

A Practical Approach to Human/Multi-Robot Teams

Elizabeth Sklar¹, Simon Parsons¹, A. Tuna Özgelen², M. Q. Azhar²,
Todd Flyr², Eric Schneider¹, and Jeffery Raphael¹

¹Dept of Computer Science, University of Liverpool, UK

²Dept of Computer Science, Graduate Center, City University of New York, USA

contact authors: *e.i.sklar@liverpool.ac.uk*, *s.d.parsons@liverpool.ac.uk*

Abstract. Practical issues are often ignored in multiagent systems research, but can cause debilitating failure when some methodologies are implemented on physical multi-robot teams. For example, limited network bandwidth and intermittent connectivity can hinder communication between team members; memory and computational constraints of on-board controllers can render complex techniques inaccessible; and the need for quick decisions in dynamic environments favours reactive architectures and circumvents all but simple deliberation. In addition, when considering human-robot interfaces for multi-robot teams, further issues come to the fore, such as lack of understanding between an inexperienced human operator and robotic systems or information overload on the part of a human interacting with even a moderate number of robots streaming data in real time. The work presented here provides an overview of several approaches currently under investigation as part of the *HRTeam* project. The focus is on two key challenges: *coordinated task allocation* across a range of mission classifications, evaluated on low-end robot platforms; and *cooperative decision-making* between human and robot team members, evaluated with inexperienced human subjects.

1 Introduction

Our research addresses issues that are well-studied in *virtual* or *simulated multiagent systems (MAS)*, but present difficulties when implemented in *physical multi-robot systems (MRS)*. Our focus is on tasks that require coordinated exploration, in situations that could benefit from cooperative decision making and settings that need to be robust to dynamic changes in team composition. Our long term goal is to identify MAS approaches that are well-suited to MRS settings, as well as to devise approaches that address particular MRS challenges. Example application environments for our work include search and rescue [3, 29, 47] and humanitarian de-mining [8, 13]. We investigate some of the problems faced in these environments using a modified version of the Treasure Hunt Game [15].

Practical considerations dominate our research approach. We deploy a team of low-end robots and distribute exploration tasks amongst team members. Some

of the problems that are especially prevalent with inexpensive robot platforms include poor quality of images obtained from low-end cameras; limited speed and memory for processing on-board small platforms; and slow and intermittent network connectivity for team communication. Because transfer from a laboratory setting to the “real world” will require methods that perform well in the face of such practical difficulties, we take advantage of these opportunities to investigate solutions that are robust to such challenges, rather than try to eliminate them in our lab by using more expensive equipment.

We have developed the *HRTeam* software and hardware framework to support our experimental research in both physical and simulated settings [37]. MAS approaches can be rapidly prototyped, then calibrated and assessed in simulation; and finally validated on physical robots. Figure 1(a) shows an overhead view of the HRTeam experimental arena. A network of cameras is suspended over the arena to provide information about robots’ positions. Figure 1(b) shows a Surveyor Blackfin¹ robot. Each robot is mounted with a “hat”, marked with a Braille character that can be used to identify the robot from above [38]. A customised *robot controller*, built with Player [11], enables autonomous behaviours for each robot. A *visual debugger* window, shown in Figure 1(c), can be used to observe the current path planned by each robot controller. A 2D simulated version of this set-up is implemented with Stage [11], and a 3D simulated version, implemented in Blender², is under development and illustrated in Figure 1(d).

Our research focuses on two key challenges: *coordinated task allocation* across a range of mission classifications³ and *cooperative decision-making* between human and robot team members. Section 2 describes work on coordinated task allocation. Section 3 discusses work cooperative decision making. We close by outlining directions of current and future research.

2 Coordinated task allocation

Our long-term aim is to test a range of market-based coordination techniques from the MAS literature, to evaluate how techniques verified theoretically and tested in simulation perform in the rough-and-ready world of low-end physical robotics. Prior research on market-based approaches to multi-robot coordination have drawn on the Contract NET Protocol [41], applied to software agents, and the concept of *market-aware agents* [46]. Market-based approaches have been widely used in multi-robot systems to optimise resource usage, communication methodology and task-completion time [10]. A primary strength of market-based approaches is their reliance on local information and the self-interest of agents to arrive at efficient solutions to large-scale, complex problems that are otherwise intractable [5, 6]. The most common instantiations of market-based approaches in MAS and MRS are *auctions*, typically used for distributing (scarce) resources,

¹ <http://www.surveyor.com>

² <http://www.blender.org>

³ Missions are classified according to number of robots, need for repetition (e.g., maintenance vs one-time tasks), inter-task dependence, and robots-per-task requirements.

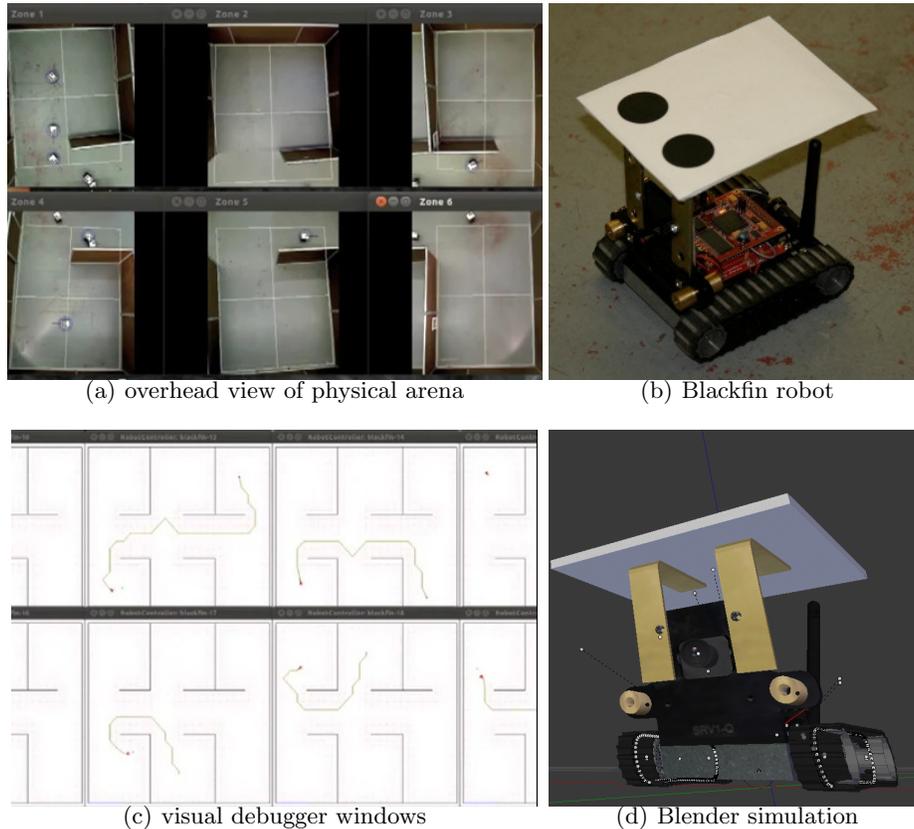


Fig. 1. HRTeam framework

where tasks or roles are treated as commodities and “auctioned” to agents. Based on individual preferences for a particular commodity, agents place “bids”. A significant body of work analyses, with respect to overall solution quality, the effects of different auction mechanisms [2, 20, 21, 41, 48], bidding strategies [24, 36, 44], dynamic task re-allocation or swapping [12], and levels of commitment to the contracts [26]. The application domains include loosely-coupled tasks, such as exploration [10, 23, 48], and tightly-coupled tasks [10, 17], such as box pushing and formation control [28]. In domains where there is a strong synergy between items, single-item auctions can result in suboptimal allocations [2]. Much of this prior work is evaluated in simulation and not with physical robots.

2.1 Approach

Our recent work has concentrated on the use of a simple market-based mechanism applied to a specific task: efficient exploration of a region with a team of

autonomous robots [30]. Given a set of n robots, we consider how best to have them visit m different positions, which we call *interest points*. The designation of interest points is an abstraction for a range of location-specific tasks, such as extinguishing a fire or clearing debris. We assume that by visiting an interest point, a robot has performed the task; future work will evaluate completing the task as well, but currently we focus on just getting there.

The baseline allocation method we consider is a greedy mechanism that takes each robot in turn and assigns it an interest point, in the same way that taxi ranks are often managed at airports. Both taxis and customers form a queue, and the first customer is assigned to the first taxi, the second customer to the second taxi, and so on. If the number of taxis exceeds the number of customers, then the surplus taxis wait in the queue until more customers arrive. Similarly if the number of customers exceeds the number of taxis, then the surplus customers wait. This mechanism might be satisfactory in a homogeneous multi-robot system where all robots start in roughly the same location and their destinations are evenly distributed, as with airport taxis on average. However, this mechanism exhibits clear inefficiencies in cases where the distributions of numbers and locations of robots and/or interest points are lumpy. For example, if $m > n$, implying that some robots would have to explore several points, a balanced mechanism should allocate points so that all robots travel approximately the same total distance.

We compare this “greedy-taxi” mechanism to a simple auction mechanism that attempts to balance, across the team, the robots’ estimated costs to complete all tasks. Given robots’ initial locations and a list of interest-point locations, robots bid for points. Bids are determined by the robots’ distances to the points, as calculated by an A* path-planner with a map of the area. Points are allocated one at a time, in a sequence of auctions. In each auction, a robot considers its last position in its bid for its next point; for its first point, this will just be its starting location, whereas for subsequent points, this will be the last point it “won” in a previous auction. Such a robot estimates the cost to travel from its most-recently-allocated interest point to the new point. A robot not yet allocated any points estimates its travel cost from its initial position, before the start of the first auction. In some situations, this approach will clearly be less efficient in total distance traveled than a *combinatorial* auction which would allow robots to bid on bundles of locations [19]. The simple auction, however, will likely be more efficient in terms of computational effort, given the well-documented computational cost associated with combinatorial auctions [9, 35]. Hence, simple auctions are likely more practical in a real-time, dynamic environment.

2.2 Experimental results

We ran a series of experiments, measuring results in five different scenarios. A *scenario* is a mission defined by a specific set of parameters: the number of robots on the team (n), the starting locations for the robots, the number of interest points to visit (m), and the locations of the interest points. All scenarios tested here involved $n = 4$ robots, $m = 20$ interest points, one fixed set of

starting locations, and five different sets of interest point locations. Experiments were conducted with each scenario using the two task allocation mechanisms discussed above (simple auction and greedy-taxi, labelled A and G, respectively), five times in the physical environment and five times in the simulated environment (labelled P and S, respectively)—100 experimental trials in all.

Each experiment recorded the following metrics. **Run time** is the elapsed time between the start of an experiment and the time that the last robot completed the tasks allocated to it. Run time varies with the scenario and the way that tasks are distributed amongst robots. **Deliberation time** is the elapsed time between the start of an experiment and the completion of the task allocation mechanism, i.e., the point at which the robots began to execute tasks (move to the interest points allocated to them). Deliberation time varies with the scenario and the allocation mechanism. **Distance travelled** is the total length of the path traveled by the robot, i.e., the sum of the Euclidean distances between consecutive position updates. Some robots complete their assigned tasks before other robots. **Idle time** is the time that elapses between when a robot completes its last task and when all robots have completed all the tasks. Idle time varies from one robot to the next; it is affected by the way that tasks are distributed amongst the robots and also by the scenario. It provides an estimate of the spread of tasks between robots. Finally, with several robots moving around a limited space, robots can easily get in each other’s way. Our system detects situations where a collision is likely: then stops one robot and gives the other the right of way. The number of such **near collisions** is tallied, as well as the amount of **delay time**, the total time that robots spend stationary as a result of collision avoidance. Delay time varies from one robot to the next, as well as with the way the tasks are distributed.

Figure 2, which considers the metrics summed over all the different experimental scenarios, provides the headline results. Since our main interest is in the performance of coordination mechanisms on physical robots, what we are most interested in is the comparison between the PA (physical robots, task allocation by simple auction) and PG (physical robots, task allocation by greedy-taxi) results. The simple auction mechanism takes slightly longer to deliberate and provide an allocation, illustrated in Figure 2(b). However, on all other metrics, the simple auction mechanism outperforms the greedy-taxi mechanism. Since the auction considers the distance that the robots travel, one might well expect that the run time, illustrated in Figure 2(a), for PA would be faster than the run time for PG. The total distance traveled, illustrated in Figure 2(c), is greater for greedy-taxi than simple auction because the bidding strategy employed for the auction mechanism attempts to minimise distance. More interestingly, the idle time, illustrated in Figure 2(d), also improved, which suggests that the tasks are more evenly distributed. (Recall that idle time shown here is the sum across the entire team.) In addition, the auction mechanism leads to less delay time, illustrated in Figure 2(e), and fewer near collisions, illustrated in Figure 2(f). (Delay time and near collisions are related since a near collision leads to one robot stopping, and hence an increase in delay time.) This suggests that the

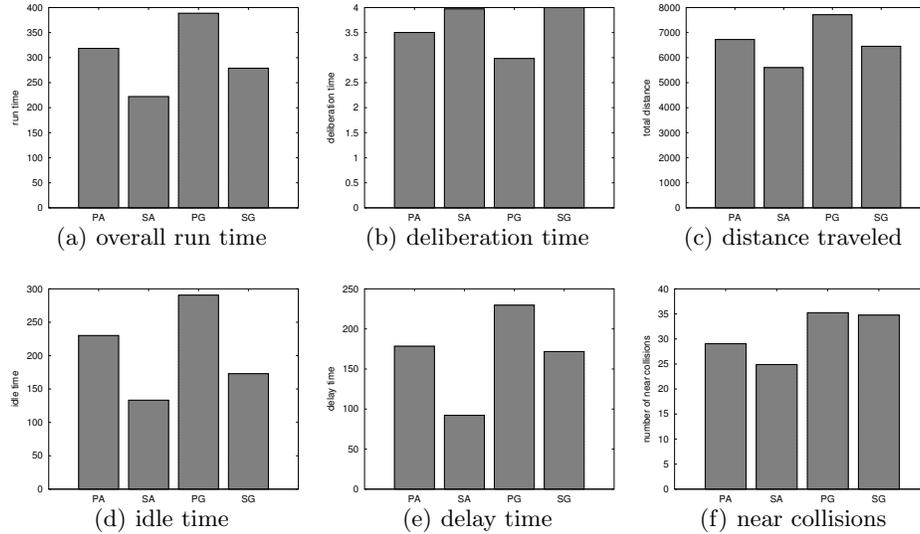


Fig. 2. Combined results for all scenarios. Values for each metric distinguish between the physical (P) and simulated (S) experiments, and between the simple auction (A) and greedy-taxi (G) task allocation.

auction mechanism spreads the robots more evenly throughout the space. Note that the results computed on the runs conducted in simulation reflect the same findings.

Figure 3 illustrates paths traveled by the robots in the physical and simulated environments for a single scenario using both allocation mechanisms. The top two graphs show the simple auction allocation mechanism and the bottom two graphs show the greedy-taxi allocation mechanism. Comparing the two graphs in the top row, where physical and simulated robots received the same distribution of interest points, it becomes obvious why the physical robots travelled further: the paths are clearly more irregular, while the simulated robots’ paths are smoother. This is due to errors in the camera-based localisation and noise from the robot’s motors that cause the robots to move in jerky motions, especially when turning. Comparing the pairs of graphs in each column, especially for the simulated runs (rightmost column), it is obvious that the robot team using the simple auction allocation mechanism traced shorter and more efficient paths than the team using the greedy-taxi allocation mechanism.

3 Cooperative decision making

Humans interact with each other in a range of relationships, some of which are subordinate, such as boss-employee or parent-child, and others are collabora-

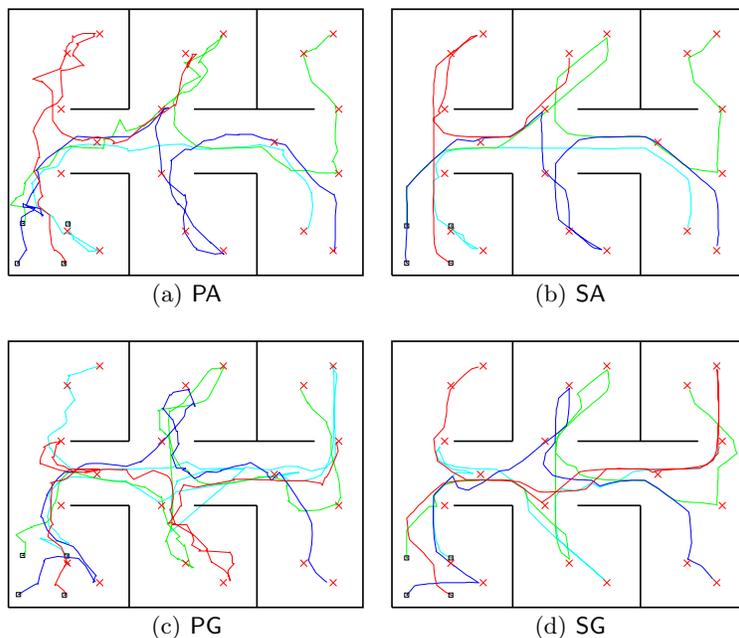


Fig. 3. Sample robot trajectories in the physical (left) and simulated (right) environments, using task allocation by simple auction (top) and greedy-taxi (bottom) mechanism.

tive, such as two lumberjacks each holding one end of a two-man cross-cut saw or two software engineers engaged in pair-programming. In many productive human-human relationships, the skills of each human complement each other, for example, a graphic designer and a web programmer collaborating on building a web site, or a composer and lyricist collaborating on writing a hit song. Each of these relationships, to be successful and productive, relies on some amount of communication—*dialogue*—in which each party presents their ideas, which are discussed together, and a shared conclusion is agreed upon by both parties.

In contrast, the vast majority of human-robot relationships are ones in which the human is the master and tells the robot what to do. For example, a workplace robot might be asked to deliver letters or packages in an office building, a nurse robot might be asked to dispense medication to a patient, or a search-and-rescue robot might be asked to explore a region inside a damaged building. In each of these cases, the need for communication is limited to the human commanding or requesting a robot to perform a task, and the robot reporting its findings to the human. There is no exchange of ideas, hence no dialogue in the formal sense described above.

This type of limited exchange puts many restrictions on the human-robot interaction. For example, if a robot fails at its assigned task, it can only report to the human that it has failed; it cannot discuss the reasons for failure or possible follow-on courses of action—as two humans would when collaborating on a task. Or, if a robot disagrees with its assigned task, perhaps because it knows of a reason why the task may fail, *a priori*, before even attempting the task, the robot has no way of explaining this to the human. Additionally, the human-robot team is constrained by the human’s scope of information and ideas: the robot cannot recognise new or unexpected opportunities and interrupt its task to suggest an alternate activity.

Dialogue that is founded on unscripted and opportunistic exchange of ideas does not exist in today’s *human-robot interaction (HRI)* systems. The current focus in most human-robot dialogue work is on natural language architectures [25] or delivery methods [18, 22, 31, 43], rather than dynamic content selection. For human-robot systems to be truly collaborative, participants need to be able to engage in constructive dialogue that can adjust dynamically as the dialogue and situation unfolds. *Argumentation* is a well-founded theoretical method that can support such needs. *Argumentation-based dialogue* can be used to handle the kinds of example situations mentioned above: recovering from failure, pre-empting failure, and revising plans dynamically.

Our work on cooperative decision making fills in missing details for combining and implementing theoretical models of logical argumentation and argumentation-based dialogue in a dynamic, real-time, human-robot setting. While there is a large literature discussing the theoretical underpinnings of argumentation [33, 34] and dialogue [27], the only implemented systems are off-line decision-making tools [7, 16]. In contrast, our system contains an end-to-end solution, which is necessary for an actual implementation, and addresses questions such as how and when to update the beliefs of an agent engaged in a real-time dialogue, and how and when to initiate a dialogue.

3.1 Approach

We are concerned with situations in which a robot and human cooperatively decide on goals to achieve and plans to achieve their goals. Our approach engages human and robot in an argumentation-based *dialogue game* [27], where they exchange locutions according to a protocol in order to reach agreement about goals and plans. The argumentation-based dialogue protocols we have implemented include: *information-seeking* [45], where one agent asks another agent a question that it believes the other can answer; *inquiry* [14], where two agents collaboratively answer a question that neither knew the answer to beforehand; and *persuasion* [32], where an agent tries to alter the beliefs of another agent. The argumentation theory underlying our implementation, which we call *ArgHRI*, has been documented in [39, 40].

To test ArgHRI, we have built a user interface, illustrated in Figure 4, which is integrated as a module in HRTeam. The interface does not interpret or process natural language. Instead, dialogue is supported through labels and icons

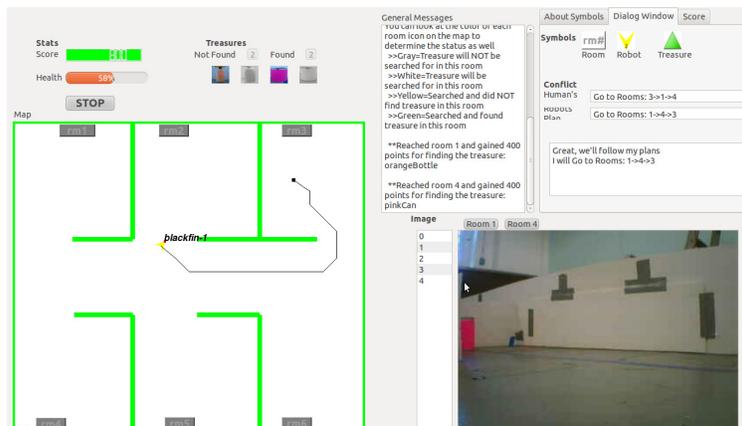


Fig. 4. ArgHRI user interface

in the interface. For testing purposes, the interface supports two modes: **minimal** dialogue, where the user selects a plan and sends it to the robot; and **full** dialogue, where the user creates a plan in collaboration with the robot, through a series of dialogue panels. Using the **belief** dialogue panel, the user can inform the robot about her beliefs (related to the mission and the robot’s environment). Using the **plan** dialogue panel, the user can recommend a plan for addressing mission goals. The robot also formulates its own plan, typically based on its current location and battery power level. After the user and robot have created plans, ArgHRI invokes the *ArgTrust* [42] argumentation engine to determine if any conflicts exist. If so, then the **conflict** dialogue panel is displayed, where the user can receive explanatory information from the robot, to resolve discrepancies in beliefs and/or plans. If, after seeing the arguments for why the robot has selected its plan, the human agrees with the robot’s reasoning, then the robot executes its own plan. Otherwise, the robot executes the human’s plan.

ArgHRI sends an XML file to ArgTrust, which describes an agent’s set of beliefs about the world, about its environment, about other agent(s) with whom it is interacting, represented as facts, predicates and rules; and a query to address with respect to the agent’s beliefs. ArgTrust outputs another XML file with an answer to the input query, highlighting any conflicts that exist in the set of beliefs relevant to the query. ArgTrust uses *argument labeling* [4]; any query is answered “yes” if the conclusion supporting the query is IN, and “no” if the argument is OUT. Since it is typically not useful to know only that an argument was rebutted, in order to properly inform the user about why there is an issue with their plan, the ArgTrust output file includes which rules or beliefs caused conflict(s).

Our experimental setup employs a modified version of the Treasure Hunt game using the arena shown in Figure 1. The human operator, or “player”, is

physically separated from the robots so that she cannot view the arena and the only information she has is provided by the robots sensors (e.g., cameras and range sensors). Multiple “treasures” are placed in the arena in secret locations, and the human-robot team’s mission is to find all the treasures. The human and robot can interact to decide which rooms in the arena to search. The robots can transmit pictures of the rooms back to the human, to identify treasures. The game is scored using a point system: when the robot searches a room, the team (human+robot) earns points if a treasure is found, but loses otherwise. Other constraints are that the robot has a limited amount of “health” points (i.e., battery power), which run out as the robot travels; so the robot will not be able to travel to all the rooms.

3.2 Pilot study results

We conducted a pilot study to test the operation of our initial system and assess participants’ views about the helpfulness of argumentation-based dialogue in this game [1]. The study was conducted as follows. One robot was placed in the physical HRTeam arena, and the human participant interacted with the robot, using the interface shown in Figure 4. Each trial followed the same procedure: (1) provide instruction to human participant; (2) human participant completes pre-survey online; (3) human and robot play the Treasure Hunt game using minimal dialogue; (4) human completes mid-survey online; (5) human and robot play the Treasure Hunt game using full dialogue; and (6) human completes post-survey. Six undergraduates, inexperienced with the ArgHRI system and unfamiliar with the related research, participated in the study.

Each pre-, mid- and post-survey consisted of multiple questions, phrased in different ways, asking about five different topics. The questions were designed to investigate participants’ perceptions about each topic, as follows: **collaboration** questions investigated if a robot with full dialogue mode is a more reliable teammate, as compared to minimal dialogue mode; **trust** questions investigated if a robot with full dialogue mode increases the robot’s credibility as a teammate, as compared to minimal dialogue mode; **dialogue** questions investigated if a robot with full dialogue mode offers a satisfying means of communication between teammates, as compared to minimal dialogue mode; **performance** questions investigated if a robot with full dialogue mode improves the team’s performance, as compared to minimal dialogue mode; and **effort** questions investigated if a robot with full dialogue mode makes it easier for the human to collaborate with the robot, as compared to minimal dialogue mode. Each survey was administered online, as part of ArgHRI, and answers were supplied on a 20-point Likert scale.

Survey results are shown numerically in Figure 5. All participants favoured full dialogue mode over minimal dialogue mode. For all five topics covered in the surveys, positive changes were recorded for full dialogue mode as compared to minimal dialogue mode. The largest positive increase during full dialogue mode was in the collaboration (i.e., 24.58%) followed by dialogue (i.e., 22.79%). The participants also found the robot more trustworthy in full dialogue mode

compared to minimum dialogue. As a whole, the study confirmed the operation of ArgHRI and indicated that users understood how to use the system and felt comfortable employing full dialogue mode to resolve conflicts during human-robot collaboration.

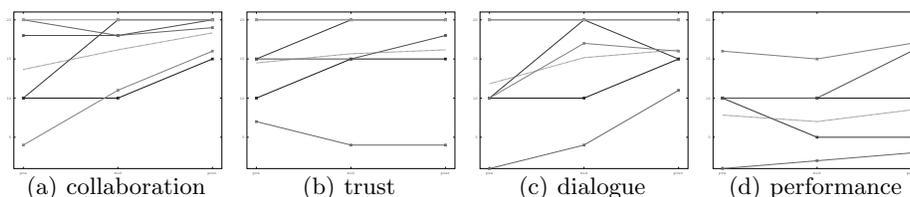


Fig. 5. Pilot study results. Each line plots the responses from one participant. The y-axis contains ratings: from 1 (worse) to 20 (better). The x-axis contains three points: pre-survey, mid-survey and post-survey, from left to right, respectively. Individual numbers don't matter—the trends are important.

4 Summary

Practical issues are often ignored in multiagent systems research, but present significant challenges to multi-robot applications. The work described here implements and evaluates proven agent-based and multiagent techniques on physical robots: market-based mechanisms for coordinated task allocation and argumentation-based dialogue for cooperative decision making. The latter was assessed in a human-robot domain and tested with inexperienced human subjects.

Research activities in both focus areas—coordinated task allocation and cooperative decision making—is ongoing. New results that implement and assess a broader base of auction mechanisms are forthcoming, as is a more significant user study with a larger number of human subjects and a more complex set of missions. Other related work in the HRTeam project investigates novel approaches to learning in multi-robot and human-robot environments.

Acknowledgments

This work was supported by the National Science Foundation under #CNS-0851901, #CNS-0520989, #IIS-1116843 and #IIS-1117000. Funding for this work was provided in part by the National Science Foundation under #11-16843 and by the US Army Research Office under the Science of Security Lablet grant. The authors would like to thank Susan L. Epstein for collaboration on architectural aspects of HRTeam not discussed here; Ofear Balas for his improvements to our path planner and auction manager software; and David Lettier for his work on the Blender simulator.

References

1. Azhar, M.Q., Schneider, E., Salvit, J., Wall, H., Sklar, E.: Evaluation of an argumentation-based dialogue system for human-robot collaboration. In: Proceedings of the Workshop on Autonomous Robots and Multirobot Systems (ARMS) at Autonomous Agents and MultiAgent Systems (AAMAS). St Paul, MN, USA (May 2013)
2. Berhaut, M., Huang, H., Keskinocak, P., Koenig, S., Elmaghraby, W., Griffin, P.M., Kleywegt, A.: Robot Exploration with Combinatorial Auctions. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, Las Vegas. pp. 1957–1962 (2003)
3. Burke, J.L., Murphy, R.R.: Human-robot interaction in user technical search: Two heads are better than one. In: International Workshop on Robot and Human Interactive Communication. pp. 307–312 (2004)
4. Caminada, M.W.A.: On the issue of reinstatement in argumentation. In: Proceedings of the 10th European Conference on Logic in Artificial Intelligence. pp. 111–123 (2006)
5. Dias, M.B., Stentz, A.: A Free Market Architecture for Distributed Control of a Multirobot System. In: Proceedings of the 6th International Conference on Intelligent Autonomous Systems (IAS-6). pp. 115–122 (July 2000)
6. Dias, M.B., Zlot, R., Kalra, N., Stentz, A.: Market-based multirobot coordination: A survey and analysis. Tech. Rep. CMU-RI-TR-05-13, Carnegie Mellon University (2005)
7. Fabregues, A., Sierra, C.: Dipgame: a challenging negotiation testbed. *Engineering Applications of Artificial Intelligence* 24(7), 1137–1146 (2011)
8. Freese, M., Matsuzawa, T., Oishi, Y., Debenest, P., Takita, K., Fukushima, E.F., Hirose, S.: Robotics-assisted demining with gryphon. *Advanced Robotics* 21(15) (2007)
9. Fujishim, Y., Leyton-Brown, K., Shoham, Y.: Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In: Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI). Stockholm (1999)
10. Gerkey, B.P., Mataric, M.J.: Sold!: Auction methods for multi-robot control. *IEEE Transactions on Robotics and Automation Special Issue on Multi-Robot Systems* 18(5) (2002)
11. Gerkey, B., Vaughan, R.T., Howard, A.: The Player/Stage Project: Tools for Multi-Robot and Distributed Sensor Systems. In: Proc of Int'l Conf on Advanced Robotics (ICAR) (2003)
12. Golfarelli, M., Maio, D., Rizzi, S.: A Task-Swap Negotiation Protocol Based on the Contract Net Paradigm. Tech. Rep. 005-97, DEIS, CSITE - Università di Bologna (1997)
13. Habib, M.K.: Humanitarian Demining: Reality and the Challenge of Technology. *International Journal of Advanced Robotic Systems* 4(2), 151–172 (2007)
14. Hulstijn, J.: Dialogue Models for Inquiry and Transaction. Ph.D. thesis, Universiteit Twente (2000)
15. Jones, E.G., Browning, B., Dias, M.B., Argall, B., Veloso, M., Stentz, A.: Dynamically Formed Heterogeneous Robot Teams Performing Tightly-Coordinated Tasks. In: Proc of Int'l Conf on Robotics and Automation (ICRA) (2006)
16. Judson, P.N., Fox, J., Krause, P.J.: Using new reasoning technology in chemical information systems. *Journal of Chemical Information and Computer Sciences* 36, 621–624 (1996)

17. Kalra, N., Stenz, A.: A Market Approach to Tightly-Coupled Multi-Robot Coordination: First Results. In: Proceedings of the ARL Collaborative Technologies Alliance Symposium (2003)
18. Kirby, R., Broz, F., Forlizzi, J., Michalowski, M., Mundell, A., Rosenthal, S., Sellner, B., Simmons, R., Snipes, K., Schultz, A., Wang, J.: Designing robots for long-term social interaction. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 2199–2204 (2005)
19. Koenig, S., Keskinocak, P., Tovey, C.: Progress on agent coordination with cooperative auctions. In: Proceedings of the AAAI Conference on Artificial Intelligence (2010)
20. Koenig, S., Tovey, C., Lagoudakis, M., Kempe, D., Keskinocak, P., Kleywegt, A., Meyerson, A., Jain, S.: The Power of Sequential Single-Item Auctions for Agent Coordination. In: Proceedings of National Conference on Artificial Intelligence (AAAI). pp. 1625–1629 (2006)
21. Kraus, S.: Automated negotiation and decision making in multiagent environments. In: Proceeding of MultiAgent Systems and Applications (ACAI-EASSS). pp. 150–172. Springer-Verlag (2001)
22. Krestin, F.: How people talk with robots: Designing dialogue to reduce user uncertainty. *AI Magazine* 32(4) (2011)
23. Lagoudakis, M., Berhault, M., Koenig, S., Keskinocak, P., Kelywegt, A.: Simple auctions with performance guarantees for multi-robot task allocation. In: Proceedings of the International Conference on Intelligent Robotics and Systems (IROS) (2004)
24. Lagoudakis, M., Markakis, V., Kempe, D., Keskinocak, P., Koenig, S., Kleywegt, A., Tovey, C., Meyerson, A., Jain, S.: Auction-based multi-robot routing. In: Proceedings of Robotics: Science and Systems Conference (2005)
25. Lemon, O., Gruenstein, A., Peters, S.: Collaborative activities and multi-tasking in dialogue systems. *Traitement automatique des langues (TAL)* 43(2), 131–154 (2002)
26. Mataric, M., Sukhatme, G., Ostergaard, E.: Multi-robot task allocation in uncertain environments. *Autonomous Robots* (2003)
27. McBurney, P., Parsons, S.: Games that agents play: A formal framework for dialogues between autonomous agents. *Journal of Logic, Language, and Information* 11(3) (2002)
28. Michael, N., Zavlanos, M.M., Kumar, V., Pappas, G.J.: Distributed multi-robot task assignment and formation control. In: Proceedings of the IEEE International Conference on Robotics and Automation (2008)
29. Nourbakhsh, I.R., Sycara, K., Koes, M., Yong, M., Lewis, M., Burion, S.: Human-Robot Teaming for Search and Rescue. *Pervasive Computing* January-March (2005)
30. Özgelen, A.T., Schneider, E., Sklar, E., Costantino, M., Epstein, S.L., Parsons, S.: A first step toward testing multiagent coordination mechanisms on multi-robot teams. In: Proceedings of the AAMAS Workshop on Autonomous Robots and Multirobot Systems (ARMS) (2013)
31. Peltason, J., Wrede, B.: The curious robot as a case-study for comparing dialog systems. *AI Magazine* 32(4) (2011)
32. Prakken, H.: Formal systems for persuasion dialogue. *Knowledge Engineering Review* 21(2) (2006)
33. Prakken, H.: An abstract framework for argumentation with structured arguments. *Argument and Computation* 1 (2010)

34. Rahwan, I., Simari, G.R. (eds.): *Argumentation in Artificial Intelligence*. Springer Verlag (2009)
35. Sandholm, T.: Algorithm for Optimal Winner Determination in Combinatorial Auctions. *Artificial Intelligence* 135, 1–54 (2002)
36. Sariel, S., Balch, T.: Efficient bids on task allocation for multi-robot exploration. In: *Proceedings of the Nineteenth Florida Artificial Intelligence Research Society Conference* (2006)
37. Sklar, E.I., Özgelen, A.T., Muñoz, J.P., Gonzalez, J., Manashirov, M., Epstein, S.L., Parsons, S.: Designing the HRTeam framework: Lessons learned from a rough-'n-ready human/multi-robot team. In: *Proceedings of the AAMAS Workshop on Autonomous Robots and Multirobot Systems (ARMS)* (2011)
38. Sklar, E.I., Özgelen, A.T., Schneider, E., Costantino, M., Muñoz, J.P., Epstein, S.L., Parsons, S.: On transfer from multiagent to multi-robot systems. In: *Proceedings of the AAMAS Workshop on Autonomous Robots and Multirobot Systems (ARMS)* (2012)
39. Sklar, E., Azhar, M.Q., Flyr, T., Parsons, S.: A Case for Argumentation to Enable Human-Robot Collaboration. In: *Proceedings of the Workshop on Argumentation in Multiagent Systems (ArgMAS) at Autonomous Agents and MultiAgent Systems (AAMAS)*. St Paul, MN, USA (May 2013)
40. Sklar, E., Azhar, M.Q., Parsons, S., Flyr, T.: A Case for Argumentation to Enable Human-Robot Collaboration (Extended Abstract). In: *Proceedings of Autonomous Agents and Multiagent Systems (AAMAS)*. St Paul, MN, USA (May 2013)
41. Smith, R.G.: The Contract Net Protocol: High-Level Communication and Control in a Distributed Problem Solver. In: Bond, A.H., Gasser, L. (eds.) *Distributed Artificial Intelligence*, pp. 357–366. Morgan Kaufmann Publishers Inc. (1988)
42. Tang, Y., Cai, K., McBurney, P., Sklar, E., Parsons, S.: Using argumentation to reason about trust and belief. *Journal of Logic and Computation, Special Issue on Agreement Technologies* 22(5), 959–1018 (2012)
43. Thomaz, A., Chao, C.: Turn taking based on information flow for fluent human-robot interaction. *AI Magazine* 32(4) (2011)
44. Tovey, C., Lagoudakis, M.G., Jain, S., Koenig, S.: Generation of Bidding Rules for Auction-Based Robot Coordination. In: *Proceedings of the 3rd International Multi-Robot Systems Workshop, Naval Research Laboratory, Washington DC* (March 2005)
45. Walton, D.N., Krabbe, E.C.W.: *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. State University of New York Press, Albany, NY, USA (1995)
46. Wellman, M.P., Wurman, P.R.: Market-Aware Agents for a Multiagent World. *Robotics and Autonomous Systems* 24, 115–125 (1998)
47. Yanco, H., Baker, M., Casey, R., Keyes, B., Thoren, P., Drury, J.L., Few, D., Nielsen, C., Bruemmer, D.: Analysis of Human-Robot Interaction for Urban Search and Rescue. In: *Proceedings of the IEEE International Workshop on Safety, Security and Rescue Robotics (SSRR)* (2006)
48. Zlot, R., Stentz, A., Dias, M.B., Thayer, S.: Multi-robot exploration controlled by a market economy. In: *Proceedings of the IEEE Conference on Robotics and Automation* (2002)