

Analyzing the impact of human bias on human-agent teams in resource allocation domains

(Extended Abstract)

Praveen Paruchuri, Pradeep Varakantham*, Katia Sycara, Paul Scerri
Robotics Institute, Carnegie Mellon University, {paruchur,katia,pscerri}@cs.cmu.edu
*School of Information Systems, Singapore Management University, pradeepv@smu.ed.sg

ABSTRACT

A range of exciting proposed applications will involve many humans and agents working along side each other to achieve a complex objective. Domains for such applications include search and rescue[5], disaster response [8], military applications[1] and commerce[3]. Researchers envision automating the allocation of shared resources using algorithms such as distributed constraint optimization algorithms (DCOPs)[7]. For example, access to satellites, robots, computation or space may be automatically assigned using DCOPs. Due to the computational load and communication intensity of these algorithms, humans in the team will need to communicate their preferences and utilities to a proxy that executes the algorithm on their behalf. If there are many resources these preferences may be communicated incompletely or approximately. However, agents participating in the allocation process will be able to precisely and completely specify their preferences for all resources. The question addressed by this paper is what happens to the quality of the overall resource allocation when human and agent preference specifications differ in this way.

Preference elicitation is known to be a difficult problem that takes a lot of time and effort for humans[2]. When the preferences correspond to utilities, e.g., the value to the team to assign a particular robot to a particular human, a considerable computational burden is placed on the human since the utility of many resources must be computed by planning for each set of resources they might get and compare against other possible resource allocations. Many well known human biases will come into play, either consciously or sub-consciously, when reporting the utility of a resource. In this paper, we use a combination of empirical and theoretical analysis to understand the impact of these biases using a DCOP algorithm.

We study the effect of two commonly known models of human biases [4] on two different resource allocation problems using two allocation algorithms. The resource allocation problems chosen are representative of most resource allocation problems addressed in literature. To be able to compute solutions to large problems quickly, we used the Distributed Stochastic Algorithm (DSA) [10]. Our expectation is that there will be very little change to the overall utility on average. This is because the solution provided by DSA is dependent on the local ranking of resources and not on the absolute utilities of the resources to the agent.

Cite as: Analyzing the impact of human bias on human-agent teams in resource allocation domains (Extended Abstract), Praveen Paruchuri, Pradeep Varakantham, Katia Sycara and Paul Scerri, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. XXX-XXX.
Copyright © 2010, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

Categories and Subject Descriptors

I.2 [Computing Methodologies]: Artificial Intelligence

General Terms

Human factors

Keywords

Teamwork, DCOPs

1. PROBLEMS AND MODELS

In this section, we provide a brief background of the motivating problems, models used to represent the problems and the approach employed to solve the models. Specifically, we describe two generic distributed resource allocation problems from literature [9, 6] and then explain the DCOP framework used to represent these problems. We then formally represent the biases introduced by humans in specifying the preference function within the utility function of DCOP and finally present the DSA algorithm which is used to solve the DCOP models. In both these domains, the constraint functions represent both hard constraints (avoid scheduling one meeting in two different time slots etc.) and soft constraints (prefer to meet in mornings than afternoons etc.).

1.1 Problem 1: Discrete Resource Allocation

In this domain, resources need to be allocated to a group of agents and humans, E based on their preferences [9]. For ease of explanation, we assume only one resource is allocated to one agent or human. We model this domain as a DCOP as follows: Each agent/human has a variable, represented as e . The values that the variable can take correspond to the resource allocated to the agent. This domain for variable e is specified as D_e and belongs to the set $1, 2, \dots, R$. The utility for e when a resource, $r_e \in d_e$, is allocated to e is

$$u_e(r_e) = L * I_{r_e \neq 0} * (1 - \prod_{\tilde{e} \in E, \tilde{e} \neq e} I_{r_e \neq r_{\tilde{e}}}) + p_e(r_e), \quad (1)$$

where $p_e(r_e)$ is the preference value of agent e for resource r (soft constraint) and L is a large negative number that represents the penalty for allocating one resource to two different entities (hard constraint) and

$$I_{condition} = 1, \text{ if, } condition = true \\ = 0, \text{ otherwise}$$

1.2 Problem 2: Distributed Event Scheduling

In this domain, M meetings need to be scheduled for humans and agents (acting on behalf of humans). A meeting, m can require multiple humans, E_m and a human, e could be part of multiple

meetings, M_e . Furthermore, each human/agent, e has preferences over the time slots, p_e^t and meetings, p_e^m . Given this information, the goal is to compute a schedule which maximizes the utility of the team. This domain is modeled as a DCOP by [6] as follows:

- (a) Each agent/human, e has multiple variables, M_e , corresponding to all the meetings where e is required.
- (b) The values that each variable can take, correspond to any of the time slots where the meeting can be scheduled. This domain for variable $m \in M_e$ is specified as D_m^e and belongs to set $1, 2, \dots, T$, where T is the set of time slots available in which a meeting can be scheduled.
- (c) If a meeting, m for e is scheduled at t_e^m , then the utility for e is defined as

$$u_e^m(t_e^m) = L * (1 - \prod_{\tilde{t} \leq T, \tilde{t} \neq t_e^m, t_e^m \neq 0} I_{\tilde{t} \neq t_e^m}) * (1 - \prod_{\tilde{e} \in E_m, \tilde{e} \neq e} I_{t_e^m = t_{\tilde{e}}^m}) + (p_e(m) + p_e^m(t_e^m)) \quad (2)$$

Here, L is a large negative number that represents a penalty for scheduling the same meeting at two different time slots and other such cases. $(p_e(m) + p_e^m(t_e^m))$ represents the preferences of the humans (soft constraints).

1.3 Distributed Constraint Optimization

DCOP [7] is a popular framework for representing coordination in multiagent systems. A DCOP consists of n variables, v_1, v_2, \dots, v_n . Variable v_i can take on any value from the discrete finite domain D_i . The goal is to find an assignment, A of values to variables, such that the sum over a set of binary constraints and associated payoff or utility functions, $u_{i,j} : D_i \times D_j \rightarrow N$, is maximized. Formally, we maximize $\sum_{v_i, v_j} u_{i,j}(d_i, d_j)$, where $d_i \in D_i, d_j \in D_j$ and $\{v_i \leftarrow d_i, v_j \leftarrow d_j\} \in A$.

1.4 Bias 1: Simplification of preferences

As shown in [4], humans tend to simplify preference values when faced with problems where multiple factors need to be considered. One popular way of simplifying preferences is thresholding. Formally, this involves approximating the preference function over a variable v , $p_e^v()$ as a function that has zero corresponding to all but the top “k” values in its range. In general, it could be a “multi-step” step function and defined as follows (with ordering of thresholds given as $thres_1 > thres_2 > \dots$):

$$\begin{aligned} \tilde{p}_e^v(d) &= \max_{val} p_e^v(\hat{d}), \quad \text{if } p_e^v(d) > thres_1 \\ &= thres_1, \quad \text{if } thres_2 \leq p_e^v(d) \leq thres_1 \\ &= \dots \end{aligned}$$

Instead of a threshold, humans may simply limit themselves to specifying some number of top preferences. Specifically, we consider the scenario where humans specify preferences for only the most important resources while ignoring the rest. The modified function for this preference simplification is defined as:

$$\begin{aligned} \tilde{p}_e^v(d) &= p_e^v(d), \text{ if } d \in \max K(p_e^v()) \\ &= 0, \text{ otherwise} \end{aligned}$$

1.5 Bias 2: Preference Exaggeration

This bias of humans arises due to exaggeration of the importance of certain features over others. This involves increasing the preference function value of a variable v , $p_e^v()$ for the top preferred value assignments. The modified function for this type of preference exaggeration is defined as:

$$\begin{aligned} \tilde{p}_e^v(d) &= \mathcal{S} * p_e^v(d) + \mathcal{A}, \text{ if } d \in \max K(p_e^v()) \\ &= 0, \text{ otherwise} \end{aligned}$$

$\max K(p_e^v())$ provides the set containing the top k values in the range of the function $p_e^v()$, \mathcal{S} and \mathcal{A} are scaling and addition factors of exaggeration.

2. ALGORITHM FOR SOLVING DCOPS

Distributed Stochastic Algorithm (DSA) is the algorithm that we employ for solving DCOPs. It should be noted that DSA is the only algorithm that can realistically scale to the type of problems considered in this paper. The following key property of the DSA algorithm that will be very useful for our analysis.

PROPOSITION 1. *For DCOP problems 1 and 2, the solution provided by DSA is reliant on ranking of preference values and not on the actual preference values.*

3. IMPACT OF BIASES

Firstly, in domains where the ordering of preference values remains the same as the original preferences after human biases, we believe that according to Proposition 1 the impact will be zero. We further believe that even if there are disruptions to the preference order (due to human biases), DSA will provide solutions that are close to the optimal solution obtained with the original preference function. A key reason for this belief is that for large problem instances of interest, we expect to find that many solutions will have approximately the same utility.

4. REFERENCES

- [1] T. Cioppa, T. Lucas, and S. Sanchez. Military applications of agent-based simulations. In *Proceedings of the 36th conference on Winter simulation*, page 180, 2004.
- [2] Y. Guo, J. Muller, and C. Weinhardt. Elicitation of user preferences for multi-attribute negotiation. In *AAMAS*, page 1005, 2003.
- [3] R. Guttman, A. Moukas, and P. Maes. Agent-mediated electronic commerce: A survey. *The Knowledge Engineering Review*, 13(02):147–159, 2001.
- [4] J. Huber, D. Ariely, and G. Fischer. Expressing preferences in a principal-agent task: A comparison of choice, rating, and matching. *Organizational Behavior and Human Decision Processes*, 87(1):66–90, 2002.
- [5] H. Kitano and S. Tadokoro. Robocup rescue: A grand challenge for multiagent and intelligent systems. *AI Magazine*, 22(1):39, 2001.
- [6] R. Maheswaran, M. Tambe, E. Bowring, J. Pearce, and P. Varakantham. Taking DCOP to the real world: Efficient complete solutions for distributed multi-event scheduling. In *AAMAS*, pages 310–317, 2004.
- [7] P. Modi, W. Shen, M. Tambe, and M. Yokoo. Adopt: Asynchronous distributed constraint optimization with quality guarantees. *Artificial Intelligence*, 161(1-2):149–180, 2005.
- [8] R. Stranders, A. Farinelli, A. Rogers, and N. Jennings. Decentralised Coordination of Mobile Sensors Using the Max-Sum Algorithm. In *IJCAI*, 2009.
- [9] M. Yokoo, O. Etzioni, T. Ishida, and N. Jennings. *Distributed constraint satisfaction: foundations of cooperation in multi-agent systems*. 2001.
- [10] W. Zhang, G. Wang, Z. Xing, and L. Wittenburg. Distributed stochastic search and distributed breakout: properties, comparison and applications to constraint optimization problems in sensor networks. *Artificial Intelligence*, 161(1-2):55–87, 2005.