

Pabuviz.org: A Visualisation Platform to Explore Participatory Budgeting Elections

Demonstration Track

Markus Utke

Eindhoven University of Technology
The Netherlands

Simon Rey

ILLC, University of Amsterdam
The Netherlands

Ulle Endriss

ILLC, University of Amsterdam
The Netherlands

ABSTRACT

We present pabuviz.org, an interactive website that provides in-depth visualisation tools to compare the outcomes that would have been returned by a range of different voting rules for participatory budgeting when applied to historical election data. This information can help policy makers with choosing an appropriate voting rule.

KEYWORDS

Participatory Budgeting; Data Visualisation; Knowledge Transfer

ACM Reference Format:

Markus Utke, Simon Rey, and Ulle Endriss. 2025. Pabuviz.org: A Visualisation Platform to Explore Participatory Budgeting Elections: Demonstration Track. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025)*, Detroit, Michigan, USA, May 19 – 23, 2025, IFAAMAS, 3 pages.

1 INTRODUCTION

Participatory budgeting (PB) is an umbrella term that covers a range of democratic innovations aimed at directly involving citizens with local budgeting decisions [3, 11]. Typically, citizens are first invited to propose projects for their neighbourhood—such as planting a tree or putting up a ping-pong table in a park. Each project has a given cost and there is a limit to the available budget, meaning that not all projects can be realised. Citizens are then asked to vote, usually by selecting a small number of projects they support.

Most of the research effort in the multiagent systems community, and specifically in the computational social choice community [9], has been, and still is, devoted to finding attractive voting rules, i.e., procedures that aggregate the ballots submitted by the voters to then output a budgeting decision. Several such rules have been proposed in the literature [1, 2, 5, 6, 8, 10], each with their own characteristics, making each such rule either more or less suitable in any given context.

It can be difficult, certainly for a non-expert, to really understand the differences between the available voting rules and thus to evaluate the high-level impact the selection of any one specific rule would have on their community. And even an expert who has a good understanding of the theoretical properties of different voting rules might find it hard to predict how those theoretical differences will play out in practice in response to the peculiarities of voter preferences in a given city. Therefore, to support policy makers in

choosing a suitable voting rule, we have developed pabuviz.org, a visualisation platform for PB. It provides intuitive and visually appealing comparison tools for PB, based on real-life data from past PB elections. It can be used as a helper tool when discussing possible voting rules for PB.

In this short paper, we briefly describe the platform’s infrastructure, we then move on to illustrate its functionality, and we conclude by outlining possible use cases.¹

2 PLATFORM INFRASTRUCTURE

The platform is composed of two elements: a database and a visualisation website.

The database hosts the data used by the website. This includes real-life election data retrieved from pabulib.org [4]; the outcomes under multiple voting rules for each of these elections, computed using the `pabutools` package [4]; and a variety of metrics and statistics used for visualisation.

The database is running independently from the visualisation website and can be accessed via an API (see db.pabuviz.org/api). The code is available as an open-source GitHub repository.²

The visualisation website—available at pabuviz.org—provides interactive access to the available visualisations. It has been developed as a REACT application. The website queries the database’s API to serve the visualisations. The code is available as an open-source GitHub repository.³ We describe the various tools available below.

3 PLATFORM FUNCTIONALITY

Using the website, you can compare voting rules for PB, such as (different versions of) the Greedy Rule, the Sequential Phragmén Rule, and the Method of Equal Shares [9]. Voting rules are always compared in view of the characteristics of their outcomes. They can be compared either at the level of a single election or at the level of a collection of elections.

The website permits working with data from elections using different ballot formats (approval, ordinal, cardinal,⁴ and cumulative). Here we only discuss the case of approval ballots, where voters can indicate for each project whether or not they support it. The comparison tools for the other ballot formats are similar.

3.1 Comparing Rules for Single Elections

After the user has selected an election and a number of voting rules, the site provides the following visualisations.



This work is licensed under a Creative Commons Attribution International 4.0 License.


¹A short video presenting the platform is available at youtu.be/oir0pJFvfc4.

²github.com/COMSOC-Community/pabuviz-db

³github.com/COMSOC-Community/pabuviz-web

⁴At the time of writing, cardinal ballots are not actually displayed on the website as so far there is no data on elections using cardinal ballots in the database.


Satisfaction histogram. Rules differ in how well they distribute individual satisfaction amongst voters. Visualising these differences results in graphs such as the one shown below.



Each colour corresponds to a voting rule (as selected from a menu by the user). A point at coordinates (x, y) indicates that for $y\%$ of the voters it is the case that (around) $x\%$ of the budget has been spent on projects they individually approve of. The vertical dotted lines indicate average satisfaction.

A high y -coordinate in the lefthand part of the graph indicates that many voters are very unhappy; a high y -coordinate in the righthand part indicates that many voters are fairly happy. For instance, in our example, the rule shown in green (the Greedy Rule) performs best in terms of average satisfaction, but also leaves a very large number of voters completely unsatisfied (at $x = 0$).

Measures to compare election outcomes. We have formulated a set of measures defined on election outcomes that allow us to rank voting rules. On the website, this information is rendered in the form of a radar chart, providing an efficient and easy-to-interpret representation. An example is shown below.




Let us briefly describe the different measures:

- **Average cardinality satisfaction:** Average over all voters of the number of approved projects selected by the rule.
- **Average relative cardinality satisfaction:** Relativised version of the above, where for each voter we divide by the maximum achievable satisfaction.
- **Average cost satisfaction:** Average over all voters of the total budget spent on approved projects.
- **Average relative cost satisfaction:** Relativised version.
- **Inverted cost Gini:** Inverse of the Gini coefficient of the cost satisfaction of the voters.
- **Voters with positive satisfaction:** Percentage of voters for whom at least one approved project has been selected.

Representation of project categories. For some elections, projects are assigned to *categories* (such as ‘social’ or ‘environmental’). We can then measure the average budget distribution over the categories in the ballots submitted. We call this the *vote share*. The

outcomes returned by different rules can then be compared to this vote share, leading to visualisations such as the one shown below.



It has been argued that the closer to the vote share an outcome is, the better the rule that generated that outcome [4].

3.2 Comparing Rules at the Aggregate Level

The website also offers the possibility to compare rules based on several elections (e.g., all those for a given city, or all those of a certain size). For this tool we reuse the satisfaction histograms and the measure radar charts described earlier. The relevant values are averaged for all the elections considered.

4 TYPICAL USE CASES

The tools provided by pabuviz.org were designed with the objective of providing visual support when trying to understand and explain the effects of different voting rules for PB. We stress that pabuviz.org (unlike, for example, equalshares.net) is not intended for use by individuals who do not have any kind of formal training in social choice theory. Rather, it assumes the presence of a ‘guide’ who can explain how to interpret the visual information provided. We have two concrete use cases in mind.

Discussions with policy makers. One of the goals of research on PB is to help decision makers (e.g., the civil servants tasked with running PB in a given city) select the voting rule that best suits their needs. As a researcher, when in discussion with such a decision maker, one can use pabuviz.org to present ready-made visualisations regarding the different rules using real-life elections.⁵ It is also possible to upload your own election file to see the visualisations specific to the city you are discussing. Note that all visualisations can be exported in either PNG or JSON format and can thus easily be integrated into other presentation tools.

Exploration of data. The platform can also be used by researchers who want to explore PB rules, to better understand how these rules perform in practice. This can be done either directly via the website or by using the API to extract the relevant information.

REFERENCES

- [1] Haris Aziz, Barton E. Lee, and Nimrod Talmon. 2018. Proportionally Representative Participatory Budgeting: Axioms and Algorithms. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 23–31.
- [2] Markus Brill, Stefan Forster, Martin Lackner, Jan Maly, and Jannik Peters. 2023. Proportionality in Approval-Based Participatory Budgeting. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*.
- [3] Yves Cabannes. 2004. Participatory Budgeting: A Significant Contribution to Participatory Democracy. *Environment and Urbanization* 16, 1 (2004), 27–46.

⁵We have previously done so in discussions with representatives of the Municipality of Amsterdam, using visualisations similar to those generated by pabuviz.org [7].

- [4] Piotr Faliszewski, Jarosław Flis, Dominik Peters, Grzegorz Pierczyński, Piotr Skowron, Dariusz Stolicki, Stanisław Szufa, and Nimrod Talmon. 2023. Participatory budgeting: Data, tools, and analysis. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*. 2667–2674.
- [5] Till Fluschnik, Piotr Skowron, Mervin Triphaus, and Kai Wilker. 2019. Fair Knapsack. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*. 1941–1948.
- [6] Maaïke Los, Zoé Christoff, and Davide Grossi. 2022. Proportional Budget Allocations: Towards a Systematization. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*. 398–404.
- [7] Pelle Nelissen. 2023. An Empirical Analysis of Participatory Budgeting in Amsterdam. *arXiv preprint arXiv:2310.18033* (2023).
- [8] Dominik Peters, Grzegorz Pierczyński, and Piotr Skowron. 2021. Proportional Participatory Budgeting with Additive Utilities. In *Proceedings of the 35th Annual Conference on Neural Information Processing Systems (NeurIPS)*. 12726–12737.
- [9] Simon Rey and Jan Maly. 2023. The (Computational) Social Choice Take on Indivisible Participatory Budgeting. *arXiv preprint arXiv:2303.00621* (2023).
- [10] Nimrod Talmon and Piotr Faliszewski. 2019. A Framework for Approval-Based Budgeting Methods. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*. 2181–2188.
- [11] Brian Wampler. 2000. A Guide to Participatory Budgeting. *Third Conference of the International Budget Project* (2000).