Reg. No: 2016/21/B/ST6/01509; Principal Investigator: dr hab. inż. Piotr Faliszewski

Algorithmic and Axiomatic Analysis of Committee Scoring Rules

The goal of the project is to study multiwinner election rules, i.e., rules that pick sets of candidates of a given size (referred to as a committees). Parliamentary elections are, perhaps, the most typical multiwinner elections, but there are many other applications. For example, consider a company that wishes to pick a palette of products for its customers. The company can treat the products as candidates, the customers as voters (whose preferences the company obtains through market research), and can use a voting rule that guarantees finding a committee (a palette of products) so that for every customer there is at least one product that the customer likes. The Chamberlin–Courant's rule is one example of such a voting rule.

We plan to study committee scoring rules, a new class of multiwinner election rules that contains many very diverse rules (including, for example, the Chamberlin–Courant rules; also see Figure 1). Due to this diversity, for most election tasks one can find an appropriate committee scoring rule. Unfortunately, currently there is no strong, coherent theory that would explain which committee scoring rules have which features and, worse yet, most committee scoring rules are NP-hard to compute. This means that it is not clear which rules to use in practice, or even how to obtain their resutls (a naive algorithm could take even hundreds of years on realistically sized data). In this project we attempt to resolve these two problems as follows:

1. We aim to create a series of mathematical theorems that would describe the properties of committee scoring rules. These theorems would say, for example, which committee scoring rules are majority consistent (i.e., which of them guarantee that a coordinated majority of voters can ensure victory of a given committee), or which committee scoring rules guarantee diversity of the committee (i.e., that every voter is, in some sense, well-represented by at least one committee member). Our preliminary results show that many such properties are inconsistent with each other and, thus, it is necessary to formulate a theory that would explain which features, and when, can be implemented within committee scoring rules. These results will allow companies and communities to pick committee scoring rules most suitable to their needs.

2. Our second goal is to develop algorithms that would effectively compute the results of committee scoring rules (perhaps approximately). Our initial results show that, indeed, this appears to be possible for a wide class of rules.

Altogether, our results should provide a strong tool that would enable one to choose appropriate committee scoring rules and apply them in practice.

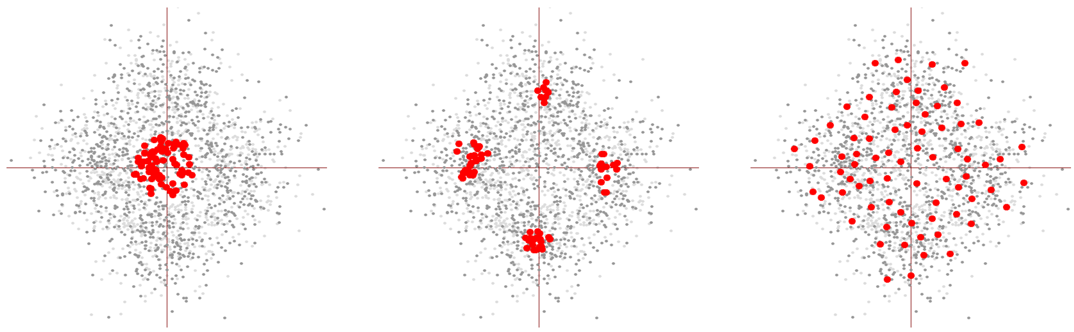

(a) k-Borda (b) Bloc (c) β-CC

Figure 1: Results of three distinct committee scoring rules. Dark gray points are candidates, light gray points are voters, and red dots are members of the winning committee. The voters form their preferences based on their Euclidean distance from the candidates (the closer a candidate is to the voter, the more this voter likes this candidate; thus the picture illustrates the space of ideas). Both voters and candidates were generated using a mixture of four Gaussian distributions (placed symmetrically). As we can see, different rules lead to very different committees. k-Borda picks a consensus committee that corresponds to the "average ideas", Bloc finds areas with high density of voters (in this case, the centers of our four Gaussians), and Chamberlin–Courant's rule (β-CC) finds a committee where every voter is close to some committee member.