Multi-agent Reinforcement Learning for Decentralized Coalition Formation Games

Kshitija Taywade
Department of Computer Science
University of Kentucky
kshitija.taywade@uky.edu

Abstract

We study the application of multi-agent reinforcement learning for game-theoretical problems. In particular, we are interested in coalition formation problems and their variants such as hedonic coalition formation games (also called hedonic games), matching (a common type of hedonic game), and coalition formation for task allocation. We consider decentralized multi-agent systems where autonomous agents inhabit an environment without any prior knowledge of other agents or the system. We also consider spatial formulations of these problems. Most of the literature for coalition formation problems does not consider these formulations of the problems because it increases computational complexity significantly. We propose novel decentralized heuristic learning and multi-agent reinforcement learning (MARL) approaches to train agents, and we use game-theoretic evaluation criteria such as optimality, stability, and indices like Shapley value.

A coalition formation game is a cooperative game in which agents form groups, or coalitions, based on their preferences. Most of the previous approaches for different types of coalition formation games have been centralized, assuming the presence of the central agency to control agents. However, this assumption is infeasible for many real-world scenarios. In our work, we model the variants of coalition formation problem as decentralized multi-agent systems and train autonomous agents to learn behavior strategies using novel decentralized learning and MARL approaches. Moreover, all the agents are autonomous, very much like most of the real-world scenarios. We focus on the three main variants of coalition formation game: Hedonic coalition formation games, Matching problems and, Coalition formation for task allocation. Hedonic games have applications in team formation and social group formation where agents can have preferences over coalitions. Matching, which is a variant of hedonic games, finds application in problems like matching between workers-employers, students-colleges, residents-hospitals, etc. Several real-world scenarios like robot teams, rescue teams, merger of organizations, inter-organization team formation can be characterized as a coalition formation for task allocation problem.

As we model the problems as decentralized multi-agent systems with fully autonomous agents, challenges arise due to the dynamicity of the environment, defined by the uncertainty and non-stationary behavior of the agents. Because of these complex challenges, it is hard to solve them with the pre-programmed agents. Instead, agents must look for the solutions on their own, using some sort of learning approach. Additively separable hedonic games (ASHGs) are one of the simplest formulations of the coalition formation problem, and finding optimal partitions for ASHGs is NP-hard in the strong sense. Therefore, we propose two learning approaches, decentralized heuristic learning and MARL. In heuristic learning, autonomous agents explore the unknown environment and collect useful information about the environment i.e. other agents present in the system, obtainable utility value, locations of other agents/coalitions, etc. Agents consistently update their information while making the rule-based decisions based on their current knowledge. We wanted to know the effectiveness of heuristic learning before proceeding to use reinforcement learning, as we considered the former simple and computationally easier than RL. A significant part of the research on multi-agent learning concerns RL techniques. As those techniques have been proved successful, we want to contribute by exploring new MARL approaches for solving a subset of problems in multi-agent settings as discussed above.

We model the problems of our interest in two main types of environments: spatial (grid world) and non-spatial (social networks, especially for bipartite matching). The spatial extension of the problems can better represent the real-world scenarios where agents must expend some efforts to seek out information about other agents.

Hedonic Coalition Formation Games

We have worked on hedonic coalition formation games, specifically ASHGs (Taywade, Goldsmith, and Harrison 2018). A hedonic game is a coalition formation game in which agents have preferences over the coalitions they can be a part of. In an additively separable hedonic game, an agent's utility for a coalition is equal to the sum of its utility for every other agent in the coalition. We used the grid world environment with agents starting off from the random grid locations. We implemented the decentralized learning approach in which agents gather knowledge about the environment by exploration, while also making the rule-based decisions based on that knowledge. Decentralized heuristic learning proved to be very effective, we got up to 95% close to the optimal results for matching problems and it also gave huge improvements on random partitions. We have also proposed an extension to our method, specifically for ASHGs, a "budding" technique which is an additional search heuristic in which new coalitions are formed by breaking apart previously formed coalitions in a locally centralized way. Our proposed budding heuristics are: Random, Greedy and Clique Detection. These techniques led to, overall, higher utility coalitions for all the agents involved.

Matching Problems

We have previously worked on optimal bipartite matching (Taywade, Goldsmith, and Harrison 2018, 2020), and roommate matching (Taywade, Goldsmith, and Harrison 2018), while more recent work includes stable bipartite matching. Bipartite matching and roommate matching are common types of hedonic games. In (Taywade, Goldsmith, and Harrison 2020), we explore the ways of modeling the environments in terms of bipartite matching (marriage problem) as there are several personal, social, and cultural factors that influence how people find potential partners. Our goal for this is to reach optimal matching for the whole system. We model the marriage problem using: a grid world environment, a small-world network graph, and an affiliation network graph. These environments represent the ways people meet each other in the real-world. In the grid world model, agents actively seek out potential partners on their own with little to no prior information about the environment. In the small-world network environment, agents are presented potential matches based on their degrees of separation. And in affiliation networks, agents are registered to matrimonial agencies/dating websites and get potential matches suggested by those agencies. In this work, we have proposed a different version of heuristic learning than in (Taywade, Goldsmith, and Harrison 2018). The overall utility of matches found through our decentralized heuristic learning approach were, on average, 84% of the utility achieved by optimal (Hungarian) centralized algorithm.

In our recent work, we proposed the use of a multi-agent reinforcement learning paradigm for a spatially formulated decentralized two-sided stable matching problem with independent and autonomous agents. This environment is very dynamic and uncertain. Moreover, agents lack the knowledge of preferences of other agents and have to explore the environment and interact with other agents to discover their own preferences through noisy rewards. We think such a setting better approximates the real-world and we studied the usefulness of our MARL approach for it. Along with conventional stable matching case where agents have strictly ordered preferences, we checked the applicability of our approach for stable matching with incomplete lists and ties. We ran experiments with various grid sizes and total number of agents in the system. We investigated our results for stability, level of instability (for unstable results), and fairness. We measured regret cost, set-equality cost and egalitarian cost. Our MARL approach mostly yields stable and fair outcomes.

Coalition Formation for Task Allocation

The problem of coalition formation for task allocation considers tasks requiring more than one agent to collaborate in order to perform them. Therefore, agents should take into account many more aspects such as tasks at hand, their own capabilities, capabilities of other agents, and agreement on the distribution of the payoffs. This is part of our future work. We intend to model this problem as a grid world, with tasks and agents placed on random locations in the initial grid configuration. In this, tasks require agents to form coalitions such that the agents in those coalitions are equipped with certain skills required by the tasks, with collective efficiencies in those skills more than the threshold values. Therefore, we represent each agent with an efficiency vector, a vector of efficiencies possessed by an agent for every skill which is part of the model. Agents are unaware of the efficiency vectors of other agents. Tasks can be performed in parallel without having overlapping coalitions and our goal is to have a partition such that for every task there is a coalition of agents that has collective efficiencies of agents above the threshold values, for the skills required by the tasks.

Our aim is to form acceptable partitions of agents, where acceptable means all tasks can be performed in parallel and also such that there is a direct correlation between Shapley values and obtained utilities. We want to propose a novel MARL approach to solve this problem. Researchers have applied RL as well as deep RL to solve coalition formation problems for performing tasks. In (Bachrach et al. 2020), authors propose a framework for training agents to negotiate and form teams using deep RL, they have also used the grid (spatially extended) environment in their experiments. Bayesian RL has also been used for coalition formation problems (Chalkiadakis and Boutilier 2004). The problem of uncertainty in the environment is more complex when we consider that the agents don't know the capabilities of other agents or even coalition values for them. We propose to address this by incorporating Bayesian learning into MARL, especially to reduce the uncertainties about the capabilities of other agents. In Bayesian learning, a prior over the unknown is updated as more information becomes available. We would also like to study the effectiveness of the incorporation of Bayesian learning by analyzing the differences in the results with and without its incorporation into MARL approach. We anticipate to make progress on MARL approach for coalition formation by the time of the workshop.

References

Bachrach, Y.; Everett, R.; Hughes, E.; Lazaridou, A.; Leibo, J. Z.; Lanctot, M.; Johanson, M.; Czarnecki, W. M.; and Graepel, T. 2020. Negotiating team formation using deep reinforcement learning. *Artificial Intelligence* 288: 103356.

Chalkiadakis, G.; and Boutilier, C. 2004. Bayesian reinforcement learning for coalition formation under uncertainty. In *IJCAI '04*, 1090–1097.

Taywade; Goldsmith; and Harrison. 2018. Decentralized Multi-agent Approach for Hedonic Games. In *EUMAS '18*.

Taywade; Goldsmith; and Harrison. 2020. Decentralized Marriage Models. In *FLAIRS '20*.