Coalition Formation in Sequential Decision-Making under Uncertainty

Doctoral Consortium

Saar Cohen
Department of Computer Science
Bar Ilan University, Israel
saar30@gmail.com

ABSTRACT

As real-world applications of coalition formation continuously evolve, the design of *new* efficient algorithms that maintain a decent and consistent solution over time is required. Specifically, when agents arrive one at a time, a moderator (i.e., an *online* algorithm) must decide to which coalition the agent should be assigned, if at all. Each agent may be further accompanied with relevant information (e.g., her set of capabilities, her preferences over the previously disclosed agents), based on which the moderator performs its decisions. Multi-agent systems further encompass uncertainties in a variety of forms: the nature of the agents' participation and arrivals may be probabilistic or even unknown. Additionally, their preferences might be not assured and even incomplete or strategic. This research will thus lay the theoretical foundations for studying the interplay between coalition formation and online, uncertain settings, while characterizing the factors which make the moderator's objective susceptible. Our methods will be further tied to practical applications, specifically ones in physical settings (e.g., task allocation in actual robots).

KEYWORDS

Coalition Formation; Online Algorithms; Probabilistic Inference; Constrained Markov Decision Processes; Reinforcement Learning

ACM Reference Format:

Saar Cohen. 2023. Coalition Formation in Sequential Decision-Making under Uncertainty: Doctoral Consortium. 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

A central issue of multi-agent systems is *cooperation* [2, 3, 10], whose objective is creating agents capable of taking joint and coordinated actions so as to improve their performance or accomplish goals that are beyond their individual capabilities. In several real-life scenarios, we thus notice the phenomenon of *coalition formation*, where each person, termed as *agent*, forms coalitions with others to get some benefit, experiencing a utility that depends on the particular set of agents she joins. For instance, in distributed vehicle routing, coalitions of delivery companies can reduce the transportation costs by sharing deliveries [15], and in information gathering several information servers come together to answer queries [11].

Prior studies typically consider coalition formation under the standard settings where a centralized deterministic entity (i.e., an *offline* algorithm) has *full* knowledge regarding the system's components and aims at deciding on a partition of the agents into coalitions. Yet, such an assumption is unrealistic in various contexts where the set of agents is unknown beforehand and its construction may demand contacting a large population. For instance, crowdsourcing markets allow task requesters to inexpensively access via a central platform a large manpower of workers with *multiple* diverse skills that sequentially arrive one by one [8]. In knowledge-intensive tasks, workers collaboratively perform complex data collection and analysis tasks, while each worker's contributions gradually increases the quality of each knowledge piece.

Another two common underlying assumptions in the coalition formation literature are that all agents are assured to participate in the game, and the nature of their collaborations and their preferences over those are certain. However, both often fail to hold for real-world problems. Generally, the agents' participation in the game and their arrivals do not depend only on choices, but also on external factors. For instance, in crowdsourcing [17], a worker may not be able to solve all the tasks due to an inability to translate certain languages or to handle time-sensitive tasks with a deadline.

Research Goal. I aim at providing a unifying framework for coalition formation in online and uncertain environments. My main goal is characterizing the factors which make the moderator's objective susceptible. I believe that among those influences are the moderator's knowledge on the arriving agents and the nature of the policies to which he commits, as well as the context at hand (e.g., the presence of strategic manipulation, cases where changing assignments is possible but expensive, or cases where the agents' capabilities are diverse). In particular, my research concentrates on theoretical analysis of online coalition formation in its most general form under online and uncertain settings. An additional objective posed by my research plan is further extending my current and future contributions to methods adapted to the physical world encompassed by other realms such as robotics.

2 ONLINE COALITIONAL SKILL FORMATION

As a first step towards mitigating the aforementioned shortcomings, we initially focused on various multi-agent applications such as crowdsourcing [17] and rescue operations [20], where agents with diverse skills emerge *dynamically* and are then assigned to tasks with heterogeneous requirements. Such scenarios can be framed as a special case of the notorious multi-agent task allocation problem [16] that attracts extensive attention – *online task allocation* [18]. When considering the aforementioned example of crowdsourcing

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.

markets [8], workers should solve problems like participatory sensing [21] and human computation [7] which require specific skills (e.g., specialized translation requires not only ability in at least two languages but also domain knowledge).

The most desirable goal of online task allocation is to *assign the most suitable agents to tasks* [14]. Namely, an agent with the skill level required by a certain task or at least a sufficient experience should be ideally assigned to that task. It is often the case that a single agent does not have the skill level needed to achieve a particular task, thus it is necessary for the agents to form *real-time* coalitions for completing certain tasks. In coalitions, agents can collaboratively complete tasks more efficiently or accurately by cooperating to meet those requirements [19]. However, adequately *modeling multiple skills* of agents, their correlations, the uncertainty about their conditional dependencies and their contributions to the formed coalitions is difficult, since different agents usually possess different skills and diverse degree of proficiency in the same skill [12]. Further, improperly measuring the coalitions' suitability (i.e., meeting the tasks' requirements as much as possible) may hinder the quality of a task's execution [9].

As such, we developed a novel framework termed as *online coalitional skill formation* (**OCSF**), for handling online task allocation from a standpoint of coalition formation [5]. The goal of the moderator is therefore to assign agents that arrive online to a coalition responsible for performing some task, so as to optimally approach the desired skill levels of all tasks. We discuss the efficiencies in considering skill domains in which the set of possible mastering levels for each skill is *discrete* instead of continuous. We then suggest different assignment algorithms based on the knowledge the moderator has on the arriving agents. When agents arrive i.i.d. according to some unknown distribution, though a *greedy* and adaptive scheme that assigns an agent to a task performs sufficiently well, the expected number of agents contacted may be arbitrarily large due to its adaptivity. If the distribution is *known*, we thus devised a novel correlation to Constrained Markov Decision Processes whose goal is maximizing the rate at which agents are assigned to each task while respecting their requirements. Applying the approach to our context, we observed that a *non-adaptive* scheme that terminates when all the tasks' requirements are met achieves the *best* performance. Finally, if the distribution is *unknown*, we provide two algorithms that *learn it online*, and whose estimation by the learning schemes was proven yield an additional factor that depends on the size of the the skill domain. Empirically, we validated that a higher diversity in skills may indeed yield suboptimal assignments.

3 COMPLEXITY OF PROBABILISTICALLY INFERRING SOLUTION CONCEPTS

In a second research trend of my thesis, we explored (offline) *Hedonic games* [6], which are a popular game-theoretic approach to the study of coalition formation whose outcome is a *partition* of the agents into coalitions, over which the agents have preferences. One of their main properties is *non-externality*: an agent minds only her own coalition, regardless of how the others aggregate. Two major underlying assumptions in hedonic games literature are that all agents are assured to participate in the game, and the nature of their collaborations and preferences over those are certain. However, both often fail to hold for real-world problems. Generally, participation in the game does not depend only on strategic choices, but also on *external factors*. For instance, in crowdsourcing markets, a worker may not be able to meet a task's deadline or have the required level of proficiency for solving it.

Focusing on various classes of *dichotomous hedonic games* (DHGs) [1], where each agent either approves or disapproves a given coalition, we thus proposed the *random extension* [4], where agents have an independent participation probability. As such, we initiated the research on the computational complexity of computing the probability that coalitions and partitions are optimal or stable [4]. While some cases admit efficient algorithms (e.g., agents approve only few coalitions), they become computationally hard (#P-hard) in their complementary scenario. We supplied a reduction to our problem from counting variants of a *new* type of manipulation in hedonic games: *constructive* control by *adding* players, whose goal is *ensuring* that a specified outcome satisfies a certain solution concept. In turn, our novel correlation allows us to drive our hardness results.

4 CONCLUSIONS AND RESEARCH PLAN

Since coalition formation in practical environments continuously evolves in time, the design of approximation algorithms that maintain a good solution over time (both theoretically and empirically) constitutes as my central challenge. I intend to devise algorithms that are not only efficient and robust to abrupt changes and uncertainty, but also make consistent decisions. When further placing ourselves in physical, online and uncertain domains as the ones imposed by the field of robotics, additional challenges arise which thus require special treatment for deriving practical algorithms for such realms. In addition to the limited sensing and computational capabilities typically owned by robots, their physical dimension further imposes special challenges (e.g., concealment disturbances), due to which appropriate schemes are required. As such, we will also perform experimental evaluations using realistic simulators (such as ARGoS [13]) and/or physical robots.

Improving algorithms via predictions is a very active research topic in recent years. As we observed in our algorithms [5], bridging between this strand and ours has the potential to bypass the worst-case lower bounds of our online problem and even devise optimal algorithms. Hence, I intend to design online algorithms enhanced with (potentially erroneous) *predictions* for our problem which have efficient tradeoffs between *consistency* (when the predictions are *accurate*, the performances of our algorithms should be close to those of an optimal offline algorithm), and *robustness* (when the predictions are *erroneous*, the performance of our algorithms should be close to those of an online algorithm without predictions). Another appealing direction for future research is considering scenarios where decisions may be postponed, agents may be reassigned to another coalition after each arrival, or both. Exploring OCSF as well as other models for online coalition formation in the presence of *strategic manipulation* and designing algorithms *robust* to manipulations also constitutes a major challenge. I aim at characterizing fairness-oriented schemes and exploring their impact on the system. Finally, there is an abundance of applications other than online task allocation [5], where it is of high interest to investigate coordination and cooperation between agents.

ACKNOWLEDGMENTS

This research was funded in part by ISF grant #1563/22. This work has been conducted in collaboration with my advisor, Prof. Noa Agmon, to whom I would like to thank for her constant guidance and support throught my research.

REFERENCES

[1] Haris Aziz, Paul Harrenstein, Jérôme Lang, and Michael Wooldridge. 2016. Boolean hedonic games. In *Proceedings of the 15th International Conference on Principles of Knowledge Representation and Reasoning*. 166–175.

[2] Saar Cohen and Noa Agmon. 2021. Recent Advances in Formations of Multiple Robots. *Current Robotics Reports* 2, 2 (2021), 159–175.

[3] Saar Cohen and Noa Agmon. 2022. Optimizing Multi-Agent Coordination via Hierarchical Graph Probabilistic Recursive Reasoning. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*. 290–299.

[4] Saar Cohen and Noa Agmon. 2023. Complexity of Probabilistic Inference in Random Dichotomous Hedonic Games. In *Proceedings of the AAAI Conference on Artificial Intelligence*.

[5] Saar Cohen and Noa Agmon. 2023. Online Coalitional Skill Formation. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (to appear)*.

[6] Jacques H Dreze and Joseph Greenberg. 1980. Hedonic coalitions: Optimality and stability. *Econometrica* (1980), 987–1003.

[7] Umair Ul Hassan, Sean O'Riain, and Edward Curry. 2013. Effects of expertise assessment on the quality of task routing in human computation. In *2nd International Workshop on Social Media for Crowdsourcing and Human Computation (SoHuman 2013) 2*. 1–10.

[8] Danula Hettiachchi, Vassilis Kostakos, and Jorge Goncalves. 2022. A survey on task assignment in crowdsourcing. *ACM Computing Surveys (CSUR)* 55, 3 (2022), 1–35.

[9] Chien-Ju Ho and Jennifer Vaughan. 2012. Online task assignment in crowdsourcing markets. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 26. 45–51.

[10] Zool Hilmi Ismail, Nohaidda Sariff, and EG Hurtado. 2018. A survey and analysis of cooperative multi-agent robot systems: challenges and directions. In *Applications of Mobile Robots*. IntechOpen, 8–14.

[11] Matthias Klusch and Onn Shehory. 1996. A polynomial kernel-oriented coalition algorithm for rational information agents. *Tokoro, ed* (1996), 157–164.

[12] Panagiotis Mavridis, David Gross-Amblard, and Zoltán Miklós. 2016. Using hierarchical skills for optimized task assignment in knowledge-intensive crowdsourcing. In *Proceedings of the 25th International Conference on World Wide Web*. 843–853.

[13] Carlo Pinciroli, Vito Trianni, Rehan O'Grady, Giovanni Pini, Arne Brutschy, Manuele Brambilla, Nithin Mathews, Eliseo Ferrante, Gianni Di Caro, Frederick Ducatelle, et al. 2012. ARGoS: a modular, parallel, multi-engine simulator for multi-robot systems. *Swarm intelligence* 6, 4 (2012), 271–295.

[14] Chenxi Qiu, Anna C Squicciarini, Barbara Carminati, James Caverlee, and Dev Rishi Khare. 2016. CrowdSelect: increasing accuracy of crowdsourcing tasks through behavior prediction and user selection. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*. 539–548.

[15] Tuomas W Sandhlom and Victor RT Lesser. 1997. Coalitions among computationally bounded agents. *Artificial intelligence* 94, 1-2 (1997), 99–137.

[16] Onn Shehory and Sarit Kraus. 1998. Methods for task allocation via agent coalition formation. *Artificial intelligence* 101, 1-2 (1998), 165–200.

[17] Aleksandrs Slivkins and Jennifer Wortman Vaughan. 2014. Online decision making in crowdsourcing markets: Theoretical challenges. *ACM SIGecom Exchanges* 12, 2 (2014), 4–23.

[18] Hanna Sumita, Shinji Ito, Kei Takemura, Daisuke Hatano, Takuro Fukunaga, Naonori Kakimura, and Ken-ichi Kawarabayashi. 2022. Online Task Assignment Problems with Reusable Resources. 36, 1 (2022), 5199–5207.

[19] Lovekesh Vig and Julie A Adams. 2006. Market-based multi-robot coalition formation. In *Distributed Autonomous Robotic Systems 7*. Springer, 227–236.

[20] Andrew K Whitten, Han-Lim Choi, Luke B Johnson, and Jonathan P How. 2011. Decentralized task allocation with coupled constraints in complex missions. In *Proceedings of the 2011 American Control Conference*. IEEE, 1642–1649.

[21] Qingwen Zhao, Yanmin Zhu, Hongzi Zhu, Jian Cao, Guangtao Xue, and Bo Li. 2014. Fair energy-efficient sensing task allocation in participatory sensing with smartphones. In *IEEE INFOCOM 2014-IEEE Conference on Computer Communications*. IEEE, 1366–1374.