Matching Algorithms under Diversity-Based Reservations

Haris Aziz
UNSW Sydney
Sydney, Australia
haris.aziz@unsw.edu.au

Sean Morota Chu
UNSW Sydney
Sydney, Australia
seanmorotachu@gmail.com

Zhaohong Sun
Cyberagent Inc.
Tokyo, Japan
sunzhaohong1991@gmail.com

ABSTRACT

Selection under category or diversity constraints is a ubiquitous and widely-applicable problem that is encountered in immigration, school choice, hiring, and healthcare rationing. These diversity constraints are typically represented by minimum and maximum quotas on various categories or types. We undertake a detailed comparative study of applicant selection algorithms with respect to the diversity goals.

KEYWORDS

Matching; diversity constraints; affirmative action; school selection

ACM Reference Format:
Haris Aziz, Sean Morota Chu, and Zhaohong Sun. 2023. Matching Algorithms under Diversity-Based Reservations. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 3 pages.

1 INTRODUCTION

How should we hire job applicants when we want to take both the overall merit as well as requirements of various departments into account? How should we decide on student intake while considering both entrance test scores and target numbers of scholarships for different categories? How should we ration healthcare resources when patients can avail resources under various categories? Which applicants should be given an immigration slot when the government has targets for various categories? These fundamental and important questions constitute a recurring theme in allocation and selection decisions. We consider a natural mathematical model for the problem that captures the main features of many of the problems discussed above. Although various choice rules and algorithms for selecting agents have been proposed, there has been little work carefully comparing the relative performance of these algorithms, especially from an experimental methodology. In this paper, we undertake one of the first detailed experimental studies to understand how well the algorithms perform with respect to capturing the intended diversity goals as well as selecting the highest priority applicants. We also try to understand the tradeoffs between merit and diversity.

We consider a very widely studied model of selection under diversity constraints. Firstly, there is a baseline ordering over the applicants. The baseline ordering could be the merit ordering in the context of school admissions, or the need for treatment in the context of healthcare rationing. If no diversity constraints are present, the selection of agents is made only with respect to the baseline priority ordering. If the diversity constraints are additionally present, then both the priority ordering and the diversity constraints are used to make selection decisions.

The diversity constraints or goals are represented by imposing minimum and maximum quotas on each of the types. In particular, given one school c, there is a lower quota of $q^1_{c,t}$ for the number of slots taken by agents for type t and there is an upper quota of $q^2_{c,t}$ for the number of slots taken by agents for type t. In the line of literature (see, e.g., Ehlers et al. [10]) both lower and upper quotas are viewed as guidelines towards reaching diversity goals: firstly, fill up slots of those types whose minimum quotas have not been reached. As a secondary consideration, fill up slots of those types whose minimum quotas have been reached, but not their maximum quotas.

Another feature of our setting is that applicants can satisfy multiple types such as being extra talented or being from a disadvantaged group. Each applicant who is selected is assumed to count towards one of the types satisfied by them. Such a type could include a general public type. This way of accounting for representation has been referred to as the one-to-one convention, which is popular in Indian college admissions [16]. Since we are not only interested in which agents are selected but also in how many target numbers of spots corresponding to relevant types are filled up, the output for our problems is not just a set of selected agents. Instead, it is a matching that matches each student to some type that the student satisfies. Such a matching not only gives information about the set of selected agents who are matched but also gives a count of how many seats of each type are used.

In this paper, we examine the following problem.

> *In selection problems under minimum and maximum quota diversity goals, how do various algorithms perform with respect to satisfying diversity goals as well as merit?*

With respect to performance on merit, we will compare the outcomes of algorithms according to various objectives, including average rank, worst rank, and best rank. When considering diversity constraints captured by lower and upper diversity quotas, a natural question is how to gauge the level of diversity captured by a given set of applicants or a matching? A natural solution was provided by Aziz and Sun [7] who viewed each type t as two ranks of slots corresponding to lower and upper quotas. A set of agents provides *maximal diversity* if there is a matching that matches the agents to the types in such a way that the number of rank 1 slots is maximized and given that the number of rank 2 slots is maximized.

One of the first algorithms for the problem was presented by Ehlers et al. [10] who assumed that each applicant can satisfy at most one type. The algorithm takes a natural greedy approach to first fill up slots corresponding to rank 1 and then to rank 2. It can suitably be extended to the case where agents may have multiple types. We will use the natural extension as one of the main

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

algorithms whose performance we examine. We will refer to the algorithm as EHYY.

Another algorithm that we consider is the *horizontal choice rule* by Sönmez and Yenmez [17] that was designed to optimally filling up seats when there is a single rank of slots. We consider two versions of the rule of Sönmez and Yenmez [17]: SY1 optimizes the use of the first ranked slots and SY2 merges the first and second ranked slots and then optimizes the use of these slots.

Aziz and Sun [7] presented algorithms that achieve maximal diversity. We will refer to the algorithm as A-S. There are several other algorithms that have been proposed or are used in real-world systems. The goal of this paper is to undertake a comparative study of various algorithms for the problem and see how they fare in terms of maximal diversity. We check how various algorithms do in terms of filling up the first ranked slots. We also check how various algorithms do in filling up the first two ranks.

From the specification, the A-S already maximizes the use of rank 1 slots and given that, it maximizes the use of rank 2 slots. One of the goals of the paper is to understand the extent to which it performs in relation to other existing approaches. We will also compare the algorithms with two baseline algorithms that predominantly care about the priority of the agents rather than diversity concerns.

2 CONTRIBUTIONS

In this paper, we present several contributions. Firstly, we present a consistent specification of various algorithms for our setting with minimum and maximum quotas or equivalently rank 1 and rank 2 seats. Secondly, we perform one of the first experimental comparisons of prominent selection algorithms in achieving optimal diversity goals as well as average merit ranking of the agents. Next, we investigate the performance of prominent selection algorithms across a variety of different environments, thereby determining the environmental parameters affecting their performance.

Some of the conclusions from the experiments include the following. The total number of reserves and the selection capacity of a problem instance influence the performance of each algorithm. As the number of reserves relative to selection capacity increases, the performance of diversity based algorithms is reduced with respect to satisfying merit compared to matching algorithms that ignore reserves. When the total number of reserves is exceeded by the selection capacity, A-S and SY2 have equivalent performance, despite having different behaviour when total reserves exceed selection capacity. Overall, A-S is the best algorithm at fulfilling reserves across two ranks but performs worse in selecting for merit compared to SY1 and SY2, which are optimal for filling the first and first two ranks of reserves respectively. The performance of EHYY is close to optimality on average when satisfying the first rank reserves, but its worst case performance is reduced when selection capacity and the number of reserves increase.

We find that, due to the various different characteristics of each algorithm, there is a necessary tradeoff between achieving merit and diversity goals, and the choice of algorithm can help negotiate between these two goals for any specific problem instance. The full paper is available on the authors' websites.

3 RELATED WORK

The literature on matching under diversity and other distributional constraints is vast. We discuss work that is closely related to our problem.

Affirmative action in two-sided matching has been considered in early work on school choice Abdulkadiroğlu [1], Abdulkadiroğlu and Sönmez [2]. In many of the diversity models, each school puts a minimum quota on each type [10, 12–14]. Ehlers et al. [10] treated the quotas in a soft manner since hard constraints can lead to infeasibility. We pursue the same approach as well. In contrast to Ehlers et al. [10], we allow agents to have multiple types.

The issue of agents having multiple 'overlapping types' has been considered in recent papers and deployed applications in the past few years, including those in Brazil, Chile, Israel, and India (see, e.g., [3, 8, 9, 11, 15]). There are two ways to perform accounting when agents have multiple types [17]. In the *one-for-all* convention, an agent is viewed as taking slots for all the types that they satisfy [6, 11]. In the *one-for-one* convention, they take the slot of one of the types they satisfy. In this paper, we pursue the one-for-one convention. This convention has the 'more widespread interpretation' [17]. The one-for-one convention has been explicitly or implicitly considered in several recent papers [3–5, 8–10, 15]. Most of these approaches do not achieve diversity optimally. In contrast, Aziz and Sun [7] presented a rule that achieves diversity optimally. When there is only one rank of reserves or when there are no maximum quotas, Sönmez and Yenmez [17] presented a rule that also satisfies diversity optimally. We will consider two extensions of the algorithm for our model.

4 DISCUSSION

We have examined the effectiveness of prominent matching algorithms in satisfying a range of performance metrics across a variety of different instances. We find that there is a necessary tradeoff when balancing performance between priority and reserves, and this tradeoff can be negotiated through our choice of selection algorithm.

When we wish to optimise our matching toward fulfilling reserves across multiple ranks, the A-S algorithm will always provide the best solution while maintaining the highest possible priority of selected agents. However, if we wish to optimise across only one rank, SY1 and SY2 can provide a solution that can achieve this while outperforming A-S in terms of priority ranking. It also becomes clear that, for most instances where q_c is not high, reserve based matching algorithms provide highly different outcomes from priority-only algorithms such as POG and POS, creating further emphasis on the tradeoff between priority and reserve satisfaction.

Therefore, when selecting an algorithm to solve a problem, we must carefully consider the following points:

(1) Whether or not the problem requires optimisation for priority or reserves.
(2) The relative importance of filling reserves according to rank against the importance of maximising priority.
(3) The value of capacity q_c relative to the number of students $|S|$.
(4) The number of reserves available relative to capacity q_c.

REFERENCES

[1] A. Abdulkadiroğlu. 2005. College admissions with affirmative action. *International Journal of Game Theory* 33, 4 (2005), 535–549.

[2] A. Abdulkadiroğlu and T. Sönmez. 2003. School Choice: A Mechanism Design Approach. *American Economic Review* 93, 3 (2003), 729–747.

[3] O. Aygün and B. Turhan. 2016. Dynamic reserves in matching markets: Theory and applications. (2016).

[4] H. Aziz. 2023. Probabilistic Rationing with Categorized Priorities: Processing Reserves Fairly and Efficiently. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS).*

[5] H. Aziz and F. Brandl. 2021. Efficient, Fair, and Incentive-Compatible Healthcare Rationing. In *EC '21: The 22nd ACM Conference on Economics and Computation, Budapest, Hungary, July 18-23, 2021.* ACM, 103–104.

[6] H. Aziz, S. Gaspers, and Z. Sun. 2020. Mechanism Design for School Choice with Soft Diversity Constraints. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020.* 153–159.

[7] H. Aziz and Z. Sun. 2021. Multi-rank Smart Reserves. In *Proceedings of the 22nd ACM Conference on Economics and Computation, Budapest, Hungary, July 18-23, EC 2021.* 105–124.

[8] S. Baswana, P. P. Chakrabarti, S. Chandran, Y. Kanoria, and U. Patange. 2019. Centralized Admissions for Engineering Colleges in India. In *Proceedings of the 20th ACM Conference on Economics and Computation.* 323–324.

[9] J. R. Correa, R. Epstein, J. Escobar, I. Rios, B. Bahamondes, C. Bonet, N. Epstein, N. Aramayo, M. Castillo, A. Cristi, and B. Epstein. 2019. School Choice in Chile. In *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019, Phoenix, AZ, USA, June 24-28, 2019.* 325–343. https://doi.org/10.1145/3328526.3329580

[10] L. Ehlers, I. E. Hafalir, M. B. Yenmez, and M. A. Yildirim. 2014. School choice with controlled choice constraints: Hard bounds versus soft bounds. *Journal of Economic Theory* 153 (2014), 648–683.

[11] Y. A. Gonczarowski, N. Nisan, L. Kovalio, and A. Romm. 2019. Matching for the Israeli "Mechinot" Gap Year: Handling Rich Diversity Requirements. In *Proceedings of the 20th ACM Conference on Economics and Computation.* 321–321.

[12] I. E. Hafalir, M. B. Yenmez, and M.A. Yildirim. 2013. Effective affirmative action in school choice. *Theoretical Economics* 8, 2 (2013), 325–363.

[13] F. Kojima. 2012. School choice: Impossibilities for affirmative action. *Games and Economic Behavior* 75 (2012), 685–693.

[14] S. D. Kominers and T. Sönmez. 2013. Designing for diversity in matching. In *EC.* 603–604.

[15] R. Kurata, N. Hamada, A. Iwasaki, and M. Yokoo. 2017. Controlled school choice with soft bounds and overlapping types. *Journal of Artificial Intelligence Research* 58 (2017), 153–184.

[16] T. Sönmez and M. B. Yenmez. 2019. Affirmative Action in India via Vertical and Horizontal Reservations. (2019).

[17] T. Sönmez and M. B. Yenmez. 2020. Affirmative Action with Overlapping Reserves. (2020). http://fmwww.bc.edu/EC-P/wp990.pdf Manuscript.