

# Towards Fairness in Collective Decision Making: Constraints and Approximations

Doctoral Consortium

Yuhang Guo  
UNSW Sydney  
Sydney, Australia  
yuhang.guo2@unsw.edu.au

## ABSTRACT

Collective decision making is a central problem in multi-agent systems, where agents hold heterogeneous preferences over a set of outcomes and a collective decision must be reached. Ensuring fairness in such settings has long been a fundamental concern, particularly in realistic environments characterized by diverse constraints and practical requirements. My doctoral research investigates the design and analysis of fair collective decision making mechanisms under these constraints, with a particular emphasis on provable approximation guarantees.

## KEYWORDS

Fairness, Facility Location, Mechanism Design, Proportional Representation, Committee Selection, Approximation Algorithms

### ACM Reference Format:

Yuhang Guo. 2026. Towards Fairness in Collective Decision Making: Constraints and Approximations: Doctoral Consortium. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 3 pages. <https://doi.org/10.65109/OEOU3703>

## 1 INTRODUCTION

Picture a research group planning a New Year celebration. the short-list is tempting: barbecue, buffet, seafood, etc. Each member arrives with a different ranking shaped by taste, diet, budget, and mood. The task sounds mundane, yet it immediately poses a principled question: how should the collective aggregate individual preferences into one choice so that the outcome is, in a meaningful sense, as *satisfying* as possible for everyone? This tension between individual desiderata and joint decision is the essence of collective decision making. It appears at every scale, from families coordinating weekend plans, to committees allocating limited resources, to institutions and nations negotiating policies with far reaching consequences. What makes these scenarios nontrivial is not merely disagreement, but the need to reconcile diverse preferences under constraints while maintaining legitimacy and robustness.

Collective decision making is therefore a central theme in multi-agent systems, where one studies how a set of agents, each with its own preference structure, can reach a joint decision that optimizes individual or social objectives. Within algorithmic game theory and

computational social choice, this umbrella covers many canonical problems. Voting and elections seek procedures that identify desirable alternatives from community preferences [18, 33]. Resource allocation investigates how to match individuals with goods or opportunities efficiently and fairly [4]. More broadly, these settings demand methods that are not only normatively compelling, but also computationally well founded. Over the past few decades, fairness has emerged as a first class concern in the design and analysis of such mechanisms. In these literature (see e.g. [3]), fairness is typically captured through axioms that articulate what outcomes should guarantee, for instance, proportion representation [5, 30], envy freeness [2, 21], or protection against systematic disadvantage. This axiomatic lens immediately crystallizes algorithmic challenges, whether outcomes meeting the prescribed criteria exist and how to compute them efficiently? More specifically, the proliferation of fairness notions prompts a deeper set of design questions. Which notion is most appropriate for a given task and application context? To what extent are these notions achievable in practice, and how do they interact with other desiderata such as efficiency and incentive compatibility? Finally, even when fair outcomes are well defined, can they be computed in polynomial time, or do they entail inherent computational intractability? In this paper, I will present several threads of my Ph.D. research on collective decision making with fairness consideration, highlighting the real-world motivations, algorithmic foundations and contributions, and potential open directions for future work.

## 2 FAIR FACILITY LOCATION GAMES

Facility location problem is the task of placing facilities in metric space to minimize total service cost or maximize coverage for a set of demand points. A substantial optimization literature studies fairness in facility location problems, where equity based measures such as distance standard deviation [24] and the Gini coefficient [23] are considered. Within algorithmic mechanism design, Procaccia and Tennenholtz [31] initiated the study of individual fairness via the maximum cost objective by designing strategyproof mechanisms with approximation guarantees. Subsequent work proposed alternative objectives. Cai et al. [11], Chen et al. [13] considered minimax envy, defined as the maximum pairwise distance gap. Ding et al. [15], Liu et al. [22] studied the envy ratio objective, comparing the best and worst agent utilities. Walsh [34] studied the Gini index objective. Fairness has also been formulated at the level of agent groups. Li et al. [20], Zhou et al. [35] considered maximum total cost and maximum average cost, capturing the worst case burden over predefined groups. Aziz et al. [7], Lam et al. [19] explored



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

*Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems ([www.ifaamas.org](http://www.ifaamas.org)). <https://doi.org/10.65109/OEOU3703>

proportional fairness with guarantees for endogenously defined groups that scale with group size.

*(Individual Fairness).* In our recent paper [6], we study the facility location mechanism design problem on a line for the envy ratio objective under the learning-augmentation framework. Facility location mechanism design with prediction has flourished in recent years. Agrawal et al. [1] initiated the study of deterministic mechanism which takes a prediction of optimal facility location as extra input, where they show a best-of-both-worlds results, designing mechanisms which perform (in terms of their approximation ratio) *consistently* well when an accurate prediction is provided, and are *robust* to entirely inaccurate predictions. In contrast to prior work, we focus on the envy-ratio objective. We establish tight approximation guarantees for learning-augmented deterministic mechanisms. For randomized mechanisms, we first resolve the open question raised by Ding et al. [15] by designing a novel randomized mechanism without prediction for the envy ratio objective, which breaks the bound of 2, and further develop learning-augmented randomized mechanism to achieve improved performance guarantees.

*(Group Fairness).* Our recent paper [16] study the facility location game in the line space, where agents belong to one or multiple groups under the general group fairness objective of *maximum group effect* which was first introduced by Marsh and Schilling [23]. The group effect aggregates, via group specific weights, either total or maximum distance to the nearest facility. We show this objective formulation subsumes the prominent objectives, including the social cost, maximum cost, as well as maximum total and average group cost. To minimize the maximum group effect, we propose unified mechanisms that achieve tight approximation guarantee for this general objective. Notably, our proposed *Balanced Mechanism* closes the approximation gap for the group fairness objectives of maximum total group cost and maximum average group cost studied by Zhou et al. [35]. A natural and challenging open problem is to determine the optimal approximation guarantees achievable by randomized mechanisms for the group fairness objective.

### 3 PROPORTIONAL FAIRNESS IN RICH DOMAIN

Proportional representation [26] is a foundational principle in electoral system design, requiring that popularly supported subcollectives such as political parties within an electorate be represented in an elected body in proportion to their level of support. It has emerged as a prominent fairness notion and has been studied extensively in committee selection [5], participatory budgeting [9], and clustering [12], among other settings. Bullinger et al. [10] considered proportional fairness in the practical problem of transit stop placement, where each agent travels regularly along a fixed route between two endpoints and seeks to minimize commuting cost. They first studied this model on the line metric and consider two objectives, Justified Representation (JR) and the core. JR was originally proposed by Aziz et al. [5] for approval based committee voting, which requires that any sufficiently large cohesive set of voters that commonly approves at least one candidate must receive at least one representative in the selected committee. In the context of transit stop placement, JR informally stipulates that no sufficiently large coalition, deemed entitled to a pair of stops, should be

able to profitably deviate to such a pair compared to the solution returned by the algorithm. Core strengthens this requirement by prohibiting any coalition from deviating to an alternative set of stops of size proportional to its population share such that every member strictly reduces their commute cost.

In our recent working paper Aziz et al. [8], we extend this line of research from the line metric to general metric spaces and connect transit stop placement to proportional fair clustering [12]. By uncovering a structural correspondence between stop placement and clustering, we show that the core need not exist in this setting. We then adapt the Greedy Capture algorithm [12] to the stop placement problem, obtaining constant factor approximations with respect to both JR and the core. To narrow the gap in the JR approximation ratio, we propose a new procedure, the Expanding Cost Algorithm (ECA), which achieves a strictly better JR guarantee. Our results reveal an inherent tension: Greedy Capture and ECA induce different trade-offs between approximating JR and approximating the core. Motivated by this, we consider a hybrid algorithm parameterized by  $\lambda \in [0, 1]$ , which controls how candidate stops are selected by interpolating between the Greedy Capture logic and the ECA logic. We show that it yields parameterized trade offs between approximation guarantees for JR and for the core.

Several interesting open problems remain. First, can the gap in the JR approximation ratio be closed. Second, since core approximation is relaxed by both coalition size and individual cost, can we obtain a complete characterization of achievable approximations in both dimensions. Third, can one incorporate efficiency and proportional fairness simultaneously, for example by minimizing total cost subject to proportional fairness constraints.

*Further Research Questions.* As noted by Micha [25], Peters [29], proportional fairness can be formalized and applied not only in voting frameworks but also across a range of AI domains where it arises naturally. For instance, in AI alignment, one may ask whether it is possible to aggregate or mix the outputs of multiple LLMs in a manner that incorporates proportional representation principles, so that the resulting outcome reflects the diverse preferences encoded in these models. In real-world applications, including sortition and federated learning, a natural question is how to select participants proportionally while respecting feasibility and resource constraints [32].

Another intriguing direction concerns veto proportionality [27], which captures the idea that a candidate should not be selected if it is vetoed by a sufficiently large group of voters. In the setting of committee selection with ordinal preferences, several fundamental questions remain open: how should veto proportionality be formally defined and efficiently computed? What is its relationship to standard proportionality axioms?

Moreover, developing randomized algorithms for committee selection that approximate proportional fairness across different environments is a promising direction for future research [14, 17, 28].

### ACKNOWLEDGMENTS

I would like to thank my advisor, Haris Aziz, for his guidance and support. I am also grateful to all of my collaborators for their instrumental contribution.

REFERENCES

[1] Priyank Agrawal, Eric Balkanski, Vasilis Gkatzelis, Tingting Ou, and Xizhi Tan. 2022. Learning-Augmented Mechanism Design: Leveraging Predictions for Facility Location. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*. 497–528.

[2] Georgios Amanatidis, Haris Aziz, Georgios Birmpas, Aris Filos-Ratsikas, Bo Li, Hervé Moulin, Alexandros A Voudouris, and Xiaowei Wu. 2023. Fair division of indivisible goods: Recent progress and open questions. *Artificial Intelligence* 322 (2023), 103965.

[3] Haris Aziz. 2020. Developments in multi-agent fair allocation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 13563–13568.

[4] Haris Aziz, Péter Biró, and Makoto Yokoo. 2022. Matching market design with constraints. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 12308–12316.

[5] Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. 2017. Justified representation in approval-based committee voting. *Social Choice and Welfare* 48, 2 (2017), 461–485.

[6] Haris Aziz, Yuhang Guo, Alexander Lam, and Houyu Zhou. 2025. Learning-Augmented Facility Location Mechanisms for the Envy Ratio Objective. In *Proceedings of the 39th Annual Conference on Neural Information Processing Systems (NeurIPS)*. forthcoming.

[7] Haris Aziz, Alexander Lam, Barton E. Lee, and Toby Walsh. 2025. Proportionality-based fairness and strategyproofness in the facility location problem. *Journal of Mathematical Economics* 119 (2025), 103129.

[8] Haris Aziz, Gai Ling, Yuhang Guo, and Jeremy Vollen. 2026. *Fair Transit Stop Placement: A Clustering Perspective and Beyond*. Technical Report.

[9] Haris Aziz and Nisarg Shah. 2020. Participatory budgeting: Models and approaches. In *Pathways Between Social Science and Computational Social Science: Theories, Methods, and Interpretations*. Springer, 215–236.

[10] Martin Bullinger, Edith Elkind, and Mohamad Latifian. 2025. Towards Fair and Efficient Public Transportation: A Bus Stop Model. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 427–435.

[11] Qingpeng Cai, Aris Filos-Ratsikas, and Pingzhong Tang. 2016. Facility location with minimax envy. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*. 137–143.

[12] Xingyu Chen, Brandon Fain, Liang Lyu, and Kamesh Munagala. 2019. Proportionally Fair Clustering. In *Proceedings of the 36th International Conference on Machine Learning (ICML) (Proceedings of Machine Learning Research, Vol. 97)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.). PMLR, 1032–1041.

[13] Xin Chen, Qizhi Fang, Wenjing Liu, Yuan Ding, and Qingqin Nong. 2022. Strategyproof mechanisms for 2-facility location games with minimax envy. *Journal of Combinatorial Optimization* 43, 5 (2022), 1628–1644.

[14] Yu Cheng, Zhihao Jiang, Kamesh Munagala, and Kangning Wang. 2020. Group fairness in committee selection. *ACM Transactions on Economics and Computation (TEAC)* 8, 4 (2020), 1–18.

[15] Yuan Ding, Wenjing Liu, Xin Chen, Qizhi Fang, and Qingqin Nong. 2020. Facility location game with envy ratio. *Computers & Industrial Engineering* 148 (2020), 106710.

[16] Yuhang Guo and Houyu Zhou. 2026. Minimizing Inequity in Facility Location Games. In *Proceedings of the 40th Annual AAAI Conference on Artificial Intelligence (AAAI)*. forthcoming.

[17] Zhihao Jiang, Kamesh Munagala, and Kangning Wang. 2020. Approximately stable committee selection. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*. 463–472.

[18] Martin Lackner and Piotr Skowron. 2023. *Multi-winner voting with approval preferences*. Springer Nature.

[19] Alexander Lam, Haris Aziz, Bo Li, Fahimeh Ramezani, and Toby Walsh. 2024. Proportional fairness in obnoxious facility location. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 1075–1083.

[20] Jiaqian Li, Minming Li, and Hau Chan. 2024. Strategyproof Mechanisms for Group-Fair Obnoxious Facility Location Problems. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*. 9832–9839.

[21] Richard J Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce*. 125–131.

[22] Wenjing Liu, Yuan Ding, Xin Chen, Qizhi Fang, and Qingqin Nong. 2020. Multiple Facility Location Games with Envy Ratio. In *Proceedings of the 14th International Conference on Algorithmic Aspects in Information and Management (AAIM)*. 248–259.

[23] Michael T Marsh and David A Schilling. 1994. Equity measurement in facility location analysis: A review and framework. *European journal of operational research* 74, 1 (1994), 1–17.

[24] Donald M McAllister. 1976. Equity and efficiency in public facility location. *Geographical analysis* 8, 1 (1976), 47–63.

[25] Evi Micha. 2025. How Proportional Representation Can Shape Artificial Intelligence. *AI Magazine* 46, 4 (2025), e70044.

[26] John Stuart Mill. 1903. *Consideration of representative government*.

[27] Hervé Moulin. 1981. The proportional veto principle. *The Review of Economic Studies* 48, 3 (1981), 407–416.

[28] Kamesh Munagala, Yiheng Shen, Kangning Wang, and Zhiyi Wang. 2022. Approximate core for committee selection via multilinear extension and market clearing. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2229–2252.

[29] Dominik Peters. 2024. Proportional representation for artificial intelligence. In *Proceedings of the 27th European Conference on Artificial Intelligence (ECAI)*. 27–31.

[30] Dominik Peters, Grzegorz Pierczyński, and Piotr Skowron. 2021. Proportional participatory budgeting with additive utilities. *Advances in Neural Information Processing Systems* 34, 12726–12737.

[31] Ariel D Procaccia and Moshe Tennenholtz. 2013. Approximate mechanism design without money. *ACM Transactions on Economics and Computation (TEAC)* 1, 4 (2013), 1–26.

[32] Bhaskar Ray Chaudhury, Linyi Li, Mintong Kang, Bo Li, and Ruta Mehta. 2022. Fairness in federated learning via core-stability. *Proceedings of the 36th Conference on Advances in Neural Information Processing Systems (NeurIPS)* 35 (2022), 5738–5750.

[33] Amartya Sen. 1986. Social choice theory. *Handbook of mathematical economics* 3 (1986), 1073–1181.

[34] Toby Walsh. 2025. Equitable Mechanism Design for Facility Location. In *Proceedings of the 34th International Joint Conference on Artificial Intelligence (IJCAI)*. 275–283.

[35] Houyu Zhou, Minming Li, and Hau Chan. 2022. Strategyproof Mechanisms for Group-Fair Facility Location Problems. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence and the 25th European Conference on Artificial Intelligence (IJCAI)*. 613–619.