Artificial Intelligence 208 (2014) 41–65

Contents lists available at ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artint

Weighted synergy graphs for effective team formation with heterogeneous ad hoc agents

Somchaya Liemhetcharat*, Manuela Veloso

School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213, USA

ARTICLE INFO

Article history:
Received 12 September 2012
Received in revised form 12 December 2013
Accepted 17 December 2013
Available online 19 December 2013

Keywords:
Multi-agent
Multi-robot
Heterogeneous
Capability
Synergy
Ad hoc
Team formation

ABSTRACT

Previous approaches to select agents to form a team rely on single-agent capabilities, and team performance is treated as a sum of such known capabilities. Motivated by complex team formation situations, we address the problem where both single-agent capabilities may not be known upfront, e.g., as in ad hoc teams, and where team performance goes beyond single-agent capabilities and depends on the specific *synergy* among agents. We formally introduce a novel weighted synergy graph model to capture new interactions among agents. Agents are represented as vertices in the graph, and their capabilities are represented as Normally-distributed variables. The edges of the weighted graph represent how well the agents work together, i.e., their synergy in a team. We contribute a learning algorithm that learns the weighted synergy graph using observations of performance of teams of only two and three agents. Further, we contribute two team formation algorithms, one that finds the optimal team in exponential time, and one that approximates the optimal team in polynomial time. We extensively evaluate our learning algorithm, and demonstrate the expressiveness of the weighted synergy graph in a variety of problems. We show our approach in a rich ad hoc team formation problem capturing a rescue domain, namely the RoboCup Rescue domain, where simulated robots rescue civilians and put out fires in a simulated urban disaster. We show that the weighted synergy graph outperforms a competing algorithm, thus illustrating the efficacy of our model and algorithms.

1. Introduction

Heterogeneous agents have varying capabilities that affect their task performance. We research on teams of such heterogeneous agents and how the performance of a team at a task relates to the composition of the team. Team performance has previously been computed as the sum of individual agent capabilities, e.g., the amount of resources an agent possesses [38,6]. In this work, we are interested in a model of team performance that goes beyond the sum of single-agent capabilities. We understand that there is *synergy* among the agents in the team, where team performance at a particular task depends not only on the individual agents' capabilities, but also on the composition of the team itself. Specific agents may have or acquire a high task-based relationship that allows them to perform better as a team than other agents with equivalent individual capabilities but a low task-based relationship. There are many illustrations of such synergy in real human teams, basically for any task. An example is an all-star sports team comprised of top players from around the

* Corresponding author. Somchaya Liemhetcharat has graduated from Carnegie Mellon University and is now at the Institute for Infocomm Research (I^{2}R), Singapore. Tel.: +65 6408 2000, fax: +65 6776 1378.

E-mail addresses: som@ri.cmu.edu (S. Liemhetcharat), veloso@cs.cmu.edu (M. Veloso).

0004-3702/$ – see front matter
http://dx.doi.org/10.1016/j.artint.2013.12.002

world, hence individual agents with high capabilities, who may have a lower synergy as a team and perform worse than a well-trained team of individuals with lower capabilities but much higher synergy.

To model task-based relationships, we introduce a connected weighted graph structure, where the vertices represent the agents, and the edges represent the task-based relationships. In such graphs, we define the level of synergy of a set of agents, as a function of the shortest path between agents. We further devise a non-binary metric of team performance based on a Gaussian model of the individual agent capabilities. Such probabilistic variables allow us to capture the inherent variability in team performance in a dynamic world. We show that our formulation of team performance captures many interesting characteristics, such as the effects of including new agents into the team.

Most existing team formation approaches assume that the agent capabilities are known *a priori* (e.g., [48]). We are motivated by research in ad hoc agents, that learn to collaborate with previously unknown teammates [40]. An ad hoc team is one where the agents in the team have not collaborated with each other. Assuming an ad hoc team, we address the team synergy learning question as: given a set of agents with unknown capabilities, how do we model and learn the capabilities and synergy of the agents through observations, in order to form an effective team, i.e., a subset of the agents? A solution to this problem will enable ad hoc teams to be applied to a variety of problems in the real world, where effective teams need to be composed from agents who may not have previously worked together.

A motivating scenario is the urban search-and-rescue (USAR) domain. Many USAR robots have been developed by different research groups, with a variety of hardware capabilities. When a disaster occurs, researchers from around the world arrive with their USAR robots. Due to safety and space constraints, only a subset of these robots may be able to be deployed to the site. Since many of these researchers have not collaborated in the past, selecting an effective team is an ad hoc problem, where the agent capabilities and synergy are initially unknown. Some of these robots may have been designed to work well with other robots developed by the same group, and in some cases, robots from different sources may have synergy in a team, e.g., a robot that clears rubble quickly so that another robot can search. Thus, it is necessary to model and learn the synergy of these robots and select the best team of robots to be deployed.

We contribute a learning algorithm that uses only observations of the performance of teams of two and three agents, in order to learn the agent capabilities and weighted graph structure of the weighted synergy graph. The learning algorithm iterates through weighted graph structures, and computes the agent capabilities using the observations. We also contribute two team formation algorithms that uses the learned weighted synergy graph to find an effective team that solves the task. Our approach does not make many assumptions about the agents, only that observations of their performance is available, and as such our approach is applicable to many multi-agent domains.

We perform extensive experiments to demonstrate that our learning algorithm effectively learns the structure of representative graph types and agent capabilities. We compare the weighted synergy graph to the unweighted synergy graph that we previously introduced [26], and demonstrate that the weighted synergy graph is more expressive and hence applicable to more domains. We apply the weighted synergy graph model to the RoboCup Rescue domain (that simulates rescue robots in a USAR scenario), and show that the learned weighted synergy graph is used to form a near-optimal team, and outperforms IQ-ASyMTRe [48], a competing algorithm.

In summary, the contributions of this work are:

1. A novel model of multi-agent team performance, the weighted synergy graph model, where agents are vertices in a connected weighted graph, edges represent how well agents work together, and agent capabilities are Normally-distributed variables;
2. The definition of the synergy of a multi-robot team as a function of the weighted synergy graph model;
3. A team formation algorithm that forms the optimal team in exponential time;
4. A team formation algorithm that approximates the optimal team in polynomial time;
5. A learning algorithm that learns a weighted synergy graph using only observations of agent teams comprising two and three agents;
6. Extensive experiments that evaluate our model and algorithms using synthetic data;
7. Application of the weighted synergy graph model to the RoboCup Rescue domain.

The article is organized as follows: Section 2 discusses related research in multi-robot task allocation, coalition formation, team formation, and ad hoc teams, and how it compares to our work. In Section 3, we formally define the weighted synergy graph model and our team formation algorithms. Section 4 contributes the synergy graph learning algorithm, while Section 5 presents extensive learning experiments. Section 6 compares the expressiveness of the weighted and unweighted synergy graph models. Section 7 details the experiments in the RoboCup Rescue domain, and Section 8 draws conclusions.

2. Related work

This section presents a review of related work, discussing the relevant domains of task allocation, coalition formation, ad hoc coordination and team formation.