Toward Human/Multi-Robot Systems to Support Emergency Services Agencies

Elizabeth I. Sklar[1] and **Eric Schneider**[1]
[1]Dept of Computer Science,
University of Liverpool, UK

A. Tuna Özgelen[2] and **M. Q. Azhar**[2]
[2]Dept of Computer Science,
The Graduate Center,
City University of New York, USA

Abstract

The ability to make decisions that balance conflicting needs and variable-quality inputs is a skill that is inherently human. In emergency situations, such capabilities are tested under pressure, as needs and inputs change—often rapidly—and deliberation must take place quickly or else opportunities are lost. This short paper identifies challenges faced when emergency services personnel are supported by human/multi-robot systems. Several strategies are proposed to address these challenges, with deployment geared toward emergency services agencies within the next 5-10 years.

Introduction

Many problems face today's emergency services organisations. Budget cuts reduce personnel and access to state-of-the-art equipment, while new technologies are emerging and agencies are under pressure to modernise. Meantime, officials are challenged by citizens' ready access to and participation in a wide range of information sources enabled through social media. When large-scale incidents occur, personnel from multiple agencies are often thrown together in *ad hoc* teams, led by senior officers who are ultimately held accountable for their decisions.

The approaches outlined here aim to support emergency services by providing intelligent systems that collect and manage data during a critical incident. Multi-robot teams can assist *operational* ("bronze") personnel at the scene by providing surveillance information and communicating with victims trapped in remote locations. Agent-based coordination mechanisms can assist *tactical* ("silver") personnel charged with real-time allocation of resources. Logical reasoning agents can assist *strategic* ("gold") personnel via intelligent, data-backed decision support tools. Our vision is that intelligent technologies will improve response time and overall success rates in the field. However, many significant challenges lie between approaches demonstrated in laboratories and practical solutions ready to be deployed by emergency services agencies. In this short paper, we identify several challenges for the AI, Robotics and Agents communities and propose strategies to address these challenges, with deployment geared toward emergency services agencies within the next 5-10 years.

Challenges

During an emergency incident, the top-priority goal is to identify and rescue victims, reducing or eliminating loss of human life. We highlight four challenges for deployment of human/multi-robot systems in real-world emergency services operations: (1) *multi-robot coordination*, (2) *human-multi-robot collaboration*, (3) *human-robot agreement*, and (4) *data-backed deliberation*.

Multi-Robot Coordination

The idea of deploying a team of robots to assist first responders in search-and-rescue operations is not new (Wong, Seet, and Sim 2011). The robotics community has investigated various aspects of multi-robot teamwork, including: *multi-robot routing* (Dias et al. 2006; Koenig et al. 2006), *collision avoidance* (Guy, Lin, and Manocha 2010; Kimmel, Dobson, and Bekris 2012), *distributed sensing* (Howard, Matarić, and Sukhatme 2002; Parker 2002; Onosato et al. 2006), and *task allocation* (Matarić, Sukhatme, and Ostergaard 2003; Liu and Nejat 2013). We focus on task allocation, for exploration and surveillance, and adapt methods derived in the multi-agent systems community, primarily market-based mechanisms which have the advantage of a strong theoretical basis (Wellman and Wurman 1998; Feigenbaum and Shenker 2002; Parsons, Rodríguez-Aguilar, and Klein 2011). We evaluate these methods using *deployment metrics*, such as the amount of time it takes to complete the tasks once they have been allocated and the amount of time robots spend avoiding collisions in a crowded environment. In contrast, assessment of mechanism design methods in the multi-agent systems community is limited to measuring computational properties of the allocation procedure. Different task allocation methods can be used by teams to self-organise and execute a range of independent and constrained (dependent) tasks in static and dynamic environments (Gerkey and Matarić 2004; Landén, Heintz, and Doherty 2012). We have demonstrated that the performance of these allocation mechanisms is sensitive to robots' starting conditions and congestion of task locations, and we have showed that our deployment-oriented performance metrics rank mechanisms differently compared with measures based on computational properties of the allocation algorithms (Schneider et al. 2014a; 2014b)—a key result that can help officers select which methods to deploy in different situations.

Human/Multi-Robot Collaboration

A small number of studies address the human factors involved in deployment of rescue robots in emergency situations (e.g., (Murphy 2004)). Some approaches focus on increasing the autonomy of the robot team (Doroodgar et al. 2010; Liu and Nejat 2013), reducing the amount of direct human control required and improving situational awareness. However, a fully autonomous multi-robot team is not realistic: in any real-world deployment, a human will ultimately be responsible for the robot team's actions, typically a senior officer at tactical or strategic level. If this officer does not understand how the robot team is organised, e.g., autonomous allocation of exploration tasks to robots, then the officer's *trust* in the solution will be diminished. We are investigating human-centric approaches to semi-autonomous task allocation in real-time, dynamic environments containing a range of independent and constrained tasks. The members of our robot team can operate autonomously to complete navigation and sensing tasks, but the distribution of tasks to robots can either be performed autonomously (e.g., using a market-based mechanism), or manually (by a human, at run-time), or *semi-autonomously* (where a human assigns some of the tasks and then lets the robots self-organise to assign the rest). We have defined a human-centric representation of complex task environments, called the *Task Assignment Graph (TAG)*, and validated this representation in a user study (Özgelen and Sklar 2014a; 2014b). Our results showed that self-reported mental demand for a human charged with manual task allocation is higher when scenario complexity increases due to the number of available choices; consequently, humans tend to assign constrained tasks first and then independent tasks.

Human-Robot Agreement

In traditional human-robot systems, the human takes the initiative by providing supervisory control (Goodrich and Schultz 2007). Mixed initiative approaches, like *adjustable autonomy* (Tambe, Scerri, and Pynadath 2002), offer a means for alternating the locus of control, but the human and robot do not share control or discuss decisions. HRI work on discussion is largely about natural language architectures (Lemon, Gruenstein, and Peters 2002), delivery methods (Bohus et al. 2011), and social intelligence (Breazeal 2003; Dautenhahn 2007). In contrast, our focus is on human and robot reaching *agreement*. We have defined a formal structure (Sklar et al. 2013), based on *logical argumentation* (Rahwan and Simari 2009; Prakken 2010) and *dialogue games* (Walton and Krabbe 1995; McBurney and Parsons 2002; Prakken 2006). Such games are characterised by rules that govern what each participant in the dialogue can say, based on their beliefs and the game state. Using our *ArgHRI* formalism, a human and robot can argue in favour of a decision (conclusion), or against the other's decision, by presenting evidence that supports (or refutes) the decision. Our mixed initiative approach allows the locus of control to shift dynamically based on the belief and game states. A preliminary user study (Azhar et al. 2013) tested implementation of our approach, and another is forthcoming.

Data-backed Deliberation

The use of new technologies for gathering data during critical incidents is becoming commonplace, but emergency services personnel are not prepared to manage the quantity and range of information streams available, for example: sensor data from multi-robot teams, status reports from operational personnel belonging to different agencies, and on-the-scene updates from citizens posting on social media. We have developed a decision-making tool that helps users balance information from different sources, of varying quality and *trustworthiness*. This tool is based on the *ArgTrust* (Tang et al. 2012) engine, which constructs data-backed arguments for and against specific conclusions, modulated by user-specified measures of trust in the provenance of the data. We tested our software with users charged with making decisions in a static humanitarian relief scenario (Salvit et al. 2014) and showed that trained users found the tool helped them manage data and make more informed decisions (than without using the tool). Our current work involves collaboration with emergency services personnel by facilitating and managing data collection in live training exercises during which large-scale emergencies are enacted, requiring (mock) rescue, decontamination and/or evacuation. Data is collected at bronze, silver and gold levels, as well as mock social media data provided by human subjects, and robot sensor data. These data can be accessed by ArgTrust and used by officers to deliberate during a crisis, while maintaining a "data trail" that can justify their decisions later.

Next Steps

Our current work involves: adapting multi-agent market-based mechanisms to handle constrained tasks in dynamic operations; using TAGs to model dynamic environments; applying argHRI to complex tasks and multi-robot settings; and extending ArgTrust to handle live data streams.

Acknowledgments

This work has been supported in part by NSF #IIS-11-16843, NSF #CNS-11-17761, ARL #W911NF-09-2-0053, a US-UK Fulbright Fellowship, and a University of Liverpool Research Fellowship.

References

Azhar, M. Q.; Schneider, E.; Salvit, J.; Wall, H.; and Sklar, E. I. 2013. Evaluation of an argumentation-based dialogue system for human-robot collaboration. In *AAMAS Workshop on Autonomous Robots and Multirobot Systems (ARMS)*.

Bohus, D.; Horvitz, E.; Kanda, T.; Mutlu, B.; and Raux, A., eds. 2011. *AI Magazine, Special Issue on Dialog with Robots*, volume 32(4).

Breazeal, C. 2003. Toward sociable robots. *Robotics and Autonomous Systems* 42:167–175.

Dautenhahn, K. 2007. A Paradigm Shift in Artificial Intelligence: Why Social Intelligence Matters in the Design and Development of Robots with Human-Like Intelligence. *50 Years of AI Journal, LNCS* 4850:288–302.

Dias, M. B.; Zlot, R.; Kalra, N.; and Stentz, A. 2006. Market-based multirobot coordination: A survey and analysis. *Proceedings of the IEEE* 94(7):1257–1270.

Doroodgar, B.; Ficocelli, M.; Mobedi, B.; and Nejat, G. 2010. The search for survivors: Cooperative human-robot interaction in search and rescue environments using semi-autonomous robots. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2858–2863.

Feigenbaum, J., and Shenker, S. 2002. Distributed Algorithmic Mechanism Design: Recent Results and Future Directions. In *Proceedings of the 6th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*, 1–13. ACM Press.

Gerkey, B. P., and Matarić, M. J. 2004. A Formal Analysis and Taxonomy of Task Allocation in Multi-Robot Systems. *Intl Journal of Robotics Research* 23(9).

Goodrich, M. A., and Schultz, A. C. 2007. Human-Robot Interaction: A Survey. *Foundations and Trends in Human-Computer Interaction* 1(3):203–275.

Guy, S. J.; Lin, M. C.; and Manocha, D. 2010. Modeling collision avoidance behavior for virtual humans. In *Proceedings of AAMAS*.

Howard, A.; Matarić, M. J.; and Sukhatme, G. S. 2002. *Mobile Sensor Network Deployment using Potential Fields: A Distributed, Scalable Solution to the Area Coverage Problem*, volume 5. Springer. 299–308.

Kimmel, A.; Dobson, A.; and Bekris, K. 2012. Maintaining team coherence under the velocity obstacle framework. In *Proceedings of AAMAS*.

Koenig, S.; Tovey, C.; Lagoudakis, M.; Kempe, D.; Keskinocak, P.; Kleywegt, A.; Meyerson, A.; and Jain, S. 2006. The Power of Sequential Single-Item Auctions for Agent Coordination. In *Proceedings of AAAI*.

Landén, D.; Heintz, F.; and Doherty, P. 2012. *Complex Task Allocation in Mixed-Initiative Delegation: A UAV Case Study*, volume 7057. Springer. 288–303.

Lemon, O.; Gruenstein, A.; and Peters, S. 2002. Collaborative activities and multi-tasking in dialogue systems. *TAL: Special Issue on Dialogue* 43(2).

Liu, Y., and Nejat, G. 2013. Robotic Urban Search and Rescue: A Survey from the Control Perspective. *Journal of Intelligent & Robotic Systems* 72(2):1–19.

Matarić, M. J.; Sukhatme, G.; and Ostergaard, E. 2003. Multi-robot task allocation in uncertain environments. *Autonomous Robots* 14(2–3):255–263.

McBurney, P., and Parsons, S. 2002. Games that agents play: A formal framework for dialogues between autonomous agents. *Journal of Logic, Language, and Information* 11(3).

Murphy, R. R. 2004. Human-robot interaction in rescue robotics. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews* 34(2):138–153.

Onosato, M.; Takemura, F.; Nonami, K.; Kawabata, K.; Miura, K.; and Nakanishi, H. 2006. Aerial Robots for Quick Information Gathering in USAR. In *Intl Joint Conf of the Society of Instrument Control Engineers and the Institute of Control, Automation and Systems Engineers (SICE-ICASE)*.

Özgelen, A. T., and Sklar, E. I. 2014a. Modeling and Analysis of Task Complexity in Single-Operator Multi-Robot Teams (Late Breaking Report). In *Proceedings of the International Conference on Human-Robot Interaction (HRI)*.

Özgelen, A. T., and Sklar, E. I. 2014b. Toward a Human-Centric Task Complexity Model for Interaction with Multi-Robot Teams. In *AAMAS Workshop on Human-Agent Interaction Design and Models (HAIDM)*.

Parker, L. E. 2002. Distributed Algorithms for Multi-Robot Observation of Multiple Moving Targets. *Autonomous Robots* 12(3):231–255.

Parsons, S.; Rodríguez-Aguilar, J. A.; and Klein, M. 2011. Auctions and bidding: A guide for computer scientists. *ACM Computing Surveys* 43(2).

Prakken, H. 2006. Formal systems for persuasion dialoge. *Knowledge Engineering Review* 21(2).

Prakken, H. 2010. An abstract framework for argumentation with structured arguments. *Argument and Computation* 1.

Rahwan, I., and Simari, G. R., eds. 2009. *Argumentation in Artificial Intelligence*. Springer Verlag.

Salvit, J.; Li, Z.; Perumal, S.; Wall, H.; Mangels, J.; Parsons, S.; and Sklar, E. I. 2014. Employing Argumentation to Support Human Decision Making: A User Study. In *AAMAS Workshop on Argumentation in Multiagent Systems*.

Schneider, E.; Balas, O.; Özgelen, A. T.; Sklar, E. I.; and Parsons, S. 2014a. An Empirical Evaluation of Auction-based Task Allocation in Multi-Robot Teams (Extended Abstract). In *Proceedings of AAMAS*.

Schneider, E.; Balas, O.; Özgelen, A. T.; Sklar, E. I.; and Parsons, S. 2014b. Evaluating auction-based task allocation in multi-robot teams. In *AAMAS Workshop on Autonomous Robots and Multirobot Systems (ARMS)*.

Sklar, E. I.; Azhar, M. Q.; Parsons, S.; and Flyr, T. 2013. A Case for Argumentation to Enable Human-Robot Collaboration (Extended Abstract). In *Proceedings of AAMAS*.

Tambe, M.; Scerri, P.; and Pynadath, D. V. 2002. Adjustable autonomy for the real world. *Journal of Artificial Intelligence Research* 17(1):171–228.

Tang, Y.; Cai, K.; McBurney, P.; Sklar, E. I.; and Parsons, S. 2012. Using argumentation to reason about trust and belief. *Journal of Logic and Computation, Special Issue on Agreement Technologies* 22(5):959–1018.

Walton, D. N., and Krabbe, E. C. W. 1995. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. SUNY Press.

Wellman, M. P., and Wurman, P. R. 1998. Market-aware agents for a multiagent world. *Robotics and Autonomous Systems, Special Issue on Multi-Agent Rationality* 24(34):115–125.

Wong, C. Y.; Seet, G.; and Sim, S. K. 2011. Multiple-Robot Systems for USAR: Key Design Attributes and Deployment Issues. *Intl Journal of Advanced Robotic Systems* 8(1).