

Applying FORR to human/multi-robot teams

Susan L. Epstein
Hunter College and
The Graduate Center
City University of New York
New York, NY USA
susan.epstein@hunter.cuny.edu

Eric Schneider
Hunter College
City University of New York
New York, NY USA
nitsuga@pobox.com

A. Tuna Ozgelen
J. Pablo Munoz
The Graduate Center
City University of New York
New York, NY USA
tunaozgelen@gmail.com
jpablomch@gmail.com

Michael Costantino
College of Staten Island
City University of New York
Staten Island, NY USA
michael.costantino@cix.csi.cuny.edu

Elizabeth I. Sklar
Simon Parsons
Brooklyn College and
The Graduate Center
City University of New York
Brooklyn, NY USA
sklar@sci.brooklyn.cuny.edu
parsons@sci.brooklyn.cuny.edu

ABSTRACT

This extended abstract briefly describes a new methodology for shared decision-making in human/multi-robot teams. We present a new application of *FORR*, a cognitive architecture that considers the opinions of others when choosing actions. While robots in our system ultimately make their own decisions about their actions, they do so based on collective advice from humans, agents and other robots.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*multiagent systems*

General Terms

Theory

Keywords

Multi-robot systems, human-robot interaction

1. INTRODUCTION

This extended abstract summarizes our work on shared decision-making in human/multi-robot teams. Our work is in the area of *mixed initiative* [2, 12] architectures, where a human operator and one or more robots work share the task of making decisions about robots' actions. This mode of control is in contrast to the other two primary architectures for human-robot systems [10]: *directly controlled*,

where decisions about robots' actions are made by human operators; and *fully autonomous*, where robots make their own decisions, without input from a human operator. One well-studied type of mixed-initiative architecture is called *adjustable autonomy*. Here, a human-robot system transfers control dynamically, back and forth, between human and robot [9, 13]. *Collaborative control* offers a dialog-based architecture in which decisions are "discussed" and made in real-time [8]. Another technique uses statistics to infer missing information in human-robot communication [11].

Our approach is founded on the notion that, while the robot ultimately makes its own decision about its actions, it can do so based on *advice* from a human collaborator, other robots or (non-embodied) agents. Just as people frequently make well-considered decisions by asking their friends and relations for opinions on what car to buy or what job to take, a robot can also ask opinions of others. We have implemented a human/multiagent/multi-robot framework where each robot in the system decides on its own actions based on the opinions of a human operator and a number of agents.

This methodology is based on *FORR* (FOR the Right Reasons), a cognitively-plausible architecture that models the development of expertise [3]. *FORR* is predicated on the theory that good decisions in complex domains are best made by a mixture of experts, called *Advisors*. Each Advisor in *FORR* is a resource-bounded procedure that represents a single rationale for decision making. *FORR* provides a common knowledge store, represented as a set of *descriptives*, that Advisors use in different ways. The *FORR* architecture is domain-independent, but the knowledge and procedures it acquires are domain-specific. To date, *FORR* has supported applications for game playing [5], simulated pathfinding [4], constraint solving [6], and spoken dialogue [7].

We have applied *FORR* to our *HRTeam* framework [14, 15]. *HRTeam* uses a modular, multiagent architecture to support shared knowledge amongst agents and to promote shared decision-making. This implementation represents the first application of *FORR* both to a team of agents and to physical robots.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

HRI/HART 2012 Boston, MA USA

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

2. EXPERIMENTAL SETUP

Our experimental testbed models the interior of a building, with a large space including six rooms and a “hallway”, in which the robots explore *interest points*. An experiment specifies a set of n robots and a set of m interest points; each interest point is to be visited by some robot. The interest points can be allocated to individual robots in a number of ways; our current method uses an auction-based mechanism. The set of interest points assigned to a given robot is referred to as that robot’s *mission*. One at a time, a robot selects an interest point from its mission and addresses it. On demand, an A* path planner produces a plan for the robot to reach the selected interest point. This plan is a sequence of subgoals (i.e., locations as (x, y) coordinates) intended to move the robot from its current position to its current interest point. Because a plan considers only a single robot, several robots may interfere along one another’s intended paths. A simple *collision avoidance* module recognizes such a situation and resolves it by pausing some robots until another robot is out of their way. Effectively, unless there is a need to avoid danger, while its mission remains incomplete, a robot selects an unachieved interest point from its mission, constructs and stores a plan to reach that interest point, and repeatedly selects and moves to subgoals along the plan trajectory until it reaches the selected interest point.

The two main components of the physical setup are the *robots* and a *global vision system*. Our multi-robot team is comprised of inexpensive, limited-function platforms. We are currently using the Surveyor SRV-1 Blackfin¹, a small tracked platform equipped with a webcam and 802.11 wireless. (See Figure 2.) Because the Blackfin has very limited on-board processing, our robots rely on off-board processing: each platform is wirelessly tethered to a remote machine which runs its Robot Controller (i.e., as shown in Figure 2).

The *global vision system* is comprised of six Logitech C600 Webcams, which are suspended 10 feet above the arena, to provide an overhead view of the testbed. Each camera’s field-of-view covers approximately one-sixth of the rectangular arena. Each camera is controlled by a Camera Agent that employs OpenCV libraries² [1] to handle image processing. The Camera Agents identify the robots and transmit their positions (2-dimensional location in global (x, y) arena coordinates and orientation θ , in degrees) to the Central Server. They function independently and only broadcast data to the Central Server; they do not receive any messages (other than acknowledgement that the messages they send have been received). Since the robots are identical, the vision system needs some help to identify them uniquely. To distinguish between them, each robot is topped by a “hat”, a white rectangle with a pattern of black dots arranged in a 2-by-3 grid; each grid square either has a dot or does not. Each hat is a character from the Braille alphabet³. Examples appear in Figure 2.

3. APPROACH

Our system architecture (Figure 2) consists of several components: a Human Interface, for operator input; the FORR engine, whose Advisors provide comments; multiple Camera Agents (one per camera) and multiple Robot Controllers

¹http://www.surveyor.com/SRV_info.html

²<http://opencv.willowgarage.com/wiki/>

³<http://www.afb.org/section.asp?SectionID=6>




Figure 1: Surveyor Blackfin and “Braille Hats”

(one per robot); and a Central Server, which handles communication amongst the system components. For each robot in the system, a Robot Controller is instantiated that sends low-level messages to the robot about how to move. The details of the initial design of our framework have been described elsewhere [15]. The new aspect detailed here is the FORR Engine, outlined in bold in Figure 2.

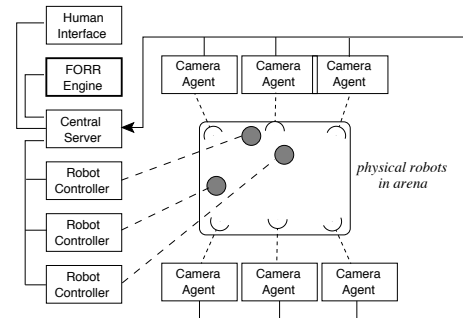


Figure 2: System architecture

As mentioned earlier, a FORR-based system uses a set of *Advisors* to make decisions. In our implementation, each Advisor uses its (relatively simple) rationale to produce *comments* about what the robot should do. A comment is represented as $\langle a, s, ID \rangle$, where ID gives the identity of the Advisor that provided the comment, a is the action that the Advisor is commenting on, and s is the numerical strength (for or against) and intensity of the Advisor’s opinion about the robot performing action a .

Five possible actions are available to the robot. It can $\text{move}(x, y, \theta)$ to location (x, y) with orientation θ . It can $\text{pause}(t)$, i.e., stop moving, until time t and then resume its plan. It can $\text{stop}(t)$ until time t and discard its plan (which is how “stop” differs from “pause”). It can $\text{resume}(t)$, indicating that the robot should resume the current plan at time t (where t can be 0 to mean “now”). Finally, it can do nothing in the current time step, indicated by a $\text{noop}()$ action, and then re-consult the decision-making process.

Two kinds of descriptives are available as input to the Advisors. *Common knowledge* is shared by all the robots and the human, although some agents may have out-of-date copies of it. Examples include a world map and the number of robots in the world. *Self-knowledge* is information the robot has about itself, including its current position and orientation, whether it is currently stopped, its current interest point, any plan it has to reach that point, and the next subgoal in that plan.

FORR organizes its Advisors into a hierarchy of three tiers and considers each tier in turn, in order to make a decision. Advisors in tier 1 are “always correct” and are pre-sequenced. As soon as the first tier-1 Advisor produces a positive com-

ment, its action is forwarded to the Robot Controller and no further Advisors in any tier are consulted. Advisors in tier 2 are reactive planners. If no tier-1 Advisor makes a positive comment and the robot has a current interest point but no plan, each tier-2 Advisor in turn has the opportunity to construct a plan for that point. The first plan output by an Advisor is accepted, and no further Advisors in any tier are consulted. Finally, if neither tier 1 nor tier 2 makes a decision, the Advisors in tier 3 combine their comments in a process called *voting* that chooses an action based both on the Advisors’ past reliability and the comments’ strengths.

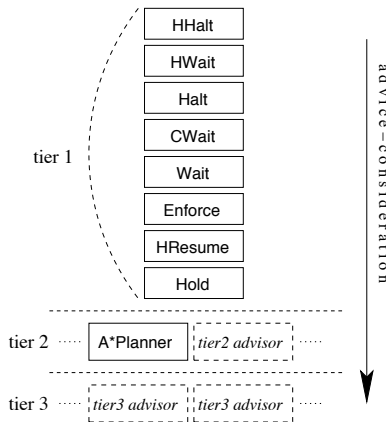


Figure 3: FORR Advisors, organized in 3 tiers

The collision-avoidance mechanism and the human team member are integrated into decision making as Advisors. *HHalt*, *HWait* and *HResume* are driven by the human operator. *CWait* is driven by the collision-avoidance mechanism. The Advisor sequence in tier 1 is shown in Figure 3. *Halt* advocates a `stop()` action for a robot that is not stopped and has a current plan, and discards the plan. *Wait* advocates a `pause()` action for a robot that is not stopped and has a current plan, but leaves the plan in place. *Enforce* checks the progress of a robot with respect to its current subgoal; if it has reached it, *Enforce* advocates a `stop()` and marks the current subgoal as reached and removes it. If the robot has deviated too far from its current subgoal, *Enforce* advocates a `stop()` and removes the current subgoal and possibly the current interest point. If the robot has a plan but no current subgoal, *Enforce* selects a new current subgoal, and advocates a `move()` action to it. *Hold* advocates a `no-op` for a robot that has no current subgoal. There is currently only one tier-2 Advisor, the A* planner, but others are under development, including one to return the robot to recharge and another to circumnavigate an obstruction. Tier-3 Advisors under development represent rationales like “travel toward the current interest point” or “move to a central location”. Control will only reach tier 3 if the robot has no plan (i.e., if tier 2 does not make a decision).

4. SUMMARY

We have briefly described our FORR-based method for shared decision-making in human/multi-robot teams. This architecture contributes to the collection of mixed-initiative approaches to human-robot interaction, by providing a new technique that is different from the standard turn-taking or negotiated approaches. Our method allows a human and

a number of agents to provide advice to a robot—and the robot makes its own decision, by considering the input of its Advisors. Experimentation using this architecture and the physical setup described is on-going.

5. ACKNOWLEDGMENTS

Funding for the work described here has been provided by the National Science Foundation under grants IIS #11-17000 and CNS #08-51901, and by City University of New York Collaborative Incentive Research Grant, #1642.

6. REFERENCES

- [1] G. Bradski and A. Kaehler. *Learning OpenCV: Computer Vision with the OpenCV Library*. O’Reilly, 2008.
- [2] J. R. Carbonell. Mixed-initiative man-computer instructional dialogues. In *BBN Tech Report*, 1971.
- [3] S. L. Epstein. For the right reasons: The forr architecture for learning in a skill domain. *Cognitive Science*, 18(3):479–511, 1994.
- [4] S. L. Epstein. Pragmatic navigation: Reactivity, heuristics, and search. *AIJ*, 100(1–2):275–322, 1998.
- [5] S. L. Epstein. *Learning to Play Expertly: A Tutorial on Hoyle*, pages 153–178. Nova Science, 2001.
- [6] S. L. Epstein, E. C. Freuder, and R. J. Wallace. Learning to support constraint programmers. *Computational Intelligence*, 21(4):337–371, 2005.
- [7] S. L. Epstein, R. J. Passonneau, J. Gordon, and T. Ligorio. The Role of Knowledge and Certainty in Understanding for Dialogue. In *AAAI Fall Symp on Advances in Cognitive Systems*, 2011.
- [8] T. Fong, C. Thorpe, and C. Baur. Multi-Robot Remote Driving With Collaborative Control. *IEEE Trans on Industrial Electronics*, 50(4):699–704, 2003.
- [9] M. A. Goodrich, D. R. Olsen, J. W. Crandall, and T. J. Palmer. Experiments in Adjustable Autonomy. In *IJCAI Wkshp on Autonomy, Delegation, and Control: Interaction with Autonomous Agents*, 2001.
- [10] M. A. Goodrich and A. C. Schultz. Human-robot interaction: a survey. *Foundations & Trends in Human-Computer Interaction*, 1(3), 2007.
- [11] J.-H. Hong, Y.-S. Song, and S.-B. Cho. Mixed-initiative human-robot interaction using hierarchical bayesian networks. *IEEE Trans on Sys, Man & Cyber*, 37(6), 2007.
- [12] E. Horvitz. Principles of mixed-initiative user interfaces. In *Proc of CHI*, 1999.
- [13] P. Scerri, D. V. Pynadath, and M. Tambe. Why the elf acted autonomously: Towards a theory of adjustable autonomy. In *AAMAS*, 2002.
- [14] E. I. Sklar, S. L. Epstein, S. Parsons, A. T. Ozgelen, and J. P. Munoz. A framework in which robots and humans help each other. In *AAAI Spring Symp: Help Me Help You: Bridging the Gaps in Human-Agent Collaboration*, 2011.
- [15] E. I. Sklar, A. T. Ozgelen, J. P. Munoz, J. Gonzalez, M. Manashirov, S. L. Epstein, and S. Parsons. Designing the HRTeam Framework: Lessons Learned from a Rough-and-Ready Human/Multi-Robot Team. In *AAMAS Workshop on Autonomous Robots and Multirobot Systems*, 2011.