Distributive and Temporal Fairness in Algorithmic Collective Decision-Making

Doctoral Consortium

Nicholas Teh
University of Oxford
United Kingdom
nicholas.teh@cs.ox.ac.uk

ABSTRACT

From dividing parliamentary seats after a national election, to scheduling conference activities for an international AI conference, or deciding how to split public budget for city-wide projects, numerous real-life scenarios necessitates a group of individuals collectively reaching a desirable outcome through a preference aggregation process. In recent years, algorithms have been deployed in many scenarios to aid humans in such collective decision-making processes, with the goal of achieving fair outcomes efficiently. My work looks at the design and analysis of algorithms for various collective decision-making settings, including (i) indivisible resource allocation in the presence of strategic agents with different entitlements, (ii) multiwinner elections with temporal considerations, and (iii) the division of time and money when agents have cardinal preferences.

KEYWORDS

Fair Allocation, Temporal Voting, Computational Social Choice

ACM Reference Format:
Nicholas Teh. 2024. Distributive and Temporal Fairness in Algorithmic Collective Decision-Making: Doctoral Consortium. 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 INTRODUCTION

A national broadcaster that needs to decide on their prime time show schedule. A company that wishes to donate a portion of their proceeds to a single charity that is chosen annually. A recently-elected government that needs to allocate manpower among various ministries. Numerous real-world scenarios call for the need to make consequential decisions—the decision made will directly impact the entity's organizational goals (e.g., profit, efficient governance, etc.). Hence, there is a need to make an informed decision that will satisfy its user base—be it the viewers, customers, or the electorate. To facilitate this process, the entity may wish to elicit the preferences of its users, and attempt to make a *fair* and *efficient* decision.

In the presence of vast amounts of preference data, recent years have seen algorithms being deployed in numerous application areas to aid society in collective decision-making. This presents opportunities for research into various application areas of algorithmic decision-making, addressing overarching questions such as (1) what does it mean to be fair, (2) how can we achieve these fair outcomes, and (3) can we compute such outcomes efficiently?

In this paper, I will highlight several areas of collective decision-making that my research has largely focused on, providing motivating applications, my specific contributions to the area, and the open questions that remain for future work.

2 WEIGHTED FAIR DIVISION

The problem of fairly allocating indivisible resources to a set of agents has been a longstanding one at the intersection of computer science and economics. In 1948, Steinhaus [27] developed a mathematical framework for the systematic study of fair division. This led to a myriad of various works into the design and analysis of fair and efficient algorithms for various scenarios involving the allocation of resources [7, 24].

However, most of the works in this area largely assumed agents have equal entitlements to the resource. While this may capture applications such as inheritance claims or divorce settlements, it is not well-suited for settings where agents may have different entitlements to the resource. For instance, when considering parties in a coalition government, it is reasonable to assume that parties who won more seats or are pivotal to the formation of the government be afforded more parliamentary seats compared to smaller parties. Recent works on *weighted fair division* thus assume agents have weights (which represent its entitlements) [10–12]. Weighted fairness can also be adapted for upholding *group fairness*, where the weights represent the number of agents in each group [26].

Both in the weighted and unweighted case, numerous works study axiomatic properties that guarantee fairness (e.g., envy-freeness, maximin share), efficiency (e.g., Pareto-optimality), monotonicity (e.g., resource-, population-monotonicity) and non-manipulability (e.g., strategyproofness); as well as algorithms and rules that satisfy a combinations of these properties.

In Suksompong and Teh [28], we study an extension of the popular *maximum Nash welfare* rule (which maximizes the geometric product of agents' utilities, and is known for its highly desirable properties in the unweighted setting [9, 21]) to the weighted case with respect to monotonicity and non-manipulability properties, when agents have binary valuations. In a follow up work, Suksompong and Teh [29] extended these guarantees to the larger class of *weighted additive welfarist rules*, and for the more general class of *matroid-rank* (or binary submodular) valuations. In Montanari et al. [23], we defined families of envy-based fairness axioms for weighted fair division in the submodular setting, and proposed algorithms that achieve these properties.

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).

Our work provides a first step into considering weighted fair division beyond additive valuation functions. One can consider other valuation classes, such as supermodular or subadditive functions (which capture other applications).

3 TEMPORAL FAIRNESS IN VOTING

In the traditional multiwinner voting model, we have a set of voters with preferences over a set of candidates. The goal is then to select a subset of candidates that is *excellent* in quality, *diverse* in attributes, and *representative* of the voter population. This model captures numerous settings, from parliamentary elections in democratic systems to product placement in online shopping platforms.

However, most works in the area focus on a single-round election, even though numerous applications in real-world are more appropriately modelled with multi-round elections. To address this, there are several strands of work in voting—namely *perpetual voting*, *sequential committee elections*, and *scheduling*—that incorporate some temporal aspect in their model. However, existing works within and across these topics are largely disjoint and there lacks a systematic way to tackle various problems associated with voting over time.

To address this, in Elkind et al. [17], we survey existing works and provide a unified framework to facilitate the systematic study of temporal fairness in multiwinner voting, discussing ways of augmenting the traditional multiwinner voting model with a temporal dimension. We highlight key challenges, consolidate existing bodies of work, and position them within our framework. We also identify gaps in the literature and directions for future research.

In Elkind et al. [18], we build on prior work looking into notions of representation (specifically, variants of *justified representation*) for the temporal setting [8, 13] and consider the complexity of verifying outcomes in temporal elections. We show that verifying proportional outcomes in temporal elections is strictly harder than in single-round elections. In Neoh and Teh [25], we look at a similar model, with a focus on the computational complexity corresponding to the decision problems of maximizing various welfare objectives, and show the general difficulty in doing so. We also highlight special cases that enable efficient computation.

In Elkind et al. [16], we consider a special setting of whereby the outcome is a sequence of unique candidates (i.e., a permutation of candidates), which can also be viewed as a scheduling problem. Each voter's preference is also a desired sequence of unique candidates. We study the complexity of maximizing various welfare objectives, providing NP-hardness results, parameterized complexity results, and efficient algorithms for restricted cases. In particular, for the special case where the number of agents is a constant, we provide a randomized algorithm that is capable of computing (if it exists) an outcome satisfying a broad class of fairness properties.

4 PORTIONING

In the *portioning* problem, the aim is to divide a resource among a set of candidates. Such a resource could be public funds, whereby a town council is tasked with allocating these funds to several public infrastructure projects. Residents have preferences over which projects they prefer, and the town council's goal is to ensure the funds are used in a fair and efficient way. This example also illustrates how *participatory budgeting* can be an application of portioning. On the other hand, one may wish to fairly divide *time*. For instance, a conference organizer may be deciding on the proportion of time at the conference that is to be allocated for research talks, invited speakers, tutorials, workshops, and social activities.

Most prior works on portioning usually assume approval ballots [5, 6, 15] or ranked preferences [3]. However, these preference formats may not be adequate or expressive enough in reflecting agents' preferences for the aggregation process to be meaningful. Thus, one can consider agents with *cardinal* preferences, as in Freeman et al. [20], which focused primarily on strategyproofness as a key desiderata. However, it is far from being the only desirable property in this setting.

In Elkind et al. [19], we conduct an extensive study of various aggregation rules and welfare-optimization with respect to numerous fairness and welfare axioms. We compare this against a rule proposed by Freeman et al. [20], and show that strategyproofness comes at a price. We also highlight the tradeoffs that exists between various other fairness axioms. Numerous avenues for future research remain, including considering other classes of aggregation rules, considering characterizations for subset(s) of the axioms, or looking at more general agent preference models.

5 OTHER TOPICS

Envy-Free House Allocation with Minimum Subsidy. House allocation refers to the special setting of the indivisible item allocation problem, where each agent receives exactly one good (or in this case, a house) [1, 2, 22, 30]. This is a classic economic problem with applications such as assigning workers to offices, or students to dormitory rooms. A common fairness notion studied in this setting is *envy-freeness*, but it may not always exist. In Choo et al. [14], we consider the problem of achieving an envy-free house allocation using *subsidies* (or money), which has been considered in the fair allocation literature, but not in the house allocation model. We show that computing an envy-free allocation with subsidy is NP-hard in general, but if the number of houses differs from the number of agents by an additive constant or if the agents have identical valuations, then such an allocation can be computed efficiently.

Better Collective Decisions via Uncertainty Reduction. Consider a city council that wishes to consult the population on the implementation of several proposals (e.g., building a park, increasing taxes, etc.). The voter community is able to decide on a range of binary issues by means of issue-by-issue majority voting. For each issue and each agent, one of the two options is objectively better than the other. However, agents may be confused about some of the issue, for instance, as they have incomplete information about the project due to finite resources (e.g., time) to do their due diligence on every issue. This may result in them voting for the option that is objectively worse for them. In Bulteau et al. [4], we formalize the abovementioned model and study the computational complexity of the decision problems faced by a benevolent external party in employing several tools to help society reach better collective decisions. Numerous research directions remain, including the study of more general proposal spaces or considering weighted proposals with different importance.

REFERENCES

[1] Atila Abdulkadiroğlu and Tayfun Sönmez. 1998. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica* 66, 3 (1998), 689–701.

[2] Atila Abdulkadiroğlu and Tayfun Sönmez. 1999. House Allocation with Existing Tenants. *Journal of Economic Theory* 88, 2 (1999), 233–260.

[3] Stéphane Airiau, Haris Aziz, Ioannis Caragiannis, Justin Kruger, Jérôme Lang, and Dominik Peters. 2023. Portioning Using Ordinal Preferences: Fairness and Efficiency. *Artificial Intelligence* 314 (2023), 103809.

[4] Shiri Alouf-Heffetz, Laurent Bulteau, Edith Elkind, Nimrod Talmon, and Nicholas Teh. 2022. Better Collective Decisions via Uncertainty Reduction. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*.

[5] Haris Aziz, Anna Bogomolnaia, and Hervé Moulin. 2020. Fair Mixing: the Case of Dichotomous Preferences. *ACM Transactions on Economics and Computation* 8, 4 (2020), 18:1–18:27.

[6] Anna Bogomolnaia, Hervé Moulin, and Richard Stong. 2005. Collective Choice under Dichotomous Preferences. *Journal of Economic Theory* 122, 2 (2005), 165–184.

[7] Steven Brams and Alan Taylor. 1996. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press.

[8] Laurent Bulteau, Noam Hazon, Rutvik Page, Ariel Rosenfeld, and Nimrod Talmon. 2021. Justified Representation for Perpetual Voting. *IEEE Access* 9 (2021), 96598–96612.

[9] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2019. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation* 7, 3 (2019), 12:1–12:32.

[10] Mithun Chakraborty, Ayumi Igarashi, Warut Suksompong, and Yair Zick. 2021. Weighted Envy-Freeness in Indivisible Item Allocation. *ACM Transactions on Economics and Computation* 9, 3 (2021), 18:1–18:39.

[11] Mithun Chakraborty, Ulrike Schmidt-Kraepelin, and Warut Suksompong. 2021. Picking sequences and monotonicity in weighted fair division. *Artificial Intelligence* 301 (2021), 103578.

[12] Mithun Chakraborty, Erel Segal-Halevi, and Warut Suksompong. 2022. Weighted fairness notions for indivisible items revisited. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*. 4949–4956.

[13] Nikhil Chandak, Shashwat Goel, and Dominik Peters. 2024. Proportional Aggregation of Preferences for Sequential Decision Making. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*.

[14] Davin Choo, Yan Hao Ling, Warut Suksompong, Nicholas Teh, and Jian Zhang. 2024. Envy-Free House Allocation with Minimum Subsidy. *Operations Research Letters* (2024), 107103.

[15] Conal Duddy. 2015. Fair Sharing under Dichotomous Preferences. *Mathematical Social Sciences* 73 (2015), 1–5.

[16] Edith Elkind, Sonja Kraiczy, and Nicholas Teh. 2022. Fairness in Temporal Slot Assignment. In *Proceedings of the 15th International Symposium on Algorithmic Game Theory (SAGT)*. 490–507.

[17] Edith Elkind, Svetlana Obraztsova, and Nicholas Teh. 2024. Temporal Fairness in Multiwinner Voting. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*.

[18] Edith Elkind, Svetlana Obraztsova, and Nicholas Teh. 2024. Verifying Proportionality in Temporal Voting. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*.

[19] Edith Elkind, Warut Suksompong, and Nicholas Teh. 2023. Settling the Score: Portioning with Cardinal Preferences. In *Proceedings of the 26th European Conference on Artificial Intelligence (ECAI)*. 621–628.

[20] Rupert Freeman, David Pennock, Dominik Peters, and Jennifer Wortman Vaughan. 2021. Truthful Aggregation of Budget Proposals. *Journal of Economic Theory* 193 (2021), 105234.

[21] Daniel Halpern, Ariel D. Procaccia, Alexandros Psomas, and Nisarg Shah. 2020. Fair Division with Binary Valuations: One Rule to Rule Them All. In *Proceedings of the 16th International Conference on Web and Internet Economics (WINE)*. 370–383.

[22] Aanund Hylland and Richard Zeckhauser. 1979. The Efficient Allocation of Individuals to Positions. *Journal of Political Economy* 87, 2 (1979), 293–314.

[23] Luisa Montanari, Ulrike Schmidt-Kraepelin, Warut Suksompong, and Nicholas Teh. 2024. Weighted Envy-Freeness for Submodular Valuations. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*.

[24] Hervé Moulin. 2003. *Fair Division and Collective Welfare*. MIT Press.

[25] Tzeh Yuan Neoh and Nicholas Teh. 2024. Welfare Maximization in Perpetual Voting (Student Abstract). In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*.

[26] Jonathan Scarlett, Nicholas Teh, and Yair Zick. 2023. For One and All: Individual and Group Fairness in the Allocation of Indivisible Goods. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 2466–2468.

[27] Hugo Steinhaus. 1948. The problem of fair division. *Econometrica* 16, 1 (1948), 101–104.

[28] Warut Suksompong and Nicholas Teh. 2022. On Maximum Weighted Nash Welfare for Binary Valuations. *Mathematical Social Sciences* 117 (2022), 101–108.

[29] Warut Suksompong and Nicholas Teh. 2023. Weighted Fair Division with Matroid-Rank Valuations: Monotonicity and Strategyproofness. *Mathematical Social Sciences* 126 (2023), 48–59.

[30] Lin Zhou. 1990. On a conjecture by Gale about one-sided matching problems. *Journal of Economic Theory* 52, 1 (1990), 123–135.