

Maximizing Resource Allocation Likelihood with Minimum Compromise

Extended Abstract

Yohai Trabelsi
Department of Computer Science
Bar-Ilan University
Ramat Gan, Israel
yohai.trabelsi@gmail.com

Abhijin Adiga
Biocomplexity Institute and
Initiative, University of Virginia
Charlottesville, VA, USA
abhijin@virginia.edu

Sarit Kraus
Department of Computer Science
Bar-Ilan University
Ramat Gan, Israel
sarit@cs.biu.edu

S. S. Ravi
Biocomplexity Institute and
Initiative, University of Virginia
and University at Albany – SUNY
Charlottesville, VA, USA
ssravi0@gmail.com

ABSTRACT

Many scenarios where agents with preferences compete for resources can be cast as maximum matching problems on bipartite graphs. Our focus is on resource allocation problems where agents may have preferences that make them incompatible with some resources. We assume that a PRINCIPAL chooses a maximum matching randomly so that each agent is matched to a resource with some probability. Agents would like to improve their chances of being matched by modifying their preferences within certain limits. The PRINCIPAL’s goal is to advise an unsatisfied agent to relax its restrictions so that the total cost of relaxation is within a budget (chosen by the agent) and the increase in the probability of being assigned a resource is maximized. We develop efficient algorithms for some variants of this budget-constrained maximization problem and establish hardness results for other variants. For the latter variants, we also develop algorithms with performance guarantees. We experimentally evaluate our methods on synthetic datasets as well as on two novel real-world datasets: a vacation activities dataset and a classrooms dataset.

KEYWORDS

Matching advice; bipartite matching; resource allocation; submodular and supermodular functions

ACM Reference Format:

Yohai Trabelsi, Abhijin Adiga, Sarit Kraus, and S. S. Ravi. 2022. Maximizing Resource Allocation Likelihood with Minimum Compromise: Extended Abstract. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022)*, Auckland, New Zealand, May 9–13, 2022, Online, IFAAMAS, 3 pages.

1 INTRODUCTION

There are many practical contexts where a set of **agents** must be suitably matched with a set of **resources**. Examples of such

contexts include matching courses with classrooms [13], medical students with hospitals [15], buyers with products [10], and customers with taxicabs [6]. In this paper, we assume that the matching process assigns at most one resource to each agent and that each resource is assigned to at most one agent. It is possible that some agents are not assigned resources and some resources are unused.

Agents have **restrictions** or **preferences** while resources have **constraints**. We assume that agents’ preferences are *soft*; that is, agents are willing to *relax* their preferences so that they can get a resource that is adequate for their purposes. An agent who is unwilling to compromise may not get any resource. However, the constraints associated with resources are *hard*; that is, they cannot be relaxed.

Example: An instructor who indicates her preference for the classroom capacity as “Capacity ≥ 70 ” may be willing to relax this preference to “Capacity ≥ 60 ” to improve her chances of obtaining a classroom. However, a classroom of size 50 imposes the hard constraint “Capacity ≤ 50 ”.

An agent is **compatible** with a resource (i.e., the agent can be matched with or assigned the resource) only when the (hard) constraints of the resource are satisfied by the agent’s preferences. The problem of assigning resources to agents can be modeled as a matching problem on the following bipartite graph, which we refer to as the **compatibility graph**: the graph has two disjoint sets of nodes corresponding to the agents and resources respectively; each edge $\{u, v\}$ in the graph indicates that the agent represented by u is compatible with the resource represented by v . A PRINCIPAL (who is not one of the agents) chooses a maximum matching in the graph to maximize the number of agents who are assigned resources. Usually, there are many such maximum matchings, each one allocating resources to a (possibly) different set of agents. For fairness, the PRINCIPAL chooses a maximum matching randomly out of a given distribution. The PRINCIPAL may use, for example, an algorithm for fair matching [5] or a straight-forward process that randomly orders the agents and uses a deterministic matching algorithm such as Hopcroft-Karp’s algorithm [8] to generate a maximum matching.

Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9–13, 2022, Online, Auckland, New Zealand. © 2022 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

It is natural for an agent, who is concerned that it will not be matched in the randomly generated matching, to seek advice from the PRINCIPAL in the form of changes to its preferences in order to increase the likelihood of getting matched. In general, agents are eager to get such advice when there are several rounds of matching and they failed in previous ones. Such a situation arises, for example, in the case of medical students who were not matched during the first round of the residency matching process [9] and in hot-desking [16]. Informally, the question of developing such recommendations can be modeled as the following budget-constrained optimization problem: find a set of modifications to an unmatched agent’s preferences under a budget constraint so that the likelihood of the agent being matched to a resource is maximized, given the resource compatibility information for the other agents. (A rigorous definition of the problem is given in the next section.)

Several recommendation systems in environments where agents compete for resources are similar to our notion of a PRINCIPAL. As an example, GPS navigation apps provide advice to an agent (driver) without taking into account possible changes in the behaviors of other agents due to similar recommendations or other reasons. These recommendations often lead to undesirable consequences that are commonly referred to as the price of anarchy [17].

2 SUMMARY OF RESULTS

1. The matching advice problem. We represent the agent-resource relationship discussed above using an $\mathbb{X}\mathbb{Y}$ -bipartite graph called the *compatibility graph* $G(\mathbb{X}, \mathbb{Y}, E)$, where \mathbb{X} and \mathbb{Y} denote the set of agents and resources respectively and the edge $\{x, y\} \in E(G)$ iff the agent represented by x is compatible with the resource represented by y . The special agent who is seeking advice is denoted by x^* . When agent x^* relaxes a subset R of its restrictions, it incurs a cost $\rho(R)$, which is the sum of the costs of relaxing each restriction in R . In general, this relaxation adds edges to G , resulting in a new compatibility graph G' . We use $g(R)$ to denote the gain in probability of x^* being matched in G' over that in G . A formal definition of the main problem considered in our work is as follows.

Problem MATCHINGADVICE.

Given: A bipartite compatibility graph $G(\mathbb{X}, \mathbb{Y}, E)$, an agent $x^* \in \mathbb{X}$ seeking advice, its set of restrictions R , and a budget β .

Requirement: A subset of restrictions $R^* \subseteq R$ with $\rho(R^*) \leq \beta$ such that removal of R^* maximizes the gain in probability $g(R^*)$.

We study several forms of restrictions arising from agent preferences and resource properties in real-world applications.

2. Algorithms for and complexity of improving the likelihood of matching. We develop efficient algorithms for the above budget-constrained optimization problem for several classes of restrictions. For other classes, we show that the problem is NP-hard, develop approximation algorithms and establish their performance guarantees. These results rely on the properties of sub- and super-modular functions under constraints.

3. Experimental study. We study the performance of our recommendation algorithms on both synthetic data sets as well as two real-world data sets. The latter data sets arise in the contexts of assigning classrooms to courses and matching children with activities. We evaluate our algorithms under different cost schemes.

The insights gained from this study can inform the PRINCIPAL (e.g., university administration) on issues such as adding, removing or modifying resources to cater to the needs of agents.

Related work. Resource allocation in multi-agent systems has been studied by a number of researchers (e.g., [1–3, 7, 14]). The general focus of this work is on topics such as how agents express their resource requirements, algorithms for allocating resources to satisfy those requirements and evaluating the quality of the resulting allocations. Nguyen et al. [11] discuss some complexity and approximability results in this context. Zahedi et al. [18] study the problem of how the task allocator can respond to queries dealing with counterfactual allocations.

Motivated by e-commerce applications, Zanker et al. [19] discuss the design and evaluation of constraint-based recommendation systems that allow users to specify soft constraints regarding products of interest. These constraints are in the form of rank ordering of desired products and they consider the problem of relaxing the constraints so that a maximum number of users can obtain one of their top- k desired products, for a specified value of k . Felfernig et al. [4] provide a discussion on the design of constraint-based recommendation systems and the technologies that are useful in developing such systems. Parameswaran et al. [12] discuss the development of a recommendation system that allows university students to choose courses. The system has the capability to handle complex constraints specified by students as well as those imposed by courses. Zhou and Han [20] propose an approach for a graph-based recommendation system that groups together agents with similar preferences to allocate resources. The approach also allows users to get additional feedback regarding the allocation. To our knowledge, the problem studied in our paper, namely advising agents to modify their preferences to improve their chances of obtaining resources, has not been addressed in the literature.

3 LIMITATIONS AND FUTURE WORK

One limitation of our advice framework is the assumption that only one agent seeks advice from the PRINCIPAL while the preferences of the other agents remain unchanged. Thus, a natural direction for future work is to extend the framework to allow changes to the preferences of multiple agents. Further, our work assumes that each agent is matched to a single resource. So, another direction is to extend the advice framework by allowing agents to specify the number of resources needed. In such a case, when an agent does not receive the requested number of resources, the agent may be advised to either change her preferences or reduce the number of requested resources. We note that our framework can be easily extended to many scenarios where shared resources exist. In such cases, for a shared resource, one can simply create copies of resources with identical properties.

ACKNOWLEDGMENTS

We thank the AAMAS 2022 reviewers for their feedback. This work was supported in part by University of Virginia (UVA) Strategic Investment Fund SIF160, NSF Grant OAC-1916805 (CINES), and by the Israel Science Foundation under grant 1958/20 and the EU Project TAILOR under Grant 992215.

REFERENCES

- [1] Y. Chevaleyre, P. E. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. A. Padget, S. Phelps, J. A. Rodriguez-Aguilar, and P. Sousa. 2006. Issues in Multiagent Resource Allocation. *Informatica (Slovenia)* 30, 1 (2006), 3–31.
- [2] D. A. Dolgov and E. H. Durfee. 2006. Resource Allocation Among Agents with MDP-Induced Preferences. *J. Artif. Intell. Res.* 27 (2006), 505–549.
- [3] Paul E Dunne, Sarit Kraus, Efrat Manisterski, and Michael Wooldridge. 2010. Solving coalitional resource games. *Artificial Intelligence* 174, 1 (2010), 20–50.
- [4] A. Felfernig, G. Friedrich, D. Jannach, and M. Zanker. 2011. Developing Constraint-based Recommenders. In *Recommender Systems Handbook*. Springer, New York, NY, 187–215.
- [5] David García-Soriano and Francesco Bonchi. 2020. Fair-by-design matching. *Data Mining and Knowledge Discovery* 34, 5 (2020), 1291–1335.
- [6] Keivan Ghoseiri, Ali Haghani, Masoud Hamed, et al. 2010. *Real-time rideshare matching problem*. Technical Report. Mid-Atlantic Universities Transportation Center.
- [7] V. I. Gorodetski, O. Karsaev, and V. Konushy. 2003. Multi-agent System for Resource Allocation and Scheduling. In *Proc. CEEMAS*. Springer, New York, NY, 236–246.
- [8] J. E. Hopcroft and R. M. Karp. 1973. An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs. *SIAM J. Computing* 2, 4 (1973), 225–231.
- [9] Dafna Izenberg, Seema Marwaha, and Joshua Tepper. 2018. Medical students who don't match through CaRMS: "It's like a scarlet letter". *Healthy Debate*: <https://healthydebate.ca/2018/03/topic/medical-students-carms/>.
- [10] Linyuan Lü, Matúš Medo, and Y-C Zhang. 2009. The role of a matchmaker in buyer-vendor interactions. *The European Physical Journal B* 71, 4 (2009), 565–571.
- [11] T. T. Nguyen, M. Roos, and J. Rothe. 2013. A survey of approximability and inapproximability results for social welfare optimization in multiagent resource allocation. *Annals of Mathematics and Artificial Intelligence* 68, 1-3 (2013), 65–90.
- [12] A. G. Parameswaran, P. Venetis, and H. Garcia-Molina. 2011. Recommendation systems with complex constraints: A course recommendation perspective. *ACM Trans. Inf. Syst.* 29, 4 (2011), 20:1–20:33.
- [13] A. E. Phillips, H. Waterer, M. Ehrigott, and D. M. Ryan. 2015. Integer programming methods for large-scale practical classroom assignment problems. *Computers & Operations Research* 53 (2015), 42–53.
- [14] Hanan Rosemarin, Ariel Rosenfeld, Steven Lapp, and Sarit Kraus. 2021. LBA: Online Learning-Based Assignment of Patients to Medical Professionals. *Sensors* 21, 9 (2021), 3021.
- [15] Alvin E Roth. 1986. On the allocation of residents to rural hospitals: a general property of two-sided matching markets. *Econometrica: Journal of the Econometric Society* 54, 2 (1986), 425–427.
- [16] S. Varone and C. Beffa. 2019. Dataset on a problem of assigning activities to children, with various optimization constraints. *Data in brief* 25 (2019), 104168.
- [17] Leonard M Wapner. 2020. GPS navigation apps and the price of anarchy. *The Mathematical Gazette* 104, 560 (2020), 235–240.
- [18] Z. Zahedi, S. Sengupta, and S. Kambhampati. 2020. 'Why not give this work to them?' Explaining AI-Moderated Task-Allocation Outcomes using Negotiation Trees. arXiv: 2002.01640.
- [19] M. Zanker, M. Jessenitschnig, and W. Schmid. 2010. Preference reasoning with soft constraints in constraint-based recommender systems. *Constraints An Int. J.* 15, 4 (2010), 574–595.
- [20] W. Zhou and W. Han. 2019. Personalized Recommendation via User Preference Matching. *Information Processing and Management* 56, 3 (2019), 955–968.