Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence (IJCAI-23)
Doctoral Consortium Track

Fairness and Stability in Complex Domains

Julian Chingoma
Institute of Logic, Language and Computation, University of Amsterdam
j.z.chingoma@uva.nl

Abstract

Fairness and stability are normative concepts that have been investigated for many social choice domains. Recently, increasing attention has fallen on richer, and more complex, settings and we look to develop, and study in depth, these fairness and stability notions in a variety of such complex domains.

1 Introduction

With the many instances of collective decision-making tasks that occur in the real world, the many efforts to determine the best ways to make these collective choices have spawned the field of social choice theory [Arrow *et al.*, 2002]. Social choice theory is dedicated to studying the methods to aggregate the opinions of the group with the classical example being the voting in elections to elect a single winning candidate.

Beyond the case of single-winner elections, an area that has garnered significant attention from researchers is that of multiwinner voting (MWV) [Faliszewski *et al.*, 2017; Lackner and Skowron, 2023] where instead of there being a single election winner, one models scenarios where a committee with multiple candidates is to be selected. Applications include that of the apportionment problem, where a fixed number of parliamentary seats are to be distributed to political parties, or in shortlisting tasks where a fixed, or possibly variable, number of candidates are to be selected. A substantial part of the work done in MWV has been dedicated to studying *fairness* and *stability* properties for aggregation methods (I will briefly detail these notions in Section 3). In my work, I will investigate the extent to which fairness and stability properties can developed for certain *complex domains*.

What are these *complex* domains? In recent years, there has been an uptick in research done on a richer variant of the MWV model, that being the model of *participatory budgeting (PB)* [Lackner and Skowron, 2023]. PB is the process where citizens collectively decide on the public projects, each project coming with a cost, that are to be implemented whilst respecting a budget limit. As a generalisation of the MWV problem, where each candidate would come with a cost, we say that PB represents a more complex domain (with respect to MWV at least). We explore *other* such complex domains. Consider the following examples: selecting a committee, of either a fixed or variable size, where there are constraints over the candidates that may be selected; or aggregating opinions over yes/no decisions on binary issues that also have constraints on the yes/no choices in the outcome. We identify these, amongst others, as domains where desirable normative properties should be established, as they represent in a similar vein to PB, more complex variants of the MWV problem.

In this abstract, I highlight the main approaches, both conceptually and methodologically, that my work will consist of.

2 Using Judgment Aggregation

As we look to establish notions that are well-studied in domains such as MWV and PB, within other complex domains, we deem it important to understand how various domains relate to each other. In this regard, we wish to investigate the extent that one domain may model another, using the model of *Judgment Aggregation (JA)*. JA is a general framework fitted with logical constraints that allows one to model a wide variety of collective choice problems within it [Endriss, 2016]. Specifically, in an approach that extends work done in embedding single-winner voting rules in the JA model [Endriss, 2018], we demonstrated that some of the classical MWV rules can be embedded within the JA model (in particular, the JA model which differentiates between the constraints on the allowed inputs from agents, and the feasible outputs of the aggregation method) [Chingoma *et al.*, 2022].

3 Axiomatic Study

This section details the main notions that I will study.

3.1 Fairness

Fairness, as a feature of an aggregation process, is a property that is difficult to argue against. However, there are various ways to define what is fair, and careful thought is required to develop notions that are both natural, and achievable. This is particularly true when one moves away from more 'simple' settings such as the apportionment model.

Proportionality is one notion that has drawn the focus of many researchers. Intuitively, a group of agents that are in agreement, in some sense, and represent an α fraction of the total population, are entitled to influence, loosely speaking,

an α fraction of the collective outcome. Ensuring proportionally representative collective outcomes has been tackled in settings of increasing complexity: from the apportionment problem, to the more general problem of MWV, and also a step further, in the even richer PB setting [Lackner and Skowron, 2023]. So, proportionality is the main fairness notion that we study as we look to import it into other complex domains.

To begin with, we initiated our study of proportionality within the JA framework [Chingoma *et al.*, 2022]. Here, we adopted an interpretation of a MWV problem with a variable number of winners. Proportional representation in this model has seen recent focus [Freeman *et al.*, 2020; Skowron and Górecki, 2022], however, our work also explored the possibilities of proportional representation with the presence of logical constraints over candidates. In particular, our work saw the design of JA rules that satisfy our proportionality property for this domain, to varying degrees.

While the main focus will be on proportionality, we do not dismiss the need to consider, and systematically study, other fairness concepts. These may come from those already seen in the literature such as *diversity* [Lackner and Skowron, 2023; Faliszewski *et al.*, 2017] (ensuring as many agents as possible obtain some kind of representation within the outcome), or we may develop novel fairness notions, tailored to each domain.

3.2 Stability

Another concept that has been commonly studied in social choice theory is that of *stability*. Inspired by the game-theoretic notion of stability, it has been applied in many social choice contexts as a measure of quality of a collective outcome. Intuitively, a stable outcome is one where no group of agents has an incentive to deviate to another outcome, subject to restrictions on the permitted deviations.

A notable instantiation of this notion is that of *weak Gehrlein stability* [Faliszewski *et al.*, 2017] which deems an unstable committee to be one where a majority of the agents prefer some unelected candiate to a current committee member. During the first steps in studying stability, we have shown that through the use of the aforementioned JA model, one can simulate known MWV rules (from the ordinal setting where agents rank the candidates), as well as gain access to a class of novel MWV rules, that satisfy this property of weak Gehrlein stability [Chingoma *et al.*, 2022]. Stability has also been studied in the context of proportional representation for MWV models where agent groups may deviate from outcomes where they receive insufficient representation [Lackner and Skowron, 2023; Faliszewski *et al.*, 2017]. There is then room to further explore various stability notions within the complex domains.

4 Computational Efficiency

We do not only aim to develop aggregation methods that produce fair and stable outcomes. An important consideration in our work is the potential for practical implementation of these methods on real-world instances of the complex domains. So, finding *efficient* aggregation methods is a task at the core of our work. This is even more relevant within complex domains where, generally speaking, aggregation methods are known to be computationally intractable, such as in the JA model [Endriss, 2016].

5 Conclusion and Future Work

The broad theme of my research, transferring the notions of fairness and stability to more complex domains, has been detailed in this paper. Beyond establishing these properties in more complex domains, there are many avenues for future work. Another path is extending our work on Gehrlein stability in the JA model, to see which other stability notions can be embedded within JA, as this may shed more light on the nature of stable rules. Also of interest is developing our initial study of proportionality in JA in a search for a fairness concept for the JA framework in its full generality.

References

[Arrow *et al.*, 2002] Kenneth J. Arrow, Amartya Sen, and Kotaro Suzumura. *Handbook of Social Choice and Welfare*, volume 1. Elsevier, 2002.

[Chingoma *et al.*, 2022] Julian Chingoma, Ulle Endriss, and Ronald de Haan. Simulating multiwinner voting rules in judgment aggregation. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2022)*. IFAAMAS, May 2022.

[Endriss, 2016] Ulle Endriss. Judgment aggregation. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 17. Cambridge University Press, 2016.

[Endriss, 2018] Ulle Endriss. Judgment aggregation with rationality and feasibility constraints. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2018)*. IFAAMAS, July 2018.

[Faliszewski *et al.*, 2017] Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. Multiwinner voting: A new challenge for social choice theory. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 2, pages 27–47. AI Access, 2017.

[Freeman *et al.*, 2020] Rupert Freeman, Anson Kahng, and David M. Pennock. Proportionality in approval-based elections with a variable number of winners. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI-2020)*, pages 132–138. AAAI Press, 2020.

[Lackner and Skowron, 2023] Martin Lackner and Piotr Skowron. *Multi-Winner Voting with Approval Preferences*. Springer Briefs in Intelligent Systems. Springer, 2023.

[Skowron and Górecki, 2022] Piotr Skowron and Adrian Górecki. Proportional public decisions. In *36th AAAI Conference on Artificial Intelligence, AAAI 2022*, pages 5191–5198. AAAI Press, 2022.