Towards a competence-based approach to allocate teams to tasks

Extended Abstract

Athina Georgara*
IIIA-CSIC
Barcelona, Spain
ageorg@iiia.csic.es

Juan A. Rodríguez-Aguilar
IIIA-CSIC
Barcelona, Spain
jar@iiia.csic.es

Carles Sierra
IIIA-CSIC
Barcelona, Spain
sierra@iiia.csic.es

ABSTRACT

We tackle the problem of allocating teams to tasks. For that, we propose a novel allocation problem that matches teams with tasks based on the similarity between the competencies required by the tasks and the competencies offered by teams. We formally cast our problem as an optimisation problem, characterise the size of its search space, and outline a heuristic approach to solve it.

KEYWORDS

Team Formation; Task allocation; Optimisation; Heuristics

1 INTRODUCTION

Team formation is a line of research that has received much attention within the artificial intelligence and multiagent systems' literature. For example, [6] studies the formation of robot teams for search and rescue missions, and [11] forms teams for unmanned aerial vehicles for surveillance missions. There are several challenges linked to team formation, such as forming a single or multiple teams that satisfy some desirable properties, or are stable under a game theoretic point of view. In this work we focus on the problem of allocating teams to tasks requiring multiple competencies. The existing literature has studied different types of matching problems between teams and tasks. [1, 2, 9] study the formation of a single team to be matched with a single task. Forming a single team to resolve multiple tasks is studied in [7]; while [4] studies the formation of multiple teams to resolve the same task. Finally, [6] and [5] address how to form multiple teams to resolve multiple tasks while admitting *overlaps* of different sorts. This means that an agent i might be part of multiple teams, a team might tackle multiple tasks, and multiple tasks might be assigned the same task. Allocating teams to multiple tasks disregarding overlaps has been studied in [8] and [12]. However, both [8] [12] exhibit practical limitations: limited scalability, and not expressive enough competence models. In this work, we focus on this last problem and work towards a

competence-based approach to solve it without significant practical limitations.

2 ALLOCATING TEAMS TO TASKS PROBLEM

Within a known multiagent system, we can determine a fixed set of *competencies*, denoted by C; where each competence represents a specified capability, skill, characteristic or piece of knowledge. Each agent is characterised by their *competencies*. Formally, given a set A of agents, with $|A| = n$, the competencies of a are $C_a \subseteq C$. Similarly, we describe each task by means of its *required competencies*, i.e., the necessary competencies a team must count on to successfully fulfil the task. We complete the specification of a task with the relative importance of each competence in the task and the required team size. Formally, task τ is given by $\langle id_\tau, C_\tau, w_\tau, s_\tau \rangle$, where id_τ is a unique identifier, $C_\tau \subseteq C$ are the required competencies, $w_\tau : C_\tau \to (0, 1]$ is a weighted function, and $s \in \mathbb{N}_+$ is the required team size. The set of all tasks is denoted by T, with $|T| = m$. Given task $\tau \in T$, we denote the set of all size-compliant teams for τ as $\mathcal{K}_\tau = \{K \subseteq A : |K| = s_\tau\}$.

2.1 Suitability of a team for a task

Whether a size-compliant team is suitable for some task, and the degree of this suitability, is determined via competencies. That is, on the one hand a task τ specifies a set of required competencies C_τ that are necessary for the task to be successfully solved. On the other hand, a team of agents, $K \subseteq A, |K| = s_\tau$, collectively possess a set of competencies $C_K = \bigcup_{a \in K} C_a$. Thus, the suitability of K to τ primarily depends on the matching between τ's required competencies and those collectively offered by K.

One step further, it is natural to assume that each agent shall be responsible for some of the required competencies; and collectively, considering the responsibilities assigned to all team members, the team should cover *all* competencies required by the task. In order to assign responsibilities, we require a number of properties to be satisfied by any allocation. First and foremost, each required competence must be assigned to at least one team member. This ensures that all requirements will be covered by the team. Also, we need that each member contributes to the task by adopting at least one competence as their responsibility. Moreover, we limit the number of competencies that each agent can undertake to avoid overloading. To capture such requirements, we built on the inclusive competence allocation introduced in [3]:

*Definition 2.1 (*Fair Competence allocation Function (FCAF)).* Given a task τ and a size-compliant team of agents K, we say that $\eta_{\tau \to K} : K \to 2^{C_\tau}$ is a fair competence allocation function if it satisfies: (i) $\bigcup_{a \in K} \eta_{\tau \to K}(a) = C_\tau$; and (ii) $1 \le |\eta_{\tau \to K}(s)| \le \left\lceil \frac{|C_\tau|}{|K|} \right\rceil$.*

Therefore, given an FCAF $\eta_{\tau \to K}$, we can determine the suitability of team K to τ by aggregating the suitability of each team member to their assigned competencies. For any agent $a \in K$ and each of their assigned competencies $c \in \eta_{\tau \to K}(a)$, there is a competence c' possessed by a (i.e., $c' \in C_a$) that is the most similar to c. That is, a's suitability to τ is determined via the maximum degree of similarity between all assigned competencies, $\eta_{\tau \to K}(a)$, and those possessed by a, C_a. Thus, if $f(a, \eta_{\tau \to K})$ determines the suitability of agent a to its assigned competencies, then $\prod_{a \in K} f(a, \eta_{\tau \to K})$ is K team's suitability to τ.

2.2 Defining the optimisation problem

Given a set of different tasks T (with $|T| = m$), and a population of agents A (with $|A| = n$), the problem to solve is the following: *(i)* form non-overlapping teams from the agents in A, and *(ii)* match each task with at most one size-compliant team so that the *overall teams' suitability* is maximised. Let $g : T \to 2^A$ be an allocation function and G the family of all allocation functions. Then, we determine the overall teams' suitability in g as the product of the suitability of each team to its assigned task. Thus, the optimal allocation function g^* is such that $g^*(\tau) \cap g^*(\tau') = \emptyset$ for every pair or tasks $\tau \neq \tau'$, $|g^*(\tau)| = s_\tau$ for every $\tau \in T$, and:

$$g^* = \arg \max_g \prod_{\tau \in T} \Big(\prod_{a \in g(\tau)} f(a, \eta^*_{\tau \to g(\tau)}) \Big) \qquad (1)$$

Note that in equation 1 we use η^*, i.e., the optimum FCAF of each pair $\langle \tau, g^*(\tau) \rangle$ for $\tau \in T$. Thus, the optimum FCAF should be the one that maximises the team's suitability, i.e. maximises $\prod_{a \in g^*(\tau)} f(a, \eta_{\tau \to g^*(\tau)})$. To determine the overall team's suitability in an allocation function g, we need to solve one separate optimisation problem for each task. Each of these $|T|$ optimisation problems is a local competence allocation problem for each task. Next, we study the space of allocation functions G.

2.3 Characterising the search space

The purpose of this section is to characterise the search space defined by the optimisation problem described above. This amounts to quantifying the number of feasible team allocation functions in G. For that, we start by splitting the tasks in T into r buckets of tasks, where the tasks in the same bucket require teams of the same size. That is, we have $b_1, \ldots, b_r \subseteq T$ buckets where $b_i \cap b_j = \emptyset$, $\forall i, j = 1, \cdots, r$ and $\bigcup_{i=1}^r b_i = T$. Each bucket b_i is characterised by the required team size s_i by all tasks in b_i; i.e., task $\tau_a \in b_i$ if and only if $s_{\tau_a} = s_i$. Moreover, for any pair of buckets b_i and b_j, it holds that $s_i \neq s_j$. Next, we will distinguish three cases when counting the number of feasible team allocations in G:

Case I. We have exactly as many agents as required by all tasks in T, $\sum_{\tau \in T} s_\tau = n$. In this case, we seek for partition functions over T. The size of G is $\frac{n!}{\prod_{i=1}^r (s_i!)^{|b_i|}}$ according to Theorem 3.4.19 in [10].

Case II. We have more agents than those required by all tasks in T, namely $\sum_{\tau \in T} s_\tau < n$. Following example 3.4.20 in [10], we assume one more bucket b_{r+1} containing exactly one *auxiliary* task, which requires a team of size $s_{r+1} = n - \sum_{i=1}^r s_i \cdot |b_i|$. Now there are $|G| = \frac{n!}{\prod_{i=1}^k (s_i!)^{|b_i|} \cdot (n - \sum_{i=1}^r |b_i| \cdot s_i)!}$ feasible team allocation functions.

Case III. We have less agents than the required ones by all tasks in T, $\sum_{\tau \in T} s_\tau > n$. In this case, first we need to introduce $cover(T, A) = \{T' \subset T : \sum_{\tau \in T'} s_\tau \leq n \wedge \nexists \tau' \in T - T' : s_{\tau'} \leq n - \sum_{\tau \in T'} s_\tau\}$ as the set that contains all the subsets of tasks $T' \subset T$ that along A leads to Case I or Case II, and by adding any $\tau \notin T'$ in T' leads to Case III. In total there are: $|G| = \sum_{T' \in cover(T, A)} \frac{n!}{\prod_{i=1}^r (s_i!)^{|b_i|} \cdot (n - \sum_{i=1}^r |b_i| \cdot s_i)!}$ different feasible team allocation functions, where variables $r, b_1, \cdots, b_r$ and $s_1, \cdots, s_r$ change according to T'. The size of set $cover(T, A)$ depends on the total number of agents, and the team sizes required by the tasks in T.

Note that the number of feasible team allocation functions quickly grows with the number of tasks and agents, hence leading to vast search spaces. This is rather challenging solving our optimisation problem in real life scenarios. For this reason, next section discusses a heuristic methodology to solve our problem.

3 OUTLINING A HEURISTIC METHODOLOGY

Here we describe a heuristic methodology to solve the optimisation problem formalised by equation 1. Our methodology consist of two steps: *(a)* find an initial feasible allocation, and *(b)* successively improve the initial allocation.

Initial feasible allocation In order to find an initial feasible allocation, we use some 'analytics' on the tasks and agents at hand. That is, we determine the hardness of finding a suitable team for each task based on the competencies required by the tasks, and the competencies offered by all agents. Then, starting from the hardest task, we look for a size-compliant team that is suitable for the task. That is, we search for a team whose agents possess competencies that are the most similar to some required competence by the task.

Improve allocation. At this step we iteratively try to improve the initial team allocation. At each iteration, we perform two main operations for this purpose. On the one hand, given two randomly selected tasks $\tau, \tau' \in T$, we try to greedily optimise allocations $g(\tau)$ and $g(\tau)$ using only the agents in $g(\tau) \cup g(\tau')$ and currently unassigned agents (if they exist). On the other hand, we consider all pairs of tasks and search for improving allocations of pairs of tasks by swapping agents between allocations.

4 CONCLUSIONS & ONGOING WORK

We tackled the problem of forming multiple teams to allocate them to multiple, distinct tasks, while disregarding team overlaps. We introduced the key elements of the problem, and then we described our optimisation problem. We also characterised the search space of our problem, which grows very large even with a small number of agents and tasks. The large size of our search space motivated the introduction of a heuristic methodology for solving it. Currently we are working towards a systematic evaluation to confirm the effectiveness of the proposed methodology. Thus, we are investigating: *(i)* how to evaluate the quality of the solutions obtained by our heuristic approach, along with the time that it requires, over synthetic data; and *(ii)* the capabilities of our approach to handle and solve real-world problems. Our initial results are promising regarding the quality of the solutions obtained by our approach.

REFERENCES

[1] Aris Anagnostopoulos, Luca Becchetti, Carlos Castillo, Aristides Gionis, and Stefano Leonardi. 2010. Power in Unity: Forming Teams in Large-Scale Community Systems. *International Conference on Information and Knowledge Management, Proceedings*, 599–608. https://doi.org/10.1145/1871437.1871515

[2] Aris Anagnostopoulos, Luca Becchetti, Carlos Castillo, Aristides Gionis, and Stefano Leonardi. 2012. Online Team Formation in Social Networks. In *Proceedings of the 21st International Conference on World Wide Web* (Lyon, France) *(WWW '12)*. Association for Computing Machinery, New York, NY, USA, 839–848. https://doi.org/10.1145/2187836.2187950

[3] Ewa Andrejczuk, Rita Berger, Juan A. Rodríguez-Aguilar, Carles Sierra, and Víctor Marín-Puchades. 2018. The composition and formation of effective teams: computer science meets organizational psychology. *Knowledge Eng. Review* 33 (2018), e17. https://doi.org/10.1017/S026988891800019X

[4] Ewa Andrejczuk, Filippo Bistaffa, Christian Blum, Juan A. Rodríguez-Aguilar, and Carles Sierra. 2019. Synergistic team composition: A computational approach to foster diversity in teams. *Knowledge-Based Systems* 182, 104799 (10/2019 2019). https://doi.org/10.1016/j.knosys.2019.06.007

[5] Yoram Bachrach, Reshef Meir, Kyomin Jung, and Pushmeet Kohli. 2010. Coalitional Structure Generation in Skill Games. *24th AAAI Conference on Arti-cial Intelligence (AAAI)* 2.

[6] Luca Capezzuto, Danesh Tarapore, and Sarvapali D. Ramchurn. 2020. Anytime and Efficient Coalition Formation with Spatial and Temporal Constraints. arXiv:2003.13806 [cs.MA]

[7] Chad Crawford, Zenefa Rahaman, and Sandip Sen. 2016. Evaluating the Efficiency of Robust Team Formation Algorithms. In *Autonomous Agents and Multiagent Systems*, Nardine Osman and Carles Sierra (Eds.). Springer International Publishing, Cham, 14–29.

[8] E. Czatnecki and A. Dutta. 2019. Hedonic Coalition Formation for Task Allocation with Heterogeneous Robots. In *2019 IEEE International Conference on Systems, Man and Cybernetics (SMC)*. 1024–1029.

[9] Theodoros Lappas, Kun Liu, and Evimaria Terzi. 2009. Finding a Team of Experts in Social Networks. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Paris, France) *(KDD '09)*. Association for Computing Machinery, New York, NY, USA, 467–476. https://doi.org/10.1145/1557019.1557074

[10] R. B. Maddox. 2002. *Mathematical Thinking and Writing: A transition to abstract mathematics* (1st ed.). Academic Pr, San Diego.

[11] Sameera S. Ponda, Luke B. Johnson, Alborz Geramifard, and Jonathan P. How. 2015. *Cooperative Mission Planning for Multi-UAV Teams*. Springer Netherlands, Dordrecht, 1447–1490. https://doi.org/10.1007/978-90-481-9707-1_16

[12] Fredrik Präntare and Fredrik Heintz. 2018. An Anytime Algorithm for Simultaneous Coalition Structure Generation and Assignment. In *PRIMA 2018: Principles and Practice of Multi-Agent Systems*, Tim Miller, Nir Oren, Yuko Sakurai, Itsuki Noda, Bastin Tony Roy Savarimuthu, and Tran Cao Son (Eds.). Springer International Publishing, Cham, 158–174.