use the intuitions underlying the Earth mover’s distance [28]: We view the probability mass associated with each point (each interval) as the number of “grains of sand” that lie on this point (this interval). Moving a grain of sand from point x to point y costs $|x - y|$. The distance between two distributions is the lowest possible cost of moving the “grains of sand” needed to transform one of them into the other. While this intuition is discrete in its nature, our utopic distributions are sometimes continuous (in other words, sometimes we consider probability density functions). Instead of providing a general definition of our distance, below we describe the utopic distributions that we consider and for each we derive the appropriate distance measure based on the above intuition.¹

In the descriptions below, we let k be the committee size and $W = \{w_1, \dots, w_k\}$ be a committee, whose members have ideal points $p(w_1), \dots, p(w_k) \in [0, 1]$. We assume that these points are sorted, i.e., $p(w_1) \leq p(w_2) \leq \dots \leq p(w_k)$.

Individual Excellence (\mathcal{U}_{IE}). The individual excellence utopic distribution, \mathcal{U}_{IE} , is defined as concentrating all the probability mass in the center of the interval, at point 0.5; see Figure 2a (this is inspired by the k -Borda rule, which is regarded as very good for the excellence goal, and which chooses candidates in the center).

¹The reader can rightfully complain that we wanted to deal with arbitrary goal specifications and not only specific ones. We provide two answers. First, our reasoning, in essence, applies to all utopic distributions and what we do can be seen as deriving closed form formulas, to speed up computations. Second, instead of following our method, given a utopic distribution one may generate discrete histograms for it and for d_W and directly use the Earth mover’s distance to measure their similarity.

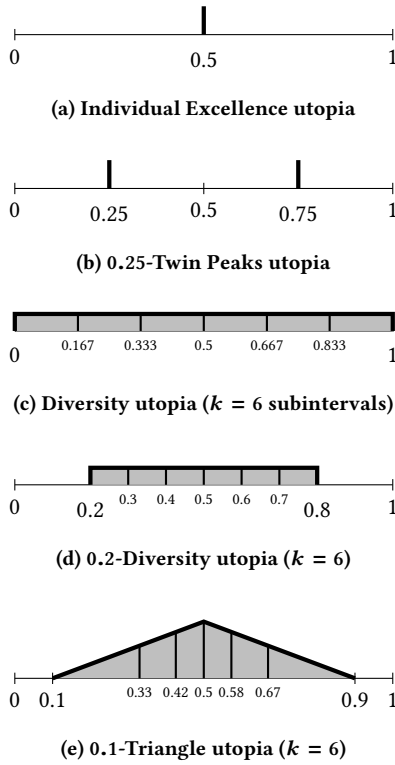


Figure 2: Several utopic distributions.

We define the distance between \mathcal{U}_{IE} and d_W to be $\text{EMD}(\mathcal{U}, d_W) = \sum_{i=1}^k 1/k |p(w_i) - 0.5|$. That is, for each member of the committee we pay the cost of moving him or her to the center of the interval (we multiply each $|p(w_i) - 0.5|$ by $1/k$ as each member of the committee is associated with probability mass $1/k$).

Twin Peaks ($\mathcal{U}_{TP}^\epsilon$). The Bloc rule motivates the study of the twin peaks utopic distributions (see Figures 1 and 2b). An ϵ -twin peaks utopic distribution for parameter ϵ , denoted $\mathcal{U}_{TP}^\epsilon$, places half of the probability mass on point ϵ and half on point $1 - \epsilon$; see Figure 2b. We let the distance between $\mathcal{U}_{TP}^\epsilon$ and d_W be:

$$\text{EMD}(\mathcal{U}_{TP}^\epsilon, d_W) = \sum_{i=1}^{k/2} 1/k |p(w_i) - \epsilon| + \sum_{i=k/2+1}^k 1/k |p(w_i) - (1 - \epsilon)|$$

(we assume that k is even; recall that the ideal points of the committee members are sorted); we assign the left half of committee members to the left peak, and the right half to the right peak.

Diversity (\mathcal{U}_D). The diversity (or, coverage) utopic distribution, denoted as \mathcal{U}_D and defined to be the uniform distribution over $[0, 1]$, models the idea that a diverse committee should cover the whole interval as uniformly as possible. Our reasoning for the distance $\text{EMD}(\mathcal{U}_D, d_W)$ is that the committee members are supposed to be distributed evenly along the interval $[0, 1]$ and, so, each of them is responsible for covering a $1/k$ -length subinterval. We assign the subintervals to committee members so that w_1 is assigned to $[0, \frac{1}{k}]$, w_2 is assigned to $[\frac{1}{k}, \frac{2}{k}]$ and so on; see Figure 2c.




Figure 3: The cases for the cost of $\text{EMD}(\mathcal{U}_D, d_W)$. The cost associated with committee member w_i is the gray area.

Let $\ell = 1/k$ be the length of the subintervals. For each w_i , $i \in [k]$, we define the cost of “spreading” his or her probability mass from d_W over the assigned subinterval $[\ell(i-1), \ell i]$ as follows:

- (1) If the committee member is to the left of his or her interval (i.e., $p(w_i) < \ell(i-1)$), then we need to pay the cost $(\ell(i-1) - p(w_i))\ell$ for moving his or her probability mass (which also is equal to ℓ) to the point $\ell(i-1)$, and then the cost $1/2\ell^2 = \int_{\ell(i-1)}^{\ell} (1 - \frac{x - \ell(i-1)}{\ell}) dx$ for “spreading” his or her weight over the interval (note that this latter cost equals to the area of a triangle; see Figure 3).
- (2) If the committee member is to the right of his or her interval (i.e., $p(w_i) > \ell i$), then we proceed analogously: Moving his or her probability weight to the point ℓi costs $\ell(p(w_i) - \ell i)$ and “spreading” this weight over the interval costs $1/2\ell^2$.
- (3) If the committee member is in his or her interval (i.e., $\ell(i-1) \leq p(w_i) \leq \ell i$), then it suffices to “spread” the $\frac{p(w_i) - \ell(i-1)}{\ell}$ fraction of his or her probability mass to the part of the interval left of him, at cost $\frac{1}{2}(p(w_i) - \ell(i-1))^2$ (analogously to the previous cases, this can be expressed as the area of a right triangle, with two sides of length $p(w_i) - \ell(i-1)$; see Figure 3), and the remaining mass, to the part of the interval to the right of him or her, at cost $\frac{1}{2}(\ell i - p(w_i))^2$.

Overall, the cost associated with w_i is:

$$\text{cost}(w_i) = \begin{cases} (\ell(i-1) - p(w_i))\ell + \frac{\ell^2}{2} & \text{if } p(w_i) \leq \ell(i-1) \\ \frac{(p(w_i) - \ell(i-1))^2 + (\ell i - p(w_i))^2}{2} & \text{if } \ell(i-1) < p(w_i) < \ell i \\ (p(w_i) - \ell i)\ell + \frac{\ell^2}{2} & \text{if } \ell i \leq p(w_i) \end{cases}$$

and we define $\text{EMD}(\mathcal{U}_D, d_W)$ to be $\sum_{i=1}^k \text{cost}(w_i)$.

Diversity/Excellence Compromises (\mathcal{U}_D^ϵ and \mathcal{U}_T^ϵ). We also consider two families of utopic distributions that achieve a certain level of compromise between the ideals of individual excellence and diversity. Let ϵ be a number in $[0, 0.5]$. We define the ϵ -diversity utopic distribution, denoted by \mathcal{U}_D^ϵ , to be the uniform distribution over the interval $[\epsilon, 1 - \epsilon]$; see Figure 2d.

Each committee member w_i is responsible for covering interval $I_i = [\epsilon + \ell(i-1), \epsilon + \ell i]$ of length $\ell = 1/k(1 - 2\epsilon)$. Using the same reasoning as for \mathcal{U}_D , we define $\text{EMD}(\mathcal{U}_D, d_W^\epsilon)$ to be $\sum_{i=1}^k \text{cost}(w_i)$, where:

$$\text{cost}(w_i) = \begin{cases} (\epsilon + \ell(i-1) - p(w_i))\frac{1}{k} + \frac{\ell}{2k} & \text{for } p(w_i) \leq \epsilon + \ell(i-1) \\ \frac{(p(w_i) - \epsilon - \ell(i-1))^2 + (\epsilon + \ell i - p(w_i))^2}{2k\ell} & \text{for } p(w_i) \in I_i \\ (p(w_i) - \epsilon - \ell i)\frac{1}{k} + \frac{\ell}{2k} & \text{for } \epsilon + \ell i \leq p(w_i) \end{cases}$$

Our second way of capturing a compromise between individual excellence and diversity is via a distribution whose probability density function, for a given $\epsilon \in [0, 0.5]$, is a triangle with a peak at

0.5, set over the interval $[\epsilon, 1 - \epsilon]$ (the area of the interval is always one). We call it the ϵ -triangle utopic distribution and denote it by \mathcal{U}_T^ϵ . Again, each committee member is associated with an interval for which the probability mass equals exactly to $1/k$; see Figure 2e. We derive the value $\text{EMD}(\mathcal{U}_T^\epsilon, d_W)$ following the same logic as in the previous cases (we omit the calculations due to limited space).

3.3 Search Algorithms

The final component of our method is an algorithm that, given a utopic distribution \mathcal{U} and one of our two families of committee scoring rules, finds a rule \mathcal{R} as close to \mathcal{U} as possible.

General Setup. For m candidates, a weakly separable rule is defined via a non-increasing vector $X = (x_1, \dots, x_m)$, such that $x_1 = 1$ and $x_m = 0$; the vector X specifies the values of the underlying single-winner scoring function γ for the possible positions in a preference order. Given a committee size k , an OWA/Borda-based rule is defined by its non-increasing OWA vector $(\lambda_1, \dots, \lambda_k)$, where $\lambda_1 = 1$. Correspondingly, given a vector V of appropriate size, we write \mathcal{R}_V to denote the rule defined by this vector (when we consider weakly separable rules, V gives the score values; when we consider OWA/Borda-based rules, it is the OWA vector). Given a vector V' , by *normalizing* it we mean sorting it, setting its first coordinate to 1, replacing all > 1 values with 1s and all < 0 values with 0, and—for weakly separable rules—setting its last coordinate to 0 (so that the vector describes a legal rule from the relevant class).

Let us fix the class of rules and the utopic distribution \mathcal{U} . Our goal is to find a vector V so that the winning committees under \mathcal{R}_V follow \mathcal{U} as closely as possible. To make this notion precise, our algorithms first compute a given number N of interval elections E_1, \dots, E_N (these are fixed throughout the whole optimization process). To evaluate the rule \mathcal{R}_V , for each election E_i we compute the winning committee W_i (if there are ties, then we break them arbitrarily). Then we compute the average distance of these committees from the utopia, $\text{EMD}(\mathcal{U}, \mathcal{R}_V) = 1/N \sum_{i=1}^N \text{EMD}(\mathcal{U}, d_{W_i})$; this value is referred to as the *score* of the rule (the lower, the better).

Brute-Force Search. This algorithm tries all rules defined by vectors with values from the set $\{0, 1/t, \dots, t-1/t, 1\}$, where t is a parameter defining the resolution of the search, and outputs the rule with the lowest score. We used it only to verify the quality of the local search algorithm below on small elections.

Local Search. We use a simple local search algorithm with the following parameters: (a) the number of iterations T , (b) the probability $\omega(i) \in [0, 1]$ of changing a given vector’s coordinate, depending on the iteration number i , (c) the range parameter $r(i) \in [0, 1]$, specifying how much vector coordinates can change, depending on the iteration number i . The algorithm works as follows:

- (1) We select a vector V uniformly at random, with coordinates from $[0, 1]$, and normalize it.
- (2) We repeat the following steps T times:
 - (a) Create a vector V' using the following procedure. Set $V' = V$. Then, for each of its coordinates v'_i , compute v''_i by adding to v'_i a value drawn uniformly at random from $[-r(i), r(i)]$. With probability $\omega(i)$, replace the value of v'_i with v''_i . Normalize V' .
 - (b) If $\text{EMD}(\mathcal{U}, \mathcal{R}_{V'}) < \text{EMD}(\mathcal{U}, \mathcal{R}_V)$ then replace V with V' . Otherwise, keep V as is.

- (3) We output the rule \mathcal{R}_V .

For weakly separable rules, we use $T = 3000$ iterations, $\omega(i) = \max(\frac{T-i}{2T}, 0.05)$, $r(i) = 0.5\omega(i)$, and $N = 400$ test elections. For OWA/Borda-based rules, we use $T = 300$ iterations, $\omega(i) = \max(\frac{T-i}{2T}, 0.1)$, $r(i) = 0.3 \cdot \max(\frac{T-i}{2T}, 0.05)$, $N = 40$ test elections. To speed up the algorithm for the case of OWA/Borda-based rules, we first run it for elections with 50 candidates, 50 voters, and committee size 10, and only then we re-run it for full-sized elections (with 100 candidates, 100 voters, and committee size 10), using the result of the first run as the input for the second one.

Naturally, we could have chosen any other algorithm for multi-dimensional optimization. We chose the simplest effective solution.

4 RESULTS

We used our search algorithm to find the best weakly separable and OWA/Borda-based rules for the individual excellence, ϵ -diversity (with $\epsilon \in \{0, 0.1, 0.2\}$), ϵ -twin peaks (with $\epsilon \in \{0.167, 0.25, 0.333\}$), and ϵ -triangle (with $\epsilon \in \{0, 0.1, 0.2\}$) utopic distributions.

The results are given in Table 1 (where we show EMD distances for the best rules we computed using our algorithm, and for the five rules from Example 2.1; we used other elections than in the optimization process) and in Figures 4–9. Each figure shows results for four utopic distributions: the individual excellence utopia, which is a border case for the other distributions, and either ϵ -diversity, ϵ -twin peak, or ϵ -triangle distributions, for appropriate values of ϵ . The largest plot on the left of each figure, marked (a), shows vectors computed for the respective four utopias. Next to it, as Plot (b), we show graphical representation of the respective utopia (drawn as a gray area over the $[0,1]$ interval) and a sample result of a single interval election (the blue dots). As Plot (c), we show the 1D histograms achieved by the computed rules (we remind the reader that different histograms have different y -axis scales, as they only show the “shape” of the result). Finally, as Plot (d), we show the scatter plots computed for disc elections according to our four rules. The vectors computed for the utopias are marked with a number (1–4) and the respective figures in (b)–(d) are marked accordingly.

EMD Results Versus Histograms. The EMD results in Table 1 show that generally our best OWA/Borda-based rules are much closer to respective utopias than our best weakly separable rules. While this is in agreement with intuition, it may sometimes be surprising that weakly separable rules achieve far more visually appealing histograms for some settings than the OWA/Borda-based ones (e.g., for the twin-peaks distributions), in spite of having worse EMD values. The reason for this discrepancy is that weakly separable rules achieve good histograms “in the aggregate” (averaged over many elections), whereas OWA/Borda rules perform well (but not great) for every election instance.

Weakly Separable Rules. Next we discuss the results for weakly separable rules. For individual excellence (see Figure 4) we obtained a nearly linear vector, very close to the Borda scoring function (vector 1). On the other extreme, for the diversity utopia, we found a rule very close to SNTV (vector 4, which is 0 for most positions, then slowly increases, and jumps to value 1 for the first position²). We view it as a negative result: our hope for finding a weakly separable

²For the elections used in the optimization process, our rule was a bit better than SNTV, but SNTV was better for other randomly chosen elections (see Table 1).

Table 1: EMD values for the best weakly separable and OWA/Borda-based rules computed for our utopias, and for the rules from Example 2.1 for the same utopias.

Utopia	Weakly Sep.	OWA/Borda	SNTV	Bloc	Borda	CC	HB
\mathcal{U}_D	0.104	0.042	0.092	0.215	0.227	0.044	0.094
$\mathcal{U}_D^{0.1}$	0.098	0.042	0.100	0.205	0.178	0.059	0.053
$\mathcal{U}_D^{0.2}$	0.073	0.044	0.126	0.210	0.130	0.100	0.044
$\mathcal{U}_{TP}^{0.167}$	0.235	0.135	0.159	0.237	0.307	0.140	0.172
$\mathcal{U}_{TP}^{0.25}$	0.185	0.060	0.144	0.222	0.224	0.124	0.116
$\mathcal{U}_{TP}^{0.333}$	0.135	0.081	0.156	0.221	0.141	0.136	0.093
\mathcal{U}_T	0.139	0.043	0.113	0.213	0.150	0.082	0.043
$\mathcal{U}_T^{0.1}$	0.136	0.045	0.137	0.220	0.127	0.116	0.051
$\mathcal{U}_T^{0.2}$	0.122	0.054	0.174	0.238	0.105	0.160	0.083
\mathcal{U}_E	0.041	0.049	0.248	0.264	0.049	0.244	0.168

rule that would robustly implement the diversity utopia turned out to be unrealistic (SNTV does not implement this utopia robustly as its results seem to be a statistical artifact: SNTV chooses candidates from areas with lower density of candidates and increased density of voters, which, statistically, happen equally often in each area of the interval).

Perhaps the most interesting results are those achieved for 0.1-diversity and 0.2-diversity (vectors 2 and 3), as they show rules that, if at all, are very rarely discussed in the literature. Both vectors 2 and 3 resemble functions of the form $\gamma(i) = (1-x)\alpha_1(i) + x\beta_m(i)$, where $x \in [0, 1]$ is a parameter (and, in our case, is close to 0.2). In other words, these functions give score 1 to position 1, score ≈ 0.2 to position 2, and then decrease linearly to 0. One could say that their Borda-score component is too small to be relevant, but this is not so. In our 1D elections we have 100 voters, which means that there are only 100 points to be distributed for being ranked on the first place, while there are ≈ 1000 points to be distributed for being ranked on the following places. This way, the rules described by vectors 2 and 3 achieve a compromise between SNTV and Borda.

The results for the ϵ -twin peaks utopic distribution (Figure 5) are quite spectacular. For each $\epsilon \in \{0.167, 0.25, 0.333\}$ we find a rule whose 1D histogram matches the respective utopic distribution very well. The twin-peaked distributions were inspired by the results for the Bloc rule (recall Figure 1); indeed, we find vectors consisting of several 1s followed by 0s (with a very rapid transition). However, as opposed to Bloc, our vectors have many more 1s (Bloc would have k of them, i.e., 10 in our case, but our rules have between 20 and 40). Apparently, this is the reason why our rules match the twin-peaked utopias better than Bloc, which selects more candidates “between the peaks”. A final remark regarding the twin peaks distributions regards individual excellence: Our results show that, as we put the peaks closer to each other, we obtain vectors of more 1s, followed by fewer 0s, but when the peaks finally coincide, we should obtain the linear function. Apparently, either there is some sort of phase transition between these two extremes, or we did not put the peaks close enough to each other to observe a smooth transition.

For the triangle utopias (see Figure 6), we seem to find rules whose scoring vectors resemble the shape of the harmonic sequence $1, 1/2, \dots, 1/m$. These rules, indeed, seem to achieve a compromise

between individual excellence (i.e., choosing winners from the center) and diversity (choosing winners from the whole interval/disc). Further, on the intuitive level, these rules are more appealing than the SNTV/Borda compromise obtained for ϵ -diversity. Indeed, this suggests that studying a variant of k -Borda that uses harmonic numbers instead of linearly decreasing scores might be interesting.

OWA/Borda-Based Rules. For individual excellence (see, e.g., Figure 7), our algorithm very quickly finds OWA vector of all 1s; thus, we obtain the k -Borda rule. For diversity (Figure 7, vector 4), we find a vector close to that of Chamberlin–Courant (a single 1 followed by 0s). For 0.1-diversity and 0.2-diversity we find, respectively, a linearly decreasing vector 2, and vector 3 that resembles (but, admittedly, quite poorly) the harmonic sequence. These two results are intriguing. First, the linear vector is a very natural solution to finding a compromise between excellence and diversity (which, in this case, would mean finding a compromise between k -Borda and Chamberlin–Courant) that has not been considered in the literature yet (even though Faliszewski et al. [19] look for rules that achieve such a compromise, they do not study this rule). On the other hand, the harmonic vector has received extensive treatment, both for Harmonic Borda [19] and for the PAV rule [1, 21, 22, 36] (which is approval-based and has motivated the study of Harmonic Borda). In fact, the OWA vectors that we have obtained for the triangle distributions (see Figure 9) seem to be closer to the harmonic sequence. This confirms the intuition of Faliszewski et al. [19] that this sequence achieves a good excellence/diversity compromise.

Finally, we consider the results for the twin peaks distributions (Figure 8). In this case we find OWA vectors that consist of 1s followed by 0s (the number of ones depends on the distance between the peaks). This means that the rules that we found are, in essence, the t -Borda rules of Faliszewski et al. [19] (for a given t , the t -Borda rule uses Borda scoring function and OWA vector of t 1s followed by 0s). Faliszewski et al. studied these rules in their search for excellence/diversity compromises, but concluded that they do not seem to work well for this case. The fact that they implement the twin peaks utopic distribution supports this conclusion.

Results for Disc Elections. Generally, our rules behave similarly on interval and disc elections, but our OWA/Borda-based rule for the 0.1-twin peaks distribution is an exception. For interval elections it outputs winners on the peaks, as well as between them, but for disc elections, it provides winners on the ring only.

5 CONCLUSIONS

We have developed a methodology for automatically designing multiwinner voting rules whose winning committees have properties specified via distributions on a 1D interval. Testing our method on weakly separable and OWA/Borda-based committee scoring rules, we confirmed many intuitions about the applicability of certain rules for certain tasks and discovered new rules to study. Our work is a proof of concept and shows that our approach is indeed feasible. A natural direction for future studies is to expand our framework for other types of goals, such as proportional representation.

ACKNOWLEDGMENTS

Piotr Faliszewski was supported by the National Science Centre, Poland, under project 2016/21/B/ST6/01509.

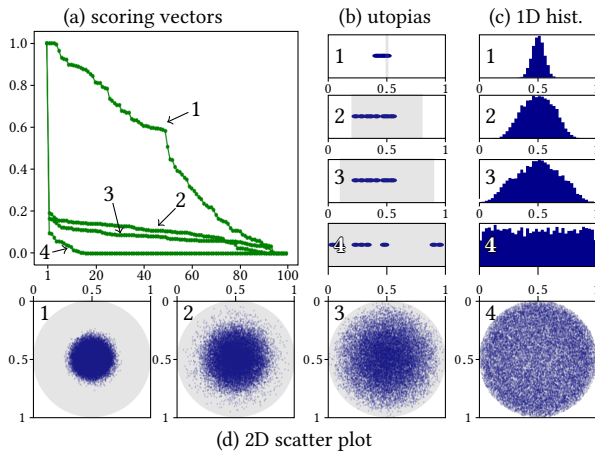


Figure 4: Results for weakly separable rules and utopic distributions (1) \mathcal{U}_{IE} , (2) $\mathcal{U}_D^{0.2}$, (3) $\mathcal{U}_D^{0.1}$, and (4) \mathcal{U}_D .

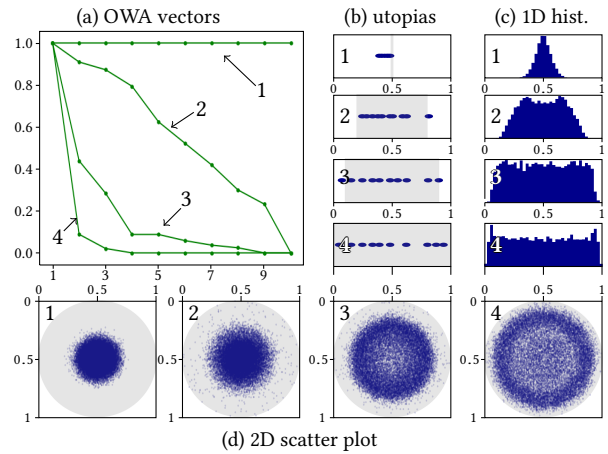


Figure 7: Results for OWA/Borda-based rules and utopic distributions (1) \mathcal{U}_{IE} , (2) $\mathcal{U}_D^{0.2}$, (3) $\mathcal{U}_D^{0.1}$, and (4) \mathcal{U}_D .

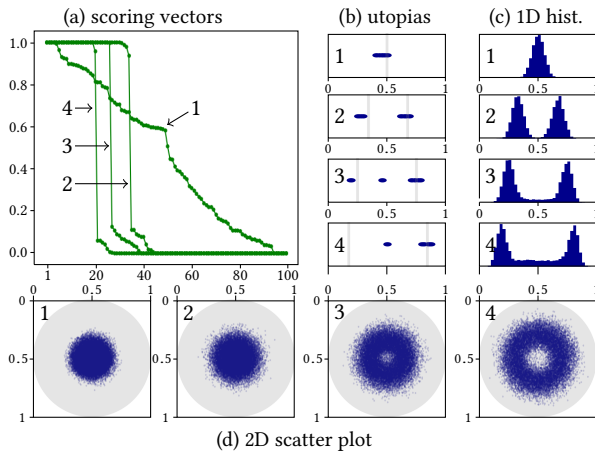


Figure 5: Results for weakly separable rules and utopic distributions (1) \mathcal{U}_{IE} , (2) $\mathcal{U}_{TP}^{0.333}$, (3) $\mathcal{U}_{TP}^{0.25}$, and (4) $\mathcal{U}_{TP}^{0.167}$.

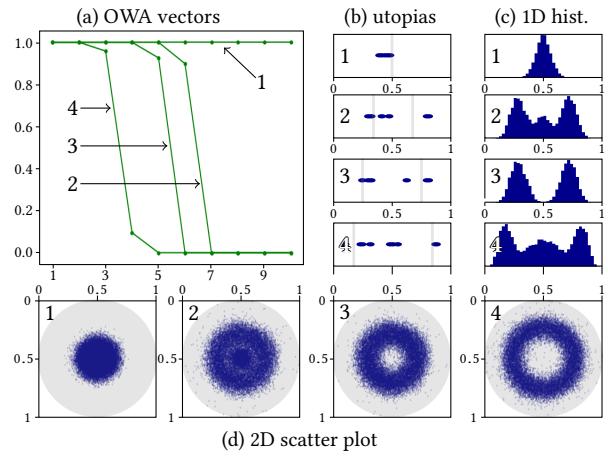


Figure 8: Results for OWA/Borda-based rules and utopic distributions (1) \mathcal{U}_{IE} , (2) $\mathcal{U}_{TP}^{0.333}$, (3) $\mathcal{U}_{TP}^{0.25}$, and (4) $\mathcal{U}_{TP}^{0.167}$.

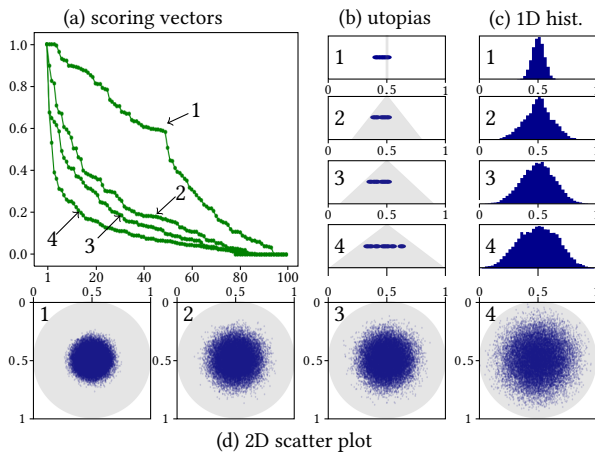


Figure 6: Results for weakly separable rules and utopic distributions (1) \mathcal{U}_{IE} , (2) $\mathcal{U}_T^{0.2}$, (3) $\mathcal{U}_T^{0.1}$, and (4) \mathcal{U}_T^0 .

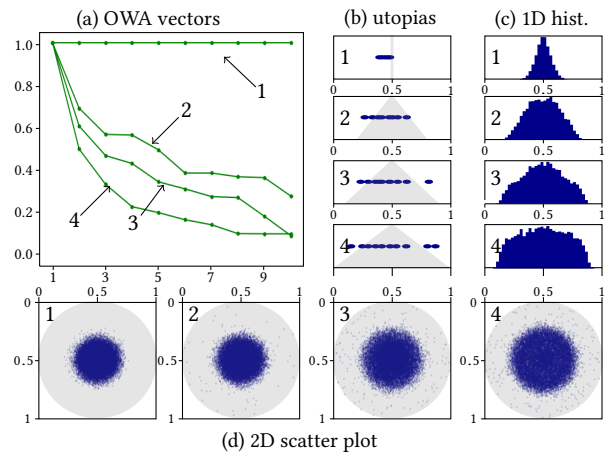


Figure 9: Results for OWA/Borda-based rules and utopic distributions (1) \mathcal{U}_{IE} , (2) $\mathcal{U}_T^{0.2}$, (3) $\mathcal{U}_T^{0.1}$, and (4) \mathcal{U}_T^0 .

REFERENCES

- [1] H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. 2017. Justified Representation in Approval-Based Committee Voting. *Social Choice and Welfare* 48, 2 (2017), 461–485.
- [2] H. Aziz, E. Elkind, P. Faliszewski, M. Lackner, and P. Skowron. 2017. The Condorcet Principle for Multiwinner Elections: From Shortlisting to Proportionality. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI-17)*. 84–90.
- [3] H. Aziz, S. Gaspers, J. Gudmundsson, S. Mackenzie, N. Mattei, and T. Walsh. 2015. Computational Aspects of Multi-Winner Approval Voting. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-15)*. 107–115.
- [4] D. Baumeister, S. Dennen, and L. Rey. 2015. Winner Determination and Manipulation in Minisum and Minimax Committee Elections. In *Proceedings of the 4th International Conference on Algorithmic Decision Theory (ADT-15)*. 469–485.
- [5] N. Betzler, A. Slinko, and J. Uhlmann. 2013. On the Computation of Fully Proportional Representation. *Journal of Artificial Intelligence Research* 47 (2013), 475–519.
- [6] I. Caragiannis, S. Nath, A. Procaccia, and N. Shah. 2017. Subset Selection Via Implicit Utilitarian Voting. *Journal of Artificial Intelligence Research* 58 (2017), 123–152.
- [7] B. Chamberlin and P. Courant. 1983. Representative Deliberations and Representative Decisions: Proportional Representation and the Borda Rule. *American Political Science Review* 77, 3 (1983), 718–733.
- [8] M. Cygan, L. Kowalik, A. Socala, and K. Sornat. 2017. Approximation and Parameterized Complexity of Minimax Approval Voting. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI-17)*. 459–465.
- [9] B. Debord. 1992. An Axiomatic Characterization of Borda’s k -Choice Function. *Social Choice and Welfare* 9, 4 (1992), 337–343.
- [10] M. Diss and A. Doghmi. 2016. Multi-Winner Scoring Election Methods: Condorcet Consistency and Paradoxes. *Public Choice* 169, 1–2 (2016), 97–116.
- [11] E. Elkind, P. Faliszewski, J. Laslier, P. Skowron, A. Slinko, and N. Talmon. 2017. What Do Multiwinner Voting Rules Do? An Experiment Over the Two-Dimensional Euclidean Domain. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI-17)*. 494–501. Full version available at <http://home.agh.edu.pl/~faliszew/2d.pdf>.
- [12] E. Elkind, P. Faliszewski, P. Skowron, and A. Slinko. 2017. Properties of Multiwinner Voting Rules. *Social Choice and Welfare* 48, 3 (2017), 599–632.
- [13] E. Elkind, P. Faliszewski, and A. Slinko. 2015. Distance Rationalization of Voting Rules. *Social Choice and Welfare* 45, 2 (2015), 345–377.
- [14] E. Elkind, J. Lang, and A. Saffidine. 2015. Condorcet winning sets. *Social Choice and Welfare* 44, 3 (2015), 493–517.
- [15] J. Enelow and M. Hinich. 1984. *The Spatial Theory of Voting: An Introduction*. Cambridge University Press.
- [16] J. Enelow and M. Hinich. 1990. *Advances in the Spatial Theory of Voting*. Cambridge University Press.
- [17] P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. 2016. Committee Scoring Rules: Axiomatic Classification and Hierarchy. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI-16)*. 250–256.
- [18] P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. 2016. Multiwinner Analogues of the Plurality Rule: Axiomatic and Algorithmic Views. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI-16)*. 482–488.
- [19] P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. 2017. Multiwinner Rules on Paths From k -Borda to Chamberlin–Courant. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI-17)*. 192–198.
- [20] P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. 2017. Multiwinner Voting: A New Challenge for Social Choice Theory. In *Trends in Computational Social Choice*, U. Endriss (Ed.). AI Access Foundation.
- [21] M. Kilgour. 2010. Approval Balloting for Multi-Winner Elections. In *Handbook on Approval Voting*. Springer. Chapter 6.
- [22] M. Lackner and P. Skowron. 2017. *Consistent Approval-Based Multi-Winner Rules*. Technical Report arXiv:1704.02453 [cs.GT].
- [23] E. Lerner and S. Nitzan. 1985. Some General Results on the Metric Rationalization for Social Decision Rules. *Journal of Economic Theory* 37, 1 (1985).
- [24] T. Lu and C. Boutilier. 2011. Budgeted Social Choice: From Consensus to Personalized Decision Making. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI-11)*. 280–286.
- [25] T. Meskanen and H. Nurmi. 2008. Closeness Counts in Social Choice. In *Power, Freedom, and Voting*, M. Braham and F. Steffen (Eds.). Springer-Verlag.
- [26] B. Monroe. 1995. Fully proportional representation. *American Political Science Review* 89, 4 (1995), 925–940.
- [27] S. Nitzan. 1981. Some Measures of Closeness to Unanimity and Their Implications. *Theory and Decision* 13, 2 (1981), 129–138.
- [28] S. Peleg, M. Werman, and H. Rom. 1989. A Unified Approach to the Change of Resolution: Space and Gray-Level. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 11, 7 (1989), 739–742.
- [29] D. Peters. 2016. *Single-Peakedness and Total Unimodularity: Efficiently Solve Voting Problems Without Even Trying*. Technical Report arXiv:1609.03537 [cs.GT].
- [30] A. Procaccia, J. Rosenschein, and A. Zohar. 2008. On the Complexity of Achieving Proportional Representation. *Social Choice and Welfare* 30, 3 (2008), 353–362.
- [31] L. Sánchez-Fernández, E. Elkind, M. Lackner, N. Fernández, J. A. Fisteus, P. Basanta Val, and P. Skowron. 2017. Proportional Justified Representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI-17)*. 670–676.
- [32] S. Sekar, S. Sikdar, and L. Xia. 2017. Condorcet Consistent Bundling with Social Choice. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-17)*. 33–41.
- [33] P. Skowron, P. Faliszewski, and J. Lang. 2016. Finding a collective set of items: From proportional multirepresentation to group recommendation. *Artificial Intelligence* 241 (2016), 191–216.
- [34] P. Skowron, P. Faliszewski, and A. Slinko. 2016. *Axiomatic Characterization of Committee Scoring Rules*. Technical Report arXiv:1604.01529 [cs.GT]. arXiv.org.
- [35] P. Skowron, M. Lackner, M. Brill, D. Peters, and E. Elkind. 2017. Proportional Rankings. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI-17)*. 409–415.
- [36] T. Thiele. 1895. Om Flerfoldsvvalg. In *Oversigt over det Kongelige Danske Videnskabsberens Selskabs Forhandlinger*. 415–441.
- [37] L. Xia. 2013. Designing social choice mechanisms using machine learning. In *Proceedings of AAMAS 2013*. 471–474.