Teams for Teams
Performance in Multi-Human/Multi-Robot Teams

We are developing a theory for human control of robot teams based on considering how control varies across different task allocations. Our current work focuses on domains such as foraging in which robots perform largely independent tasks. The present study addresses the interaction between automation and organization of human teams in controlling large robot teams performing an Urban Search and Rescue (USAR) task. We identify three subtasks: perceptual - visual search for victims, assistance-teleoperation to assist robot, and navigation-path planning and coordination. For the studies reported here, navigation was selected for automation because it involves weak dependencies among robots making it more complex and because it was shown in an earlier experiment to be the most difficult. Two possible ways to organize operators were identified as assignment of robots to particular operators or as a shared pool in which operators service robots from the population as needed. The experiment compares two member teams of operators controlling teams of 12 robots each, in the assigned robots conditions or sharing control of 24 robots in the shared pool conditions using either waypoint control or autonomous path planning. Automating path planning improved system performance. Effects of team organization were equivocal.

INTRODUCTION

Unmanned vehicle systems (UVSs), whether in the air or ground, are intrinsically complex systems and rely on remote operator guidance to accomplish different missions. For achieving the goal of military of one operator controlling multiple unmanned vehicles (UVs), the human operator's responsibility must shift from manually controlling vehicles to the supervisory control.

Current UV systems require a many-to-one ratio of operators to UVs. The Predator for example requires a pilot, payload operator, and supervisor to fly. To meet future time-critical missions such as search and rescue or perimeter patrol future systems will need to evolve into many-to-many in which teams of operators control much larger teams of UVs.

Controlling multiple robots substantially increases the complexity of the operator's task because attention must be shared among robots in order to maintain situation awareness (SA) and exert control. In the simplest case an operator controls multiple independent robots interacting with each as needed. A foraging task (Cao et al., 1997) in which each robot searches its own region would be of this category although minimal coordination might be required to avoid overlaps and prevent gaps in coverage especially if robots are in close proximity. Control performance at such tasks can be characterized by the average demand of each robot on human attention (Crandall et al., 2005). Because robots are operated independently an additional robot imposes only an additive demand on cognitive resources. Under these conditions increasing autonomy for individual robots should allow them to be neglected for longer periods of time making it possible for a single operator to control more robots. The frequent switching between robots such a control regimen requires, however, introduces difficulties in maintaining SA.

For dependent tasks the round-robin control strategy used for controlling individual robots would force an operator to plan and predict actions needed for multiple joint activities and be highly susceptible to errors in prediction, synchronization or execution. Even less tightly constrained dependent control of UVs remains substantially more difficult than independent control for teams of operators that must combine their efforts as a team (Mekdeci & Cummings 2009)

If robots are not rigidly assigned to small teams under the control of a single operator, then each event requiring cooperation will either occupy more of an operator's attention than corresponding independent tasks or require the operator to find the controller of another robot and assume the communication and coordination overhead needed to coordinate with him. These interaction times are likely to be highly variable making it difficult to schedule interactions without introducing excessive idle times. Since even moderate variability in neglect time (NT) has been shown (Goodrich et al., 2009) capable of having an operator spending 90% of the time waiting, avoiding such bottlenecks is crucial.

We are developing a general architecture for controlling robot teams based on these observations. We begin by considering how operators self organize to control robot teams and the effects of the different strategies they adopt. Two possible ways to impose organization on operators are through assigning a subset of robots to each or through a Shared Pool (Lewis et al., 2006) in which operators service robots from the full population as needed. Robot assignment has the advantage of reducing the number of robots the operator must monitor and control. The shared pool offers the scheduling advantage of load balancing in that a pool of operators are available as robots need servicing. Efficiencies such as improved SA that might result from controlling a dedicated team at a particular locale, however, must be sacrificed if operator attention is switched among robots following FIFO (first in first out) or similar discipline.

For monitoring, shared pool offers the redundant observer advantage in that a second observer with partially overlapping perceptual judgments may detect things missed by the first. We expect the effects of these advantages to interact with the types of autonomy possessed by the controlled robots. If navigation and path planning were fully autonomous, we

would expect benefits to accrue to shared pool operators due to both scheduling and redundant observer advantages. Autonomous path planning should additionally lessen the effects of loss of SA due to switching between robots because only the victim-marking subtask would be affected. If robots were able to self-reflect and report when they need assistance we might expect to see a stronger scheduling advantage for shared pool. We would additionally expect to see substantial differences between types of autonomy in the numbers of robots that could be adequately controlled.

The present experiment compares performance of robot teams navigating either autonomously or using operator supplied waypoints. The teams were controlled by pairs of operators organized through assigned robots or as a shared pool. In recent experiments (Wang et al., 2009) we have found that participants performing an Urban Search And Rescue (USAR) foraging task using waypoint control were at or over their limits when controlling 12 robots each. Participants who were asked merely to explore showed very similar performance in area covered and reported similar levels of workload on the NASA-TLX. Participants in a perceptual search condition in which the foraging task was performed without the requirement to navigate found twice the victims when monitoring 12 robots and reported substantially lower workload.

The present study uses the same robots and environment but with teams of two operators assigned to control 24 robots. These operators controlled teams of 12 robots in the assigned robots condition. In the shared pool condition operators shared control of the 24 robots. Robots were navigated by operator assigned waypoints as in (Wang et al., 2009) in the manual condition and by an autonomous path planner in the autonomy condition.

Participants were told they were members of a team and would be scored jointly according to the team's performance. Their objective was to explore as much of the map as they could, finding all victims in the areas they covered. The allocation of roles, responsibilities, and tasks was left unspecified.

Our hypothesizes were:

H1: Automating navigation will lead to improved performance

H2: Increasing automation will improve shared pool performance to a greater extent than it improves individual assignment performance.

H3: Shared pool with redundant observers will lead to greater accuracy for marking victims.

METHODS

USARSim and MrCS

The experiment reported in this paper was conducted using the USARSim robotic simulation (Lewis et al., 2007) with 24 simulated Pioneer P2-AT robots performing Urban Search and Rescue (USAR) foraging tasks. USARSim is a high-fidelity simulation of urban search and rescue (USAR) robots and environments developed as a research tool for the study of human-robot interaction (HRI) and multi-robot coordination.

MrCS (Multi-robot Control System), a multi-robot communications and control infrastructure with accompanying user interface, developed for experiments in multirobot control and RoboCup competition (Balakirsky et al., 2007) was used in this experiment. MrCS provides facilities for starting and controlling robots in the simulation, displaying multiple camera and laser output, and supporting inter-robot communication through Machinetta which is a distributed multi-agent coordination infrastructure.

Figure 1 shows the elements of the MrCS. The operator selects the robot to be controlled from the colored thumbnails at the top of the screen. In the large video window, the operator uses pan/tilt sliders to control the camera. The current locations and paths of the robots are shown on the Map Data Viewer (bottom right). Under manual control, robots are tasked by assigning waypoints on a heading-up map on the Map Viewer (bottom right) or through a teleoperation widget (upper right).
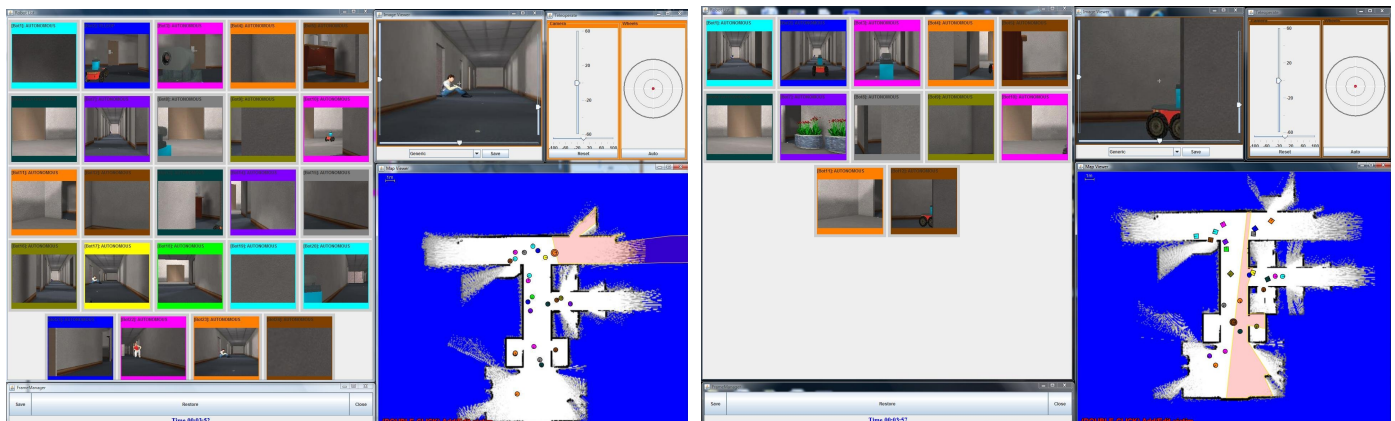

Figure 1. The MrCS user interface with 24 robots for both Shared Pool (left) and Assigned Robots (right) groups.

In the autonomous path planning condition waypoints were assigned through a distributed path planning algorithm although operators could select and control robots through the teleoperation widget as needed. In the assigned robots conditions (Figure 1 right screen) the operator's own robots were available on thumbnails to view or control. Participants could view their teammate's robots on the map, however they could not control them or see their video. In the shared pool condition the participants have equal authority to control every robot and modify the marked victim.

Experimental Conditions

A large USAR environment previously used in the 2006 RoboCup Rescue Virtual Robots competition (Balakirsky et al., 2007) was selected for use in the experiment. The environment was an office like hall with many rooms and full of obstacles like chairs, desks, and bricks. Victims were evenly distributed within the environment. The experiment followed a two by two between groups design crossing autonomy with team organization of shared pool and assigned robot groups. Twenty four robots were controlled by teams of 2 participants. In the assigned robots conditions, participants were each assigned 12 robots to control and could not control robots assigned to their teammate. In the shared pool conditions participants shared control of the 24 UGVs and viewed the same screens. Groups in the manual control conditions issued waypoints to navigate their robots. In the autonomy conditions robots generated their own waypoints using distributed path planning. Participants were able to teleoperate the in-focus robot to extricate it when it became stuck making autonomy a condition with mixed initiative navigation. In the autonomy condition operators performed a supervisory control task involving substantially greater automation than waypoint control in which the robots autonomously navigated toward their goals using autonomous path planning with the operator allowed to override by directing them through new waypoints. A fully autonomous control condition was also run to ascertain the contributions of operator assistance to exploration performance in the autonomy conditions.

Participants

120 paid participants (60 teams) were recruited from the University of Pittsburgh community balanced among conditions for genders. None had prior experience with robot control although most were frequent computer users.

Procedure

After providing demographic data participants read standard instructions on how to control robots via MrCS. In the following 30 minute training session, participants in all conditions practiced control operations. Participants were encouraged to find and mark at least one victim in the training environment under the guidance of the experimenter. After the training session, participants then began the experimental session (25 minute) in which they performed the search task controlling 24 robots in teams. After the task, the participants were asked to complete the NASA-TLX workload survey.

RESULTS

A comparison of regions explored between the experimental autonomy condition and the pure autonomy control in which human assistance was not available found no difference $t(28)=.473$, $p=.64$. This lack of difference ruled out the assistance subtask as a contributor to other effects in the experiment.
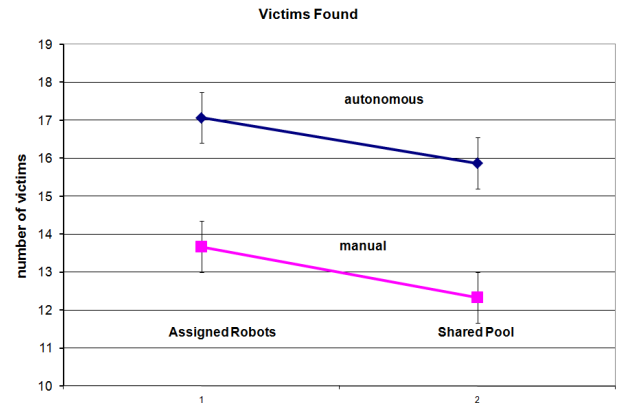

Figure 2. Victims Found

The 2x2 ANOVA for marked victims (Figure 2) comparing team organization and autonomy found a main effect for autonomy, $F_{1,56}=13.436$, $p=.001$. Tests for team organization and the interaction were not found to be significant. T-tests showed differences between autonomy and manual conditions for both assigned robots, $t(28)=2.152$, $p=.04$ and shared pool organizations, $t(28)=3.398$, $p=.002$.
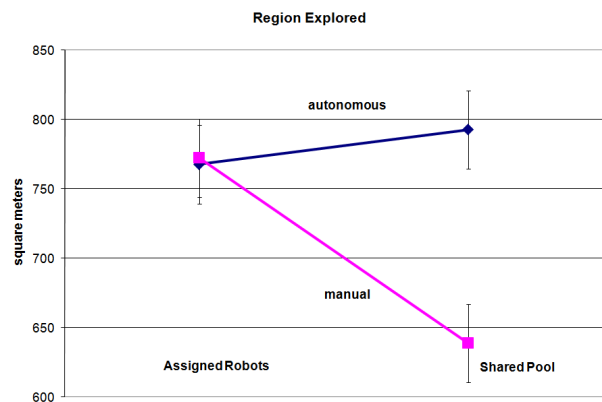

Figure 3. Region explored

The ANOVA for regions explored (Figure 3), by contrast, found main effects for autonomy, $F_{1,56}=6.982$, $p=.011$, and the interaction between autonomy and team organization, $F_{1,56}=7.878$, $p=.007$. The effect for team organization, $F_{1,56}=3.701$, $p=.059$, approached significance as well. For assigned robots there were negligible differences between manual and autonomous conditions. In the shared pool organization, by contrast, there was a substantial advantage for autonomy, $t(28)=4.771$, $p<.001$. This advantage appears due to the decline in explored area from the assigned robots to shared pool conditions for manually controlling participants, $t(28)=2.716$, $p=.015$ rather than performance in the autonomous condition which did not differ across team organization.
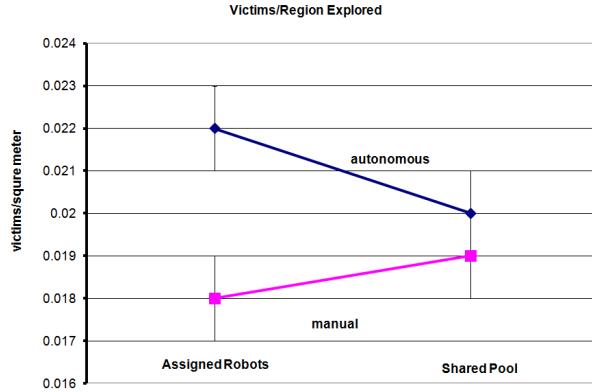
Figure 4. Victims per region explored

While participants enjoying automated path planning found more victims in the assigned robot condition for which areas explored were comparable, their overall advantage in finding victims might have resulted simply from the greater opportunity afforded by exploring larger areas. To examine this possibility we tested the adjusted measure victims/region explored. A main effect was found again for autonomy, $F_{1,56}=7.138$, p=.01, and also the interaction between autonomy and team organization, $F_{1,56}=7.138$, p=.054. No difference was found in victims/region within the shared pool condition, however, for assigned robots participants in the autonomy condition did significantly better, t(28)=3.274, p=.003. There was also a marginal difference favoring assigned robots over shared pool within the autonomy condition, t(28)=1.935, p=.063. The opposite trend for the manual condition apparent in Figure 4 did not reach significance.
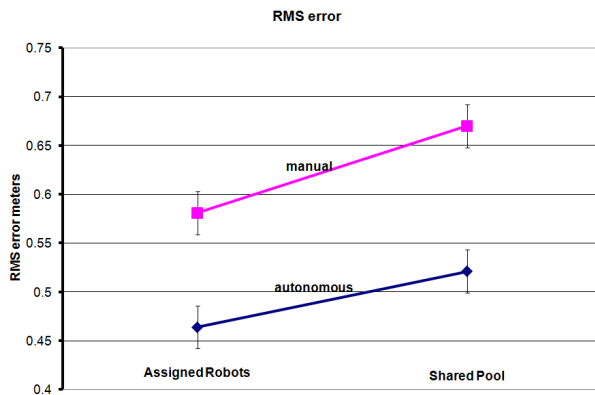

Figure 5. RMS error in marking victims

The related issue of accuracy in marking victims on the laser generated map also favored the assigned robots, $F_{1,56}=18.031$, p<.001 and autonomy conditions, $F_{1,56}=5.434$, p<.023. T-tests show smaller errors in marking for autonomy participants under both assigned robot, t(28)=2.519, p=.018, and shared pool, t(28)=3.549, p=.001 conditions. The advantage for assigned robots, however, only approaches significance, t(28)=1.879, p=.07, in the manual condition.
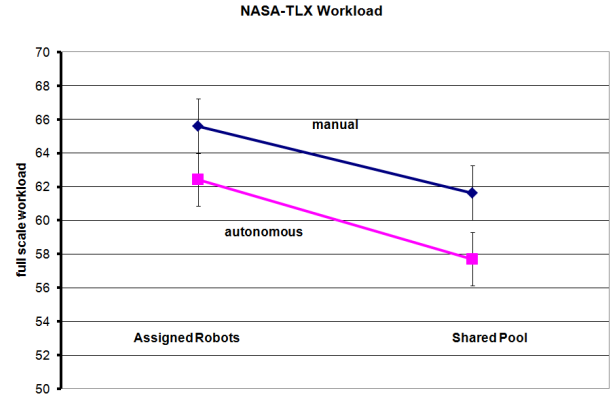

Figure 6. Subjective workload ratings

The full scale NASA-TLX workload measure showed no main effects in an overall ANOVA. When examined separately, however, a slight advantage in workload, t(118)=1.933, p=.056 was observed favoring the shared pool. Unlike earlier studies (Wang et al., 2009) no advantage was found for autonomy.

DISCUSSION

The current experiment with teams of two operators replicates the effects of automated path planning found in an earlier single operator experiment (Wang et al., 2009). In both experiments, relieving operators of the need to perform path planning and manually controlling robots led to finding more victims and marking their locations more accurately. A deterministic roadmap planner using a distributed information gain algorithm was used in the current experiment to drive the robots unlike the previous study in which paths were played back from those generated by earlier participants. Operators appeared to have little difficulty in following these algorithmically generated paths and identified approximately the same numbers of victims as in (Wang et al., 2009) following human generated paths.

The extra exploration appears principally due to autonomous robots being able to move more or less continuously with only brief pauses. In the manual condition, by contrast, an average of 6.19 robots/team were left after being given a single set of waypoints, while an average of 3.13 received no waypoints at all. The assigned robot condition shows advantages here, perhaps because of the diffusion of responsibility in the shared pool where each operator may assume a stopped robot to be the other's responsibility.

Table 1. Neglected Robots in Manual Condition

Number of Robots	Assigned	Shared Pool	Average
Totally	2.00	4.26	3.13
After the Initial Move	4.73	7.66	6.19

Replication of the accuracy advantage for automated path planning was also reassuring because studies such as (Peruch et al., 1995) suggested that this advantage might go in the other direction. Peruch et al. (1995) demonstrated that self-

controlled viewers tended to develop a rich survey knowledge more quickly than passive observers. Because operators in the manual control condition needed to match landmarks between camera views and the laser map, the active exposure to the environment offered by path planning and entering waypoints might have been expected to provide them a more detailed knowledge of the environment and hence greater accuracy in marking victims. Our replicated finding that automated path planning improves accuracy suggests that either the advantage in reduced cognitive load masks poorer survey knowledge in the autonomous condition or that the frequent switching between robots and viewpoints common to the two conditions allows autonomous participants to develop equivalent or superior survey knowledge.

Since avoiding missed targets is crucial to many foraging tasks such as de-mining or search and rescue, thoroughness may be more important than other performance gains such as widening the search area. The analysis of victims per region explored shows that in the assigned robot condition participants using automated path planning found twenty-two percent more victims. This gain is particularly significant because this group was exploring 67% of the map and coming close to matching the actual density of victims of .029/m2. Similar improvements in RMS error and reduction in reported workload suggest that substantial cognitive resources were required for navigation and became available for other subtasks improving overall performance when navigation was automated.

In the assigned robot condition humans were able to plan paths and control 12 robots each to cover the same area as the autonomous path planner. This result supports the feasibility of our information gain algorithm as a substitute for human operators (performance is no worse). We believe the divergence in area covered between manual and autonomous control in the shared pool condition was due to diffusion of responsibility. Participants in the call-center condition were confronted with a bank of videos (Figure 1) much like a security guard monitoring too many surveillance cameras. Under these conditions coordination demands were ambiguous. Participants were frequently observed to reach some form of agreement for dividing robots to be controlled at the outset (like a self-organized individual condition). Later as they performed the task some robots with apparently ambiguous assignments were never moved out from the center of the building. This diffusion of responsibility explanation is supported by reported workload which was marginally lower in the shared pool condition. Informal observation of participants suggested that in the shared pool condition, the 23 thumbnails being monitored appeared to interfere with their ability to focus on particular robots and to locate victims. When multiple robots found similar looking victims participants often became confused marking the same victim twice or leaving a similar appearing victim unmarked on the map. There also appeared to be difficulties in switching strategies where one operator might be trying to control a subset of robots while the other was organizing the task by area.

REFERENCES

Balakirsky, S., Carpin, S., Kleiner, A., Lewis, M., Visser, A., Wang, J., & Zipara, V. (2007) Toward hetereogeneous robot teams for disaster mitigation: Results and performance metrics from RoboCup Rescue. *Journal of Field Robotics.*

Cao, Y., Fukunaga, A., & Kahng, A. (1997) Cooperative mobile robotics: Antecedents and directions. *Autonomous Robots*, 4, 1-23.

Crandall, J., Goodrich, M., Olsen, D., & Nielsen, C. (2005) Validating human-robot interaction schemes in multitasking environments. *IEEE Transactions on Systems, Man, and Cybernetics*, Part A, 35(4):438–449.

Goodrich, M., Morse, B., Engh, C., Cooper, J. & Adams J. (2009) Towards Using UAVs in Wilderness Search and Rescue: Lessons from Field Trials. *Interaction Studies, Special Issue on Robots in the Wild: Exploring Human-Robot Interaction in Naturalistic Environments*, 10(3):453-478.

Lewis, M., Polvichai, J., Sycara, K., & Scerri, P.(2006) Scaling-up Human Control for Large UAV Teams, In N. Cooke (Ed.). *The Human Factors of Remotely Piloted Vehicles, New York: Elsevier*, pp. 237-250.

Lewis, M., Wang, J., and Hughes, S. (2007). USARSim : Simulation for the Study of Human-Robot Interaction, *Journal of Cognitive Engineering and Decision Making*, (1)1, 98-120.

Mekdeci, B., & Cummings, M. L. (2009). Modeling Multiple Human Operators in the Supervisory Control of Heterogeneous Unmanned Vehicles. *9th Conference on Performance Metrics for Intelligent Systems(PerMIS'09), Gaithersburg, MD,* Sep. 21-23.

Peruch, P., Vercher, J., & Guthier, G. (1995). Acquisition of Spatial Knowledge through Visual Exploration of Simulated Environments. *Ecological Psychology,* 7(1): 1-20.

Wang, H., Lewis, M., Velagapudi, P., Scerri, P., and Sycara, K. (2009) How search and its subtasks scale in N robots. *Human Robot Interaction Conference, ACM.*