EXPERIMENTS IN HUMAN-ROBOT TEAMS

Curtis W. Nielsen
Brigham Young University
Provo UT 84604
curtisn@cs.byu.edu

Michael A. Goodrich
Brigham Young University
Provo UT 84604
mike@cs.byu.edu

Jacob W. Crandall
Brigham Young University
Provo UT 84604
crandall@cs.byu.edu

Abstract In this paper, we study human-robot interaction with the goal of learning how teams of semi-independent, semi-autonomous robots are best served by human input in a map-building situation. In our experiments we use multiple combinations of three interaction schemes to control three robots as they build a topological map of an indoor environment. The results from our experiments show the tradeoffs of the various interaction schemes in terms of workload and performance.

Keywords: Adjustable interaction, Human-robot teams, topological map-building

1. Introduction

Modern research has given mobile robots the ability to solve a variety of interesting problems. By efficiently integrating human creativity, intuition, and high-level reasoning into a human-robot system, robots will be able to solve much more difficult problems. In this context, an interesting question to answer is how can human abilities best be used in a team of semi-autonomous, semi-independent robots. One purpose of our research is to develop human-robot team organization concepts that enhance a multi-agent team over all-robot-teams. The specific objective of this paper is to study how teams of semi-independent, semi-autonomous robots are best served by human input in a map-building situation. Research in this area will help identify and understand parameters that constitute effective human interaction with teams of semi-autonomous, semi-independent robots.

1.1 Previous Work

Arkin and Ali's work has been useful to our research (Ali00). In their work, they present experimental results for hundreds of test subjects of a shared-control system that allows a human to interact with a team of simple behavior-based robots. In measuring the effectiveness of human-machine interaction, much work has been done on operator workload. Of particular relevance is Boer's work relating workload and entropy (Boer99). In addition, Boer has used secondary tasks to help evaluate the cognitive workload placed on human operators. We have studied the effects of human neglect on robot performance for different types of interaction, and used these studies to design robot autonomy levels and human interfaces that facilitate high robot performance with minimal human input. We call this the study of neglect tolerance (Crandall02).

In order to perform the experiments we are interested in, we use a behavior-based topological map-building algorithm in (Nielsen02). In the behavior-based map-buliding algorithms, a robot recognizes a set of afforded behaviors, which are loosely based on navigational primitives such as "turn right" and "turn left", and uses these behaviors to define landmarks in the environment. This approach extends Sebastian Thrun's work (Thrun98) by adjusting his probabilistic map-building and localization algorithms to perform in topological environments.

2. The Experiment Framework

2.1 The Interface

We have designed an interface that permits a human to observe the movement of both individual robots and teams of robots, and then direct the control of each robot from a workstation. The interface has a view of the sonar measurements around the robot, a compass, a video feed, a list of robots currently in the system, and a mapping area. The mapping area contains information about dead-reckoning and topological maps and affords the human the ability to disambiguate topological landmarks via drag and drop algorithms. The interface we use is shown in Figure 1.

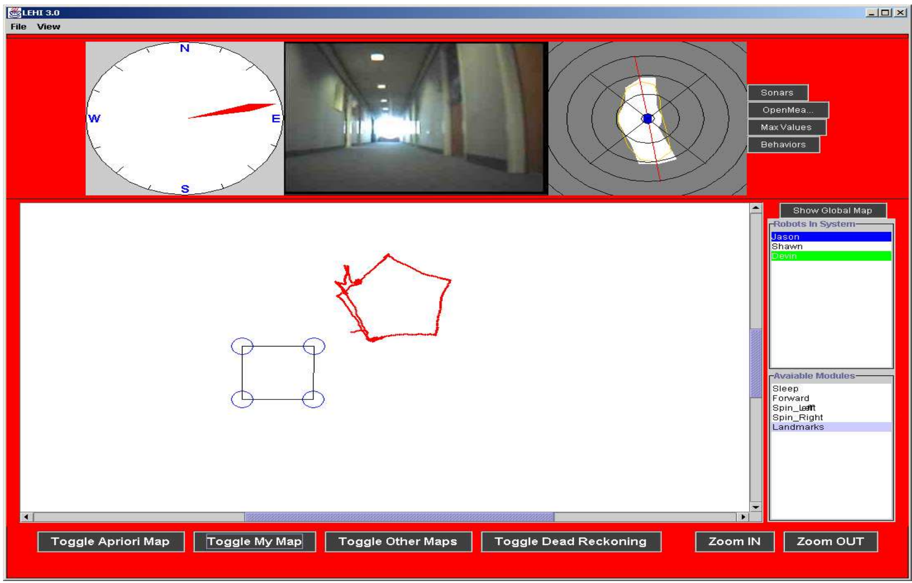

Figure 1. The interface with which we perform our experiments.

Via the interface, humans can interact with the robots by voice, mouse, keyboard, or joystick. In addition, we have implemented an adjustable interaction control system that affords humans and robots different levels of interaction. The interface acts as a centralized agent where the robots report their findings to the interface, which then passes the information to the other robots in the system. In the experiment herein, human input devices are limited to mouse and joystick.

2.2 Interaction Schemes

We are interested in experiments that study the behavior-based map-building problem in a situation where a human interacts with robots via three different interaction schemes. The human has two areas of input to the robot: 1) control of the movement of the robot, and 2) disambiguation of landmarks. The three interaction schemes we will use are now described.

2.2.1 Teleoperate and Landmark (TOL). The robot is controlled via tele-operation through the environment. When the robot reaches a landmark the user must tell the robot that it has reached a landmark by clicking a button on the joystick, similar to Sebastian Thrun's work in (Thrun98). When the user clicks a button the robot recognizes the set of afforded behaviors at the place of interest and creates a landmark that corresponds to the afforded behaviors. (An afforded behavior is a direction of possible travel.) When similar appearing places are found in the environment, it is the responsibility of the human to determine whether or not the places are the same or distinct (i.e., landmark disambiguation). If the landmarks represent the same place in the environment, the user drags the landmarks together via the graphical interface.

2.2.2 Point to Point and Human Snapper (PTP). The robot is controlled via point-to-point commands (e.g. "left at next intersection", "right at next intersection"). When the robot completes a command it resets the control to "Go forward" and the robot will proceed forward until it cannot continue, at which point it will wait for another directive from the user. As the robot moves throughout the environment it autonomously identifies afforded behaviors and uses the set of afforded behaviors to classify the landmarks. Similar to the TOL interaction scheme, in the PTP interaction scheme, the user is responsible for distinguishing between similar places in the environment.

2.2.3 Region of Interest and Sealing (ROI) . The human directs the robot to regions of interest via the graphical interface. When the robot is near the region of interest it will perform its own exploration, landmark detection, and landmark disambiguation algorithms described in (Nielsen02). In essence, the algorithms use a wall-following heuristic to estimate where the robot should move in order to learn or confirm map information about the environment. As the robot explores, it builds a map that contains sealed areas; i.e., areas that have been explored and cannot be "re-discovered" by accident. The robot uses the known part of its map for localization and can incrementally add new discoveries about the environment to the map. With the ability to localize itself on its own map, the robot can now accurately perform landmark disambiguation without human intervention.

We will conduct experiments using three robots with various combinations of these three interaction schemes.

2.3 Measuring Performance and Workload

In order to determine the effectiveness of an interaction scheme, we use two metrics: 1) the performance of the human-robot system and 2) the workload on the human.

2.3.1 Performance. The performance of the human-robot system is measured as the time it takes for the system to complete an accurate map of the environment. Other metrics could be used such as robot idle time, time of human attention to each individual robot, or re-traversal of places in the environment. However, we determined that the measurement of time to completion encompasses a number of other performance criteria, which implies that this criterion is useful for calculating performance.

2.3.2 Workload. In order to measure the workload on the user, we use four metrics: 1) entropy of the joystick, 2) velocity of the mouse, 3) number of button clicks on the mouse and joystick and 4) the number of times the user switches between robots.

The workload from joystick entropy is found by using a strategy developed by Boer et al., which calculates steering entropy for evaluating driver workload (Boer99).

The workload from mouse movement is found by calculating the average velocity of the mouse movement during an experiment. Mouse movement is calculated in pixels per second and normalized by the size of the interface to a value between 0 and 1. Instantaneous mouse velocity is weighted according to .9*previous velocity + .1*current velocity because rapid mouse movements indicate periods of high workload instead of a single instant of high workload.

The partial workload from the total clicks of the mouse and joystick and switching between robots, is added to the total workload as a square pulse measuring 10% of the maximum workload over a 10 second interval. The reason for the length of the square pulse is due to the effects of a context switch when the human changes control of robots and the dexterity involved with disambiguating landmarks and snapping similar landmarks to each other.

The average human workload for an experiment is the sum of the joystick entropy, the average mouse velocity and the average total clicks over the time for completion of the experiment. A possible area for future research is to find the best way to combine the various activities of human input to balance the workload throughout an experiment in comparison to the current model of interleaving moments of extreme workload with moments of minimal workload.

2.4 Experiments

In our experiment, we controlled the robots in the environment shown in Figure 2 with the robots starting in the middle of the map facing east. Note that the human knows the starting positions of the robot but the

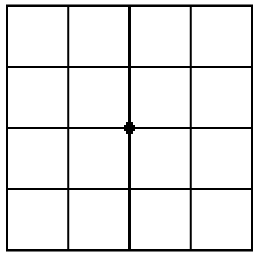

Figure 2. A topological representation of the environement we use for our experiments.

robots do not; the robots do, however, know that each robot started in the same landmark facing the same direction. This allows the robots to understand each others maps.

The goal of the robot system is to have the robots build a topological map in the shortest amount of time with minimal human workload. Table 1 shows the experiments we are interested in. Note, the TOL interaction scheme is never implemented more than once per control scheme because the workload on the human would be too high.

Control Schemes	*Interaction Scheme*	*Interaction Scheme*
All the same	PTP, PTP, PTP	ROI, ROI, ROI
Two the same one Different	PTP, PTP, TOL PTP, PTP ROI	ROI, ROI, TOL ROI, ROI, PTP
All Different	ROI, PTP, TOL	

Table 1. The seven control schemes we used in our experiments.

3. Results

3.1 Instantaneous Workload

The instantaneous workload on the human is useful for showing how the workload changes with different interaction schemes. Figures 3 to 7 show the workload for various control schemes.

It can be seen from these figures that as the number of PTP interaction schemes increase, the workload increases because of the amount of mouse movement and mouse clicks that become necessary. Additionally, when

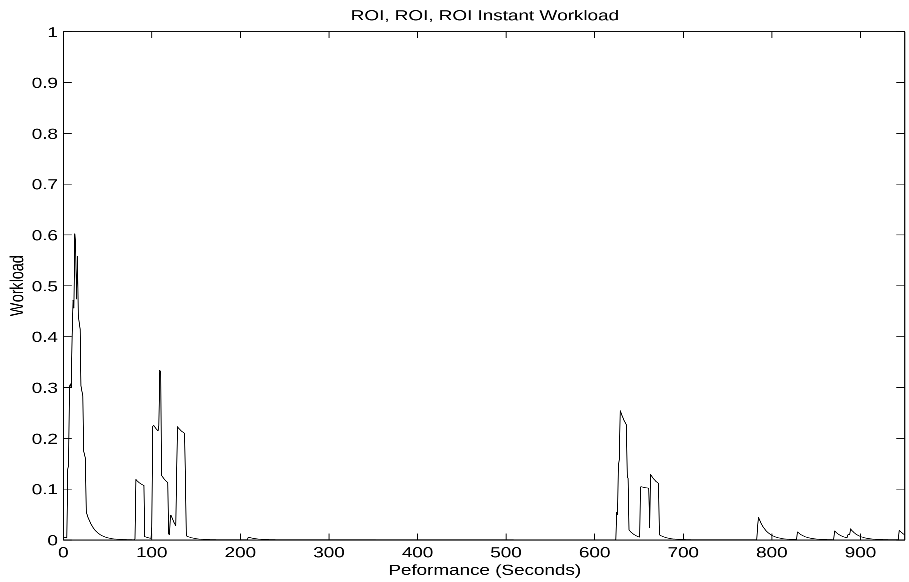

Figure 3. Typical instantaneous workload when using three ROI interaction schemes for control of the three robots.

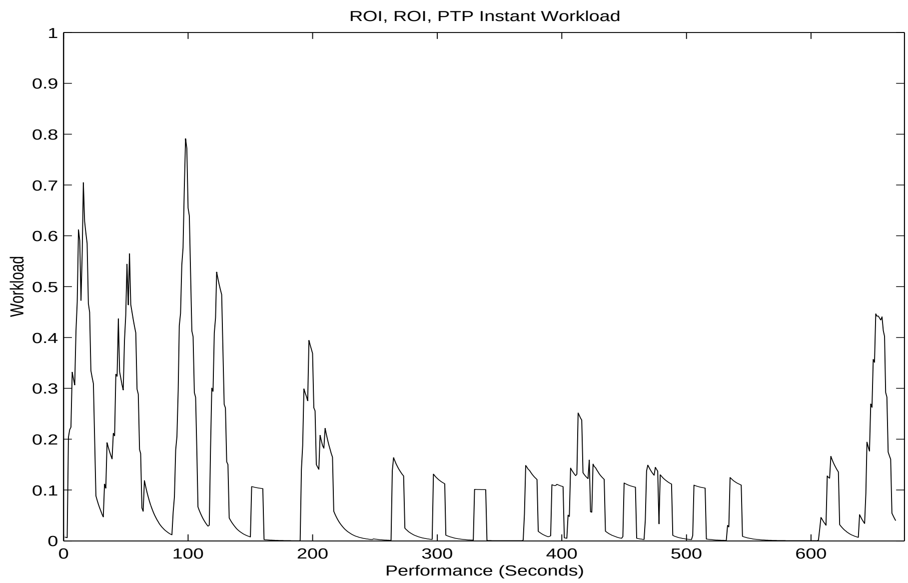

Figure 4. Typical instantaneous workload when using two ROI and one PTP interaction schemes. As we increase the number of PTP interaction schemes, we increase workload and performance.

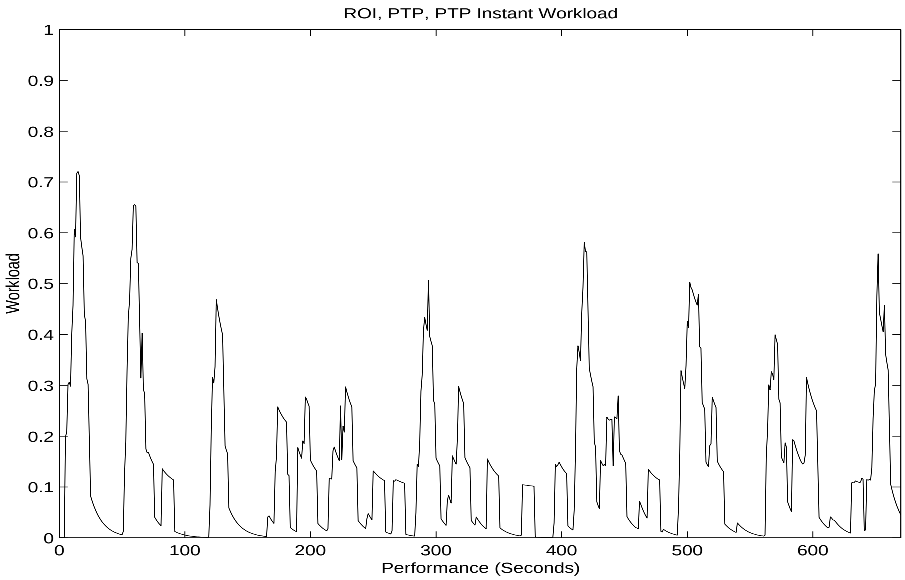

Figure 5. Typical instantaneous workload when using two PTP and one ROI interaction schemes.

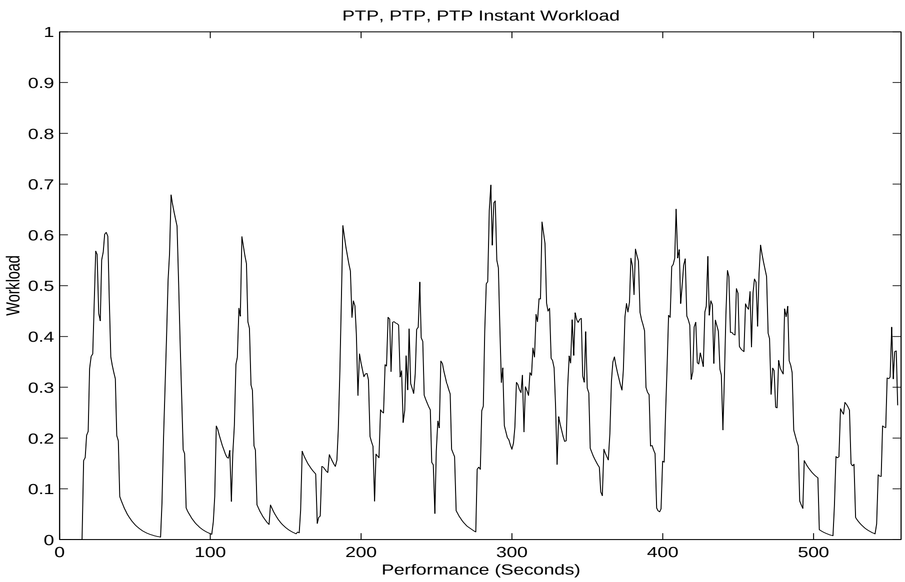

Figure 6. Typical instantaneous workload when using three PTP interaction schemes. Note the continued increase in workload and performance.

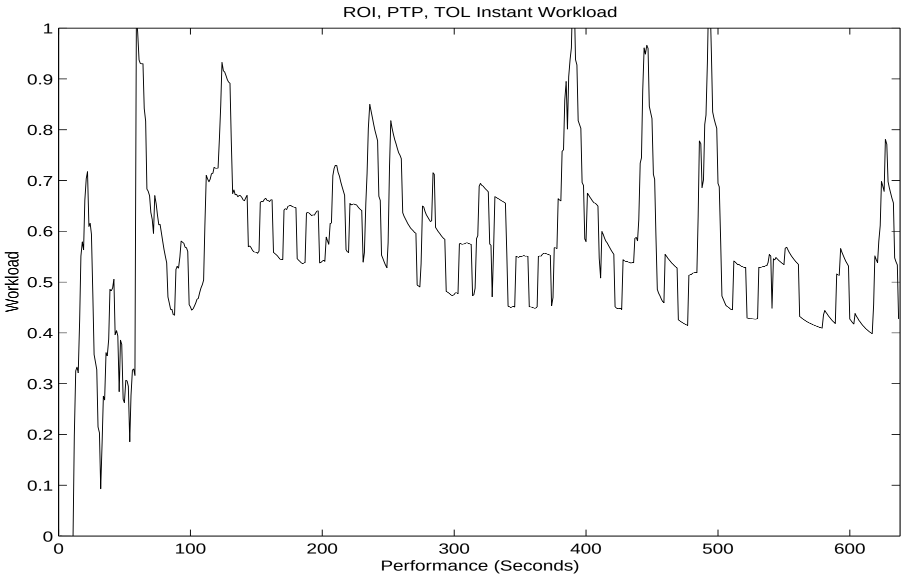

Figure 7. Typical instantaneous workload when using ROI, PTP, and TOL interaction schemes. When a TOL interaction scheme is used, there is a significant increase in workload and a slight increase in performance.

the TOL interaction scheme is used, the workload takes a significant jump because of the constant need to tend to the teleoperated robot.

The data gathered from the experiments presents clear evidence in regards to the performance and workload of the various control schemes. Figure 8 shows the average results of each control scheme with the ellipses representing one standard deviation in the performance and workload. The data used to construct Figure 8 is shown in Table 2.

Control Schemes	*Average Performance*	*Average Workload*
PTP PTP TOL	525	.6297
PTP PTP PTP	569	.2289
ROI PTP TOL	610	.6171
PTP PTP ROI	675	.2000
ROI ROI PTP	773	.0898
ROI ROI TOL	775	.5906
ROI ROI ROI	984	.0641

Table 2. The average performance, measured in seconds until completion, and average human workload for each of the control schemes.

With these experiments, the tradeoffs between the various control schemes becomes clear: As the human gives the robots more autonomy, the performance with respect to time of completion decreases and the human workload decreases. As the human takes more control of the robots, the performance is increased as the workload increases. The control scheme with the highest workload is the (PTP, PTP, TOL) control scheme, and correspondingly, this same control scheme has the best performance in relation to time of completion. Likewise, the (ROI, ROI, ROI) control scheme has the lowest workload and the worst performance. An interesting observation in the distributions of the performance and

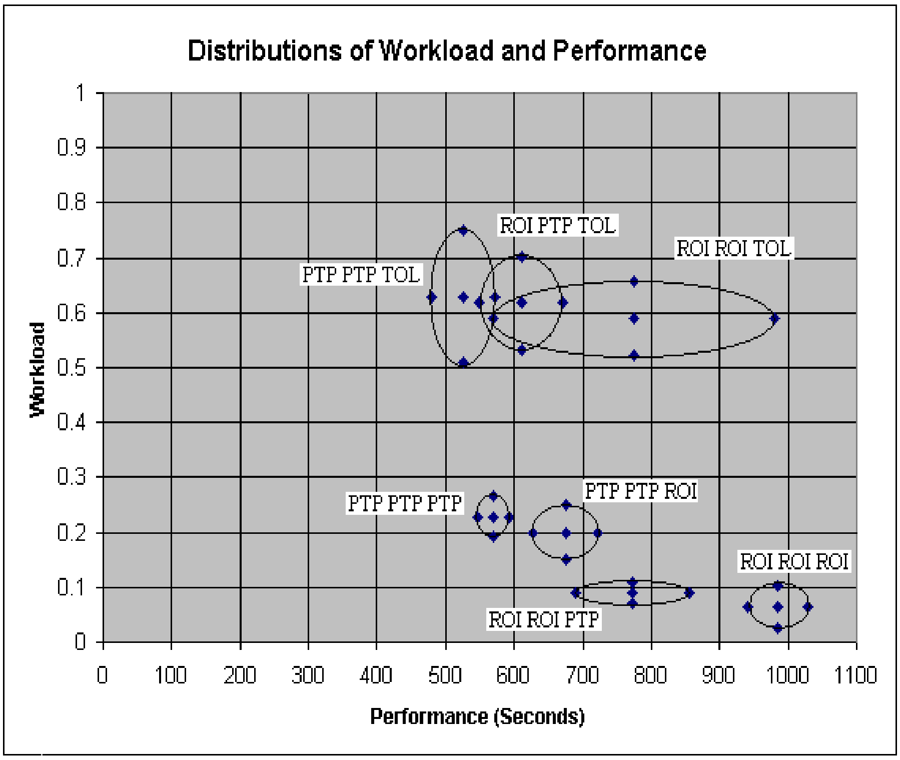

Figure 8. Distributions of the various control schemes. The ellipses represent the standard deviation in the performance and human workload over four subjects and five experiments of each control scheme.

workload is that in all the control schemes, as more PTP interaction schemes are used, the time for completion decreases fairly consistently with a similarly consistent increase in human workload. Furthermore, as more ROI interaction schemes are used, the performance decreases fairly consistently along with the workload. Additionally, when the TOL interaction scheme is used, the workload on the human increases drastically, with a slight gain in the performance of the system. It is remarkable that changes in the three interaction schemes are consistent in how they affect the performance of the system and the workload on the human.

A notable result from Figure 8 is that the variance in the time as well as the variance in workload is largest in the joystick experiments in comparison to the other experiments. Occasionally, a user less experienced with the joystick and the interface would become overwhelmed and perform poorly on the experiments. However, it is of interest that the weak performances, when they happened, only occurred when one of the interaction schemes was TOL. This is a valuable result because it shows that the PTP and ROI interaction schemes can be used with minimal instruction and an untrained human can perform comparable to a human trained in the various interaction schemes.

In essence, we have learned the following about each of the interaction schemes presented:

1 As the number of PTP interaction schemes is increased, we see an increase in performance and human workload.

2 As the number of ROI interaction schemes is increased, we see a decrease in performance and human workload.

3 When a TOL interaction scheme is used, we see a dramatic increase in workload accompanied by a slight increase in performance.

4. Conclusions

Behavior-based mapping is a topological map-building algorithm that facilitates sharing information about an environment between robots and humans. We have designed a task that uses one human and three robots to build a behavior-based map of an environment. The human interacts with the robots via three interaction schemes: 1) teleoperate and define landmarks, 2) point-to-point and human snapper, and 3) Region of interest and sealing. We have studied the performance and human workload of the human-robot system using various combinations of the three interaction schemes. The results show clear tradeoffs as the level of interaction is adjusted. When the human assumes more responsibility, the human workload increases, but performance increases as well. When the human relinquishes control to the robots, human workload and performance both decrease. Note that the performance increase happens up to a certain point dependent on the skills of the human. If the workload saturates (e.g., with TOL), team performance is sensitive to many factors and can actually decrease. We also show the instantaneous workload for a variety of interaction schemes to show the actual increase in human workload as the user assumes more responsibility.

An intersting direction for future work is to adjust the number of robots interacting with a human. As we change the number of robots, we could use the new information to improve our workload vs. performance figure. In essence, we want to learn the optimal number of robots a human can control via different control schemes and then determine the optimal ratio of interaction schemes between the robots and the human with respect to performance and human workload. This is similar to Dudenhoeffer's research (Dudenhoeffer 01) but includes human workload. This information would be valuable in determining the tradeoffs in various compositions of human-robot interaction schemes.

Additionally, we are interested in performing the experiments presented in this paper in environments of varying complexity and with more subjects in order to gain a more accurate understanding of the relationship between workload and performance of human-robot teams.

References

Ali, K. S., Arkin, R. C., (2000). "Multiagent Teleautonomous Behavioral Control," Machine Intelligence and Robotic Control.

Boer E. W., Et. Al., (1999). "Development of a steering entropy method for evaulating driver workload," International Congress and Exposition.

Crandall, J. W., Goodrich, M. A., (2002). "Characterizing Efficiency of Human Robot Interaction: A Case Study of Shared-Control Teleoperation," IEEE/RSJ International Conference on Intelligent Robots and Systems.

Dudenhoeffer, D. D., Bruemmer, D. J., Davis, M. L. (2001). "Modeling and Simulation for Exploring Human-Robot Team Interaction Requirements," Proceedings of the 2001 Winter Simulation Conference.

Nielsen, C. W. (2002). "Behavior-Based Landmarks for Topological Mapping and Navigation of Indoor Environments," Masters Thesis. Brigham Young University.

Thrun, S., Burgard, W., Fox, D. (1998). "A Probabilistic Approach to Concurrent Mapping and Localization for Mobile Robots," Machine Learning and Autonomous Robots.