Banzhaf Power in Hierarchical Voting Games

Extended Abstract

John Randolph
Brown University
Providence, United States
john_randolph@alumni.brown.edu

Amy Greenwald
Brown University
Providence, United States
amy_greenwald@brown.edu

Denizalp Goktas
Brown University
Providence, United States
denizalp_goktas@brown.edu

ABSTRACT

The Banzhaf Power Index (BPI) is a method of measuring the power of voters in determining the outcome of a voting game. Some voting games exhibit a hierarchical structure, including the US electoral college and ensemble learning methods; we call such games hierarchical voting games. It is generally understood that BPI in hierarchical voting games can be computed via a recursive decomposition of the hierarchy, which can substantially reduce the calculation's complexity. We identify a key (previously undocumented) assumption on which this decomposition is based, namely balance, meaning one group of voters has enough votes to win whenever the complementary group of voters does not, and vice versa. We then introduce a generalization of BPI that we call Extended BPI (EBPI) for all voting games, including those that are not balanced, which simplifies to BPI in balanced games. We show that BPI in unbalanced hierarchical voting games decomposes in terms of EBPI at each level in the hierarchy, which yields computational savings analogous to those achieved in the balanced case. As a sample application, we take advantage of the compositionality of language, and model the impact of individual words on a sentence's sentiment as a voting game. As the complement of a phrase in a sentence does not necessarily have the opposite sentiment, this voting game is unbalanced and requires our decomposition of BPI in terms of EBPI. Our results suggest that EBPI is an effective proxy for BPI (because the meaning of a sentence is not always 100% compositional), and demonstrate a dramatic improvement in run time.

KEYWORDS

Banzhaf Power Index; Cooperative Game Theory; Voting Games; Social Choice

ACM Reference Format:
John Randolph, Amy Greenwald, and Denizalp Goktas. 2024. Banzhaf Power in Hierarchical Voting Games: Extended Abstract. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 3 pages.

1 BACKGROUND

Elections. Elections are a form of group decision-making intended to produce outcomes that reflect the collective preferences of the group, or at least give the impression of doing so. The designers of an election might want their system to embody principles of fairness, such as "one person, one vote" [4, 11]. In complex elections, it is often unclear how much influence an individual voter has on the outcome, making it difficult to assess whether the system is fair or not. Measures of power can help us determine whether our elections meet our standards of fairness.

Elections are often modeled as a type of cooperative game called voting games [8, 28]. A voting game comprises a set of voters and a characteristic function, the latter of which takes as input a coalition, i.e., a subset of voters, and returns a binary variable. If the characteristic function outputs a 1 (resp. 0), we say that the coalition is a winning (resp. losing) coalition.

Electoral processes in representative democracies can have hierarchical structure. In the US, for example, voters vote for electoral college delegates, who then vote to elect the president. Likewise, voters in Slovenian National Council elections vote for a representative in their local councils, who in turn vote for national councillors, who in turn vote on legislation in the National Council. We call these voting games hierarchical, as they are structured as trees, with voters as the leaf nodes, such that their votes propagate up the tree to intermediate nodes, with the result at each level determined by the characteristic function at that level.

Banzhaf Power Index. In 1965, Banzhaf defended a mathematical measure of power in voting games [5] first introduced by Penrose [21] but then largely forgotten [9], now called the Banzhaf power index (BPI). BPI has since found many applications, such as the US electoral college [6], various councils and parliaments within the European Union [3, 10, 12, 22, 24], the International Monetary Fund [2], feature importance in machine learning [15], and shareholders and corporate boards [13, 16].

The formula for computing BPI is exponential: each of the 2^n coalitions must be examined to determine whether it is a winning coalition, with and without each of the n voters. This leads to a time complexity of $O(2^n)$ per voter [14], which can be prohibitive, for example, when n is the number of eligible voters in the US.

BPI in hierarchical voting games, however, decomposes in terms of BPI at each of the d levels in the hierarchy, according to a formula we call multiplicative BPI (MBPI) [19]. This decomposition provides a massive speedup when the branching factor $K \ll n$.

2 CONTRIBUTIONS

Balanced games. Our first contribution is to identify the key mathematical property of a voting game that is necessary for the decomposition of BPI into MBPI. This property is balance (or proper strong [27]), which connotes that the complement of every winning

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand. © 2024 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

	Naive (BPI)	Balanced (MBPI)	General (MEBPI)
One voter	$O(2^n)$	$O(d2^K)$	$O(ndK2^K)$
All voters	$O(n2^n)$	$O(nd2^K)$	$O(ndK2^K)$

Table 1: BPI algorithm run times where d is the depth of the tree, n is the number of voters, and K is the branching factor.

(resp. losing) coalition is a losing (resp. winning) coalition. In games that are not balanced, MBPI does not necessarily recover BPI.

Unbalanced games. Although many common voting games are balanced, others are not. For example, if a supermajority (not a simple majority) is required to win, then the corresponding voting game is not balanced. In the US Senate, 60 votes are required to pass most legislation; thus, this game is unbalanced, as every set of 41 to 59 senators is a losing coalition, as is its complement.

In search of a decomposition theorem for unbalanced hierarchical voting games, we introduce a generalization of BPI that we call extended BPI (EBPI), which is applicable to all monotone voting games, balanced or not. (In a monotone voting game, adding a voter to a winning coalition cannot turn it into a losing coalition.) Then, just as BPI can be calculated efficiently via MBPI in balanced hierarchical games, we propose multiplicative EBPI (MEBPI) to efficiently calculate EBPI in unbalanced hierarchical games. Our main theorem states that MEBPI yields a novel, more efficient means of calculating BPI for this larger class of games.

The EBPI formula, which generalizes BPI, is a factor of K slower than the naive BPI formula. Furthermore, to use the MEBPI formula to calculate the BPI of even one voter requires calculating the number of winning and losing coalitions at *every* subgame in the hierarchy, and thus involves a full tree traversal. As a result, calculating MEBPI is substantially slower than calculating MBPI, so should only be used in unbalanced games. But like MBPI for balanced games, MEBPI for unbalanced games is a vast improvement over the naive approach. We summarize the complexity of our formulas as compared to naively computing BPI in Table 1.

3 EXPERIMENTS

To demonstrate the utility of our algorithms, we compare MBPI and MEBPI with the naive approach in two experimental settings.

Slovenian national council. To demonstrate the speedup in the balanced case, we calculate the BPI of individual voters in Velenje and Ljubljana in Slovenian National Council elections.

Slovenia's national council is elected via a complex process: there are 40 indirectly elected members, 22 that represent municipalities and 18 that represent special interests, such as sports and culture, or farmers. Each of these members is elected by an electoral college: for the 22 members that represent municipalities, these electoral colleges are the local assembly, and for the other 18 members, these electoral colleges include members of the sector they represent [18]. As such, this voting game is hierarchical.

Calculating the power of any individual Slovenian voter on a piece of legislation in the Slovenian National Council by computing BPI naively is intractable, since it would involve examining all 2^n coalitions, where n is the number of voters in Slovenia (around 2 million). However, as this voting game is balanced, we can compute BPI via MBPI. We find that the BPI of a voter in Velenje is about three times that of a voter in Ljubljana.

Vocabulary Selection. We next turn our attention to the problem of vocabulary selection in sentiment analysis, which we tackle by finding the words that are most influential (as measured by approximate BPI) in determining a text's sentiment [20].

Word importance is useful in the problem of vocabulary selection in natural language processing models. A smaller vocabulary makes models more interpretable [1, 23], requires less memory [25], is more amenable to use in a resource-constrained setting [26], and is less prone to over fitting [7, 17]. BPI has been used as a heuristic to solve the vocabulary selection problem, taking the power of individual words as a proxy for their importance [20].

Language has sentiment, typically either positive, negative, or neutral. Furthermore, individual words have "power" in determining the sentiment of text. For example, a review that reads "The food is delicious." conveys a positive sentiment, stemming from the word "delicious." We can thus view sentiment analysis as a voting game in which each word is a voter contributing to overall sentiment. Moreover, because the structure of grammar is inherently hierarchical, we model word importance as a hierarchical voting game.

Our approach is only approximate, however, as sentiment is not perfectly compositional. Thus, we obtain only an approximation of BPI, albeit one that can be calculated much faster than the exact value. We compare our approximate values to exact values on small problem instances, and conclude that our approximations are acceptable given the large speedup provided by our algorithms. The run time on sentences of different lengths can be seen in Figure 1a. MEBPI is slower than MBPI, but is more faithful to BPI.

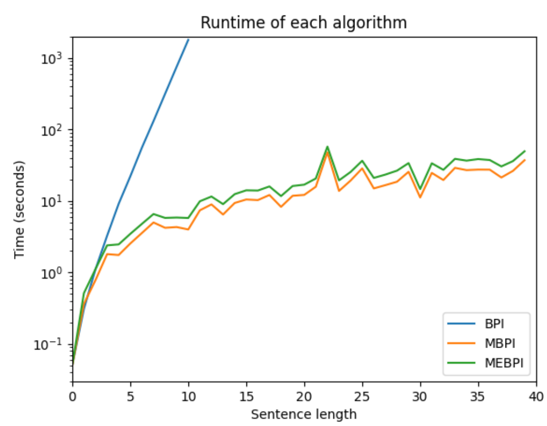

(a) The run time to compute BPI, MBPI, and MEBPI as sentence length increases. BPI quickly becomes intractable, while MEBPI is on par with MBPI, despite its much broader applicability.

4 ACKNOWLEDGMENTS

Denizalp Goktas was supported by a JP Morgan AI fellowship.

REFERENCES

[1] Amina Adadi and Mohammed Berrada. 2018. Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI). *IEEE Access* 6 (2018), 52138–52160. https://doi.org/10.1109/ACCESS.2018.2870052

[2] Fuad Aleskerov, Valeriy Kalyagin, and Kirill Pogorelskiy. 2008. Actual voting power of the IMF members based on their political-economic integration. *Annals of Economics and Statistics* 48, 9 (2008), 1554–1569. https://doi.org/10.1016/j.mcm.2008.05.020

[3] E. Algaba, J. M. Bilbao, and J. R. Fernandez. 2007. The distribution of power in the European constitution. *European Journal of Operational Research* 176, 3 (2007), 1752–1766. https://doi.org/10.1016/j.ejor.2005.12.002

[4] Carl A. Auerbach. 1964. The Reapportionment Cases: One Person, One Vote-One Vote, One Value. *The Supreme Court Review* 1964 (1964), 1–87. http://www.jstor.org/stable/3108693

[5] John F. Banzhaf. 1965. Weighted Voting Doesn't Work: A Mathematical Analysis. *Rutgers Law Review* 19 (1965), 317.

[6] Fabrice Barthelemy and Mathieu Martin. 2011. A comparison between the methods of apportionment using power indices: the case of the US presidential elections. *Annals of Economics and Statistics* 101 (2011), 87–106. https://doi.org/10.2307/41615475

[7] Wenhu Chen, Yu Su, Yilin Shen, Zhiyu Chen, Xifeng Yan, and William Yang Wang. 2019. How Large a Vocabulary Does Text Classification Need? A Variational Approach to Vocabulary Selection. *CoRR* abs/1902.10339 (2019). arXiv:1902.10339 http://arxiv.org/abs/1902.10339

[8] T.E.S. Raghavan Chih-Ru Hsiao. 1993. Shapley Value for Multichoice Cooperative Games. *Games and Economic Behavior* 5 (1993), 240–256. Issue 2. https://www.sciencedirect.com/science/article/pii/S0899825683710146

[9] Moshe Felsenthal, Dan S.; Machover (Ed.). 1998. *The Measurement of Voting Power Theory and Practice, Problems and Paradoxes*. Edward Elgar, Haifa, Israel.

[10] Imre Ferto, Laszlo A Koczy, Attila Kovacs, and Balazs R Sziklai. 2020. The power ranking of the members of the Agricultural Committee of the European Parliament. *European Review of Agricultural Economics* 47, 5 (2020), 1897–1919. https://doi.org/10.1093/erae/jbaa011

[11] Sharad Goel, Marc Meredith, Michael Morse, David Rothschild, and Houshmand Shirani-Mehr. 2020. One Person, One Vote: Estimating the Prevalence of Double Voting in U.S. Presidential Elections. *American Political Science Review* 114, 2 (2020), 456–469. https://doi.org/10.1017/S000305541900087X

[12] Werner Kirsch. 2016. Brexit and the Distribution of Power in the Council of the EU. *CEPS* (2016). https://www.ceps.eu/ceps-publications/brexit-and-distribution-power-council-eu/.

[13] Roland Kirstein. 2010. Volkswagen vs. Porsche: a power-index analysis. *International Journal of Corporate Governance* 2, 1 (2010), 1–20.

[14] Bettina Klinz and Gerhard J. Woeginger. 2005. Faster algorithms for computing power indices in weighted voting games. *Mathematical Social Sciences* 49, 1 (2005), 111–116. https://doi.org/10.1016/j.mathsocsci.2004.06.002

[15] Bogdan Kulynych and Carmela Tronscoso. 2017. Feature importance scores and lossless feature pruning using Banzhaf power indices. *NIPS Symposium on Interpretable Machine Learning* (2017). https://arxiv.org/abs/1711.04992

[16] Dennis Leech. 1988. The Relationship Between Shareholding Concentration and Shareholder Voting Power in British Companies: A Study of the Application of Power Indices for Simple Games. *Management Science* 34, 4 (1988), 509–527. https://doi.org/10.1287/mnsc.34.4.509

[17] Gurvan L'Hostis, David Grangier, and Michael Auli. 2016. Vocabulary Selection Strategies for Neural Machine Translation. *CoRR* abs/1610.00072 (2016). arXiv:1610.00072 http://arxiv.org/abs/1610.00072

[18] Ministere De L'interior. 2023. DRZAVNI SVET–THE NATIONAL COUNCIL OF SLOVENIA. http://www.senat.fr/senatsdumonde/english/slovenie.html/.

[19] N.R. Miller. 2013. A Priori Voting Power and the US Electoral College. *Power, Voting, and Voting Power* (03 2013). https://doi.org/10.1007/978-3-642-35929-3_22

[20] Roma Patel, Marta Garnelo, Ian M. Gemp, Chris Dyer, and Yoram Bachrach. 2021. Game-theoretic Vocabulary Selection via the Shapley Value and Banzhaf Index. In *NAACL-HLT*. Association for Computational Linguistics: Human Language Technologies, 2789–2798. https://doi.org/10.18653/v1/2021.naacl-main.223

[21] L.S. Penrose. 1946. The Elementary Statistics of Majority Voting. *Journal of the Royal Statistical Society* 109 (1946), 53–57. Issue 1. https://www.jstor.org/stable/2981392?origin=crossref

[22] Amnon Rapoport and Esther Golan. 1985. Assessment of Political Power in the Israeli Knesset. *American Political Science Review* 79, 3 (1985), 673–692. https://doi.org/10.2307/1956837

[23] Wojciech Samek, Thomas Wiegand, and Klaus-Robert Müller. 2017. Explainable Artificial Intelligence: Understanding, Visualizing and Interpreting Deep Learning Models. *CoRR* abs/1708.08296 (2017). arXiv:1708.08296 http://arxiv.org/abs/1708.08296

[24] J Antonio Seijas-Marcias. 2019. Power Index of Finnish Parties: Evolution of the Parliament System. *Finnish-German Yearbook of Political Economy* 41, 61 (2019), 41–61.

[25] Rico Sennrich, Barry Haddow, and Alexandra Birch. 2015. Neural Machine Translation of Rare Words with Subword Units. *CoRR* abs/1508.07909 (2015). arXiv:1508.07909 http://arxiv.org/abs/1508.07909

[26] Xing Shi and Kevin Knight. 2017. Speeding Up Neural Machine Translation Decoding by Shrinking Run-time Vocabulary. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Association for Computational Linguistics, Vancouver, Canada, 574–579. https://doi.org/10.18653/v1/P17-2091

[27] William S. Taylor, Alan D.; Zwicker. 1999. *Simple Games: Desirability Relations, Trading, Pseudoweightings*. Princeton University Press. http://www.jstor.org/stable/j.ctv18zhdsm

[28] H. Peyton Young and Shmuel Zamir (Eds.). 2015. *Handbook of Game Theory with Economic Applications* (1 ed.). Vol. 4. Elsevier. 239–326 pages. https://EconPapers.repec.org/RePEc:eee:gamhes:4