Human vs. Algorithmic Path Planning for Search and Rescue by Robot Teams

Shih-Yi Chien, Huadong Wang, Michael Lewis
School of Information Sciences
University of Pittsburgh
Pittsburgh, PA 15260 U.S.A.
shc56@pitt.edu, huw16@pitt.edu, ml@sis.pitt.edu

Substantial automation will be needed to allow operators to control the large teams of robots envisioned for search and rescue, perimeter patrol, and a wide variety of military tasks. Both analysis and research point to navigation and path planning as prime candidates for automation. When operators are isolated from robot navigation, however, there may be loss of situation awareness (SA) and difficulties in monitoring robots for failures or abnormal behavior. Operator's navigational strategies are quite complex and extremely changeable at foraging tasks in unknown environment reflecting background knowledge and expectations about human and natural environments. These considerations are missing from automated path planning algorithms leading to differences in search patterns and exploration biases between human and automatically generated paths. Effectively integrating automated path planning into multirobot systems would require demonstrating that: 1-automated path planning performs as well as humans on measures such as area coverage and 2- use of automated path planning does not degrade performance of related human tasks such as finding and marking victims. In this paper we seek to compare the divergence between human manual control and autonomous path planning at an urban search and rescue (USAR) task using fractal analysis to characterize the paths generated by the two methods. Area coverage and human contributions to mixed-initiative planning are compared with fully automated path planning. Finally, the impact of automated planning on related victim identification and marking tasks is compared for automated paths and paths generated by previous participants.

INTRODUCTION

Controlling multiple robots substantially increases the complexity of the operator's task because attention must constantly be shifted 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 search task 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 these tasks can be characterized by the average demand of each robot on human attention (Crandall et al., 2005). Because robots are controlled independently team size can be increased simply by adding additional operators, another advantage of this control regimen.

Our analysis focuses on human interaction with robots performing foraging tasks (Cao, Fukunaga, & Kahng, 1997) of this type, in which each robot searches separate regions. In earlier studies (Wang et al., 2009a) we have shown that human operators are at or over their cognitive limits when asked to control 12 robots. Performance in finding and marking victims (indirectly affected tasks), however, was substantially improved when path planning was automated. The results tended to support automation of path planning and navigation over efforts to improve automation for target recognition and offered hope that humans might be able to control even larger teams.

There are a variety of standard measures of operator effectiveness for search and rescue tasks (Steinfeld et al., 2006), for example, the number of rescued victims or the area explored. Such measures provide an appraisal of "how well" the operator has done (the performance), but fail to capture "how" the operator has achieved this performance. For example, the longest path length may not correspond to the maximum victims found, nor does the shortest path length necessarily lead to the fewest victims found. Such task independent metrics do not capture the many aspects and multitude of ways activities may be carried out — especially in adaptation to surprising or unplanned for events. (Voshell, Woods, & Phillips, 2005). Understanding "how" operators perform their tasks requires new measures sensitive to the processes they employ. Measures of process have been found to illuminate stark differences in strategy (Lewis, 1990) in a variety of domains even when performance measures show equivalence. Consideration of process measures is therefore crucial in comparisons where strategic differences are expected.

In this paper we investigate operators controlling robot teams performing a search and rescue task using either human or algorithmically generated paths. Human navigational behavior could support useful information such as orienting toward a potential victim or searching behind a desk. Such context guided navigation would not be available to a robotic path planner working solely from an occupancy grid based on laser scans. Individual variation may provide valuable information about what aspects of the environment are leading to specific behavioral selections (Gotceitas & Colgan, 1988; Neumeister et al., 2004). The paths followed by robots in a team are critical to search results especially in time-pressured domains such as search and rescue tasks, but characterizing them is complex and difficult. For instance, a straight path may cover a longer distance but fail to support situation awareness (SA) as well as a curved path that moves in an orbit around a potential victim. Full performance cannot be assessed by simply comparing outcomes, but should combine several criteria to evaluate tradeoffs, such as that between auto path planning that might increase explored area and victims found while increasing victims missed.

In this paper we compare the control behavior of human operators controlling robot teams through setting waypoints with that of a distributed path planning algorithm. Because operator involvement is required for identifying and attending to victims it is important to characterize differences in the way humans and automation went about searching and how those differences might affect operators' ability to identify and locate victims. We examine a number of measures distinguishing control behaviors of the two types and examine their effects on the full search and rescue task.

Fractal Dimension (FD) was defined by Mandelbrot in 1967 as a continuous analogue of discrete geometric dimensions (Mandelbrot, 1967; Milne, 1997). The fractal dimension provides an affordance-centered metric that captures the efficiency of exploration of space and provides a rich indication of how well a human-robot team is performing (Voshell, Woods, & Phillips 2005). FD measures movement paths at multiple spatial scales to compute the fractal value. A tortuous line, for example, has a FD value $1 \le FD \le 2$ while a completely straight line FD would equal 1 (figure 1). FD describes the variation in a surface or line when examined at varying spatial scales (Craighead, 2009). It has previously been used for estimating animal behavior, such as orientation (Benhamou, 1989) or movement paths (Nams, 2006). More recently it has been applied to characterize robot operators' path planning and search strategies, for instance, Craighead (2009) linked FD to operating skill level over training sessions while Voshell, Woods & Phillips (2005) used FD to investigate the influence of perspective folding displays on navigation. In this report we use FD as a measure to compare robot trajectories determined through autonomous path planning with those resulting from human entered waypoints. The effects of FD on two dependent measures, victims found and teleoperation duration/times were examined without reference to source of the paths. We hypothesize that since human generated paths are straighter and more likely to take into account landmarks and other aspects of the visual environment that may facilitate locating the robot and victim, they are not as likely to require teleoperating a robot to identify it on the map or maneuvering a robot to bring landmarks into view to provide context for marking victims.

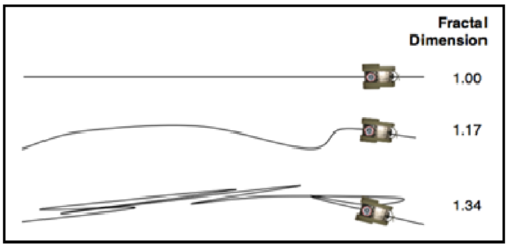

Figure 1. Example of Path Tortuosity (Craighead, 2009)

METHODS

Data Sets

Data for the comparison of human and algorithmically generated paths come from a series of experiments (Wang et al., 2009a, Wang et al., 2009b) as well as data from a subsequent condition extending Wang et al. (2009b). The initial experiment, Wang et al. (2009a), used a repeated measures, between groups design to compare operators controlling 4, 8, and 12 robots in a full task condition in which they both navigated and searched for victims or part task conditions in which they either navigated only or searched and marked victims with robots following paths determined by earlier participants in the navigation only condition. Results from this final condition provide data for search performance using human (but not self) generated paths. The second data set (Wang et al., 2009b) comes from an experiment comparing two person human teams controlling 24 robots in either an assigned robot condition in which each controlled 12 or a shared pool condition in which control of the 24 robots was shared between the two operators. Teams either controlled robots manually by assigning waypoints or performed only the search and victim marking task with robot paths determined by a distributed path planning algorithm.

Table 1 Experiment design for data sets

Wang et al. 2009a 45 participants	4 Robots	8 Robots	12 Robots
Full task			
Exploration			
Perceptual search			✓

Wang et al. 2009b 120 participants	Assigned Robots	Shared Pool
Manual	✓	
Automated	✓	

An added control for the last experiment examined navigation performance for an automation only condition in which no operator was available to assist by freeing stuck robots.

USARSim and MrCS

USARSim (Lewis, Wang & Hughes, 2007) 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 multirobot communications and control infrastructure with accompanying user interface developed for experiments in multirobot control and RoboCup competition (Balakirsky et al., 2007) was used in these experiments. MrCS provides facilities for starting and controlling robots in the simulation, displaying camera and laser output, and supporting inter-robot communication through Machinetta a distributed multi-agent system.

Figure 2 shows the MrCS user interface for 12 robots. Thumbnails of robot camera feeds are shown on the left, a video feed of interest in the top, middle. A GUI element in the top right allows teleoperation and camera pan and tilt. The bottom right shows the current map and allows operators to mark victims.

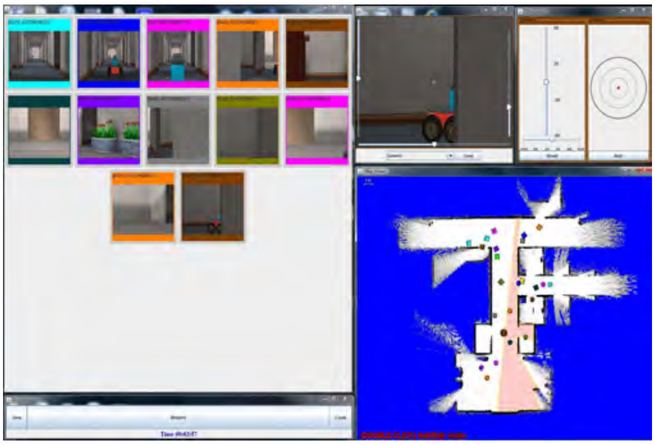

Figure 2. The MrCS user interface for Assigned 12 robots.

Path Planning

The autonomous path planning used in Wang et al. (2009b) was based on a deterministic roadmap planner (Latombe, 1991) and developed using the Carnegie Mellon Robot Navigation Toolkit (CARMEN) (Montemerlo, 2003). It used current team information about the environment to generate further paths for robots. The proposed paths might be accepted or rejected based on current and expected location information. For example, to avoid duplication a robot might reject a proposed path that brought it into the vicinity of another robot.

Human (not self) generated paths were played back from log files recorded by participants in Wang et al.'s (2009a) exploration condition. These participants were instructed to explore as much of the environment as possible but did not search for victims or mark them. In the manual condition (assigned robot teams in Wang et al., 2009b), the operator used MRCS to issue multiple way point paths to the robots (Figure 3).

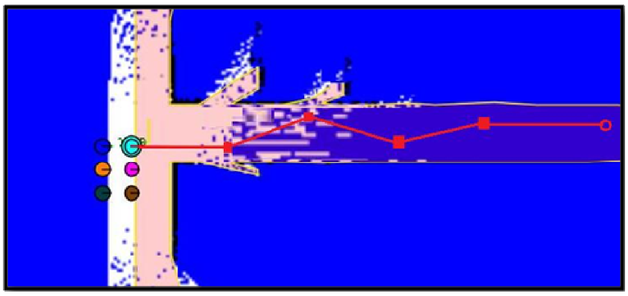

Figure 3. Issue multiple way point paths to the robots

Fractal Dimension Analysis

The reported FD values were calculated using the analysis program developed by Nams (2006), which examines the path and provides a plot of FD as calculated at all spatial scales between a min and max range in addition to an average of all calculated FD values (Craighead, 2009). Since the gross distance will be influenced by different starting points, in order to increase the precision, we used the fractal mean to compute all FD values. This approach computes FD twice, once starting at the first point and going forward through the path; once starting at the last point and going backward through the path. Because FD is highly dependent on the granularity chosen to assess "straightness", after experimentation, we chose a minimum divider of 0.5 which is the diameter of the robot and a maximum of 10 which is the half length of the map. We believe this choice is defensible because it ignores segments smaller than the robot's displacement while still able to distinguish between straight and unstraight lines at its maximum divider.

RESULTS

No difference (mean=767.76 sd= 84.18, mean=801.82 sd=79.62, p=.73) was observed in the areas covered using manual or automated path planning.

Path planning and Monitoring

The One way ANOVA for Fractal Dimension (Figure 4) shows a significant effect for autonomy, $F(1,28)=52.546$, $p<.001$.

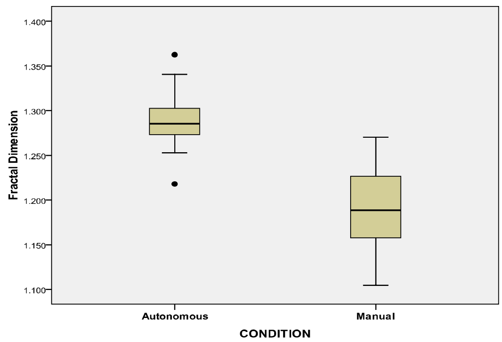

Figure 4. Fractal Dimension

The ANOVA for average accumulated path length for each robot (Figure 5), found significant effect for autonomy, $F(1,28)=10.79$, $p=.002$ as well.

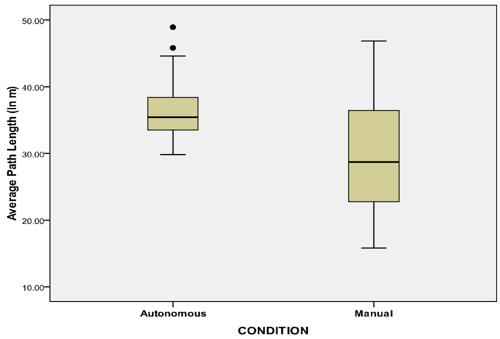

Figure 5. Average Path Length

Teleoperation to bring the robot nearer a victim may help in marking the victim accurately as well as clarify the relation between a camera view and the robot on the map. Teleoperation sequence is a count of the occasions in which marking a victim was preceded by teleoperation. The ANOVA for teleoperation frequency (Figure 6), found significant effects for autonomy, $F(1,28)=10.79$, $p=.002$ with operators in the manual condition teleoperating far less often.

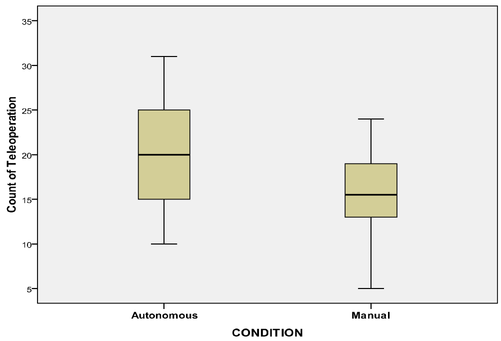

Figure 6. Teleoperation Frequency

Performance

A comparison of regions explored between the autonomy condition and the manual control condition found no difference $F(1,28)=.010$, $p=.921$. This lack of difference shows robots in both conditions covered a similar area.

Table 2. Performance Measures for a one-way ANOVA

Measures	Conditions	Mean	SD	F	p
Victims found	Autonomous	17.07	3.936	4.630	.040
	Manual	13.67	4.685		
Victim/ Region	Autonomous	.0220	.0031	10.71	.003
	Manual	.0175	.0043		
RMS Errors	Autonomous	.4637	.0938	6.344	.018
	Manual	.5812	.1544		

The one-way ANOVA for marked victims (Table 2) comparing manual and autonomous control found a significant advantage for autonomy, $F(1,28)=4.630$, $p=.040$. While participants enjoying automated path planning found more victims, 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 significant effect was found again for autonomy, $F(1,28)=10.71$, $p=.003$.

The related issue of accuracy in marking victims on the laser generated map also favored the robots with automated path planning, $F(1,28)=6.344$, $p<.001$.

Fractal Dimension, Path Length, & Performance

FD ranges from1-2, where higher FD indicates higher levels of tortuosity. In these data FD was found to correlate significantly with Path length and both correlated with each of the other performance measures (Table 3).

Table 3 Correlation between Measures

Measures	Path length	Victim Found	Area Explored	Victim/ Region	RMS	Teleop Sequence
FD	.627**	.437**	.104	.533**	-.205	.459*
Path length	1	.748**	.567**	.633**	-.505**	.372*

Note: **P<0.01 *P<0.05

DISCUSSION

In this paper we investigated characteristics and compatibility of path plans that were self generated (manual), generated by other humans (perceptual search), or generated by a distributed path planning algorithm for use by operators performing a search and rescue task. Our dominant finding was that automated path planning appeared capable of supporting the operators' search and victim marking activities without any decrement relative to human generated paths. The distributed path planning algorithm used in Wang et al. (2009b) produced paths with coverage equivalent to that of human path planners and improvements in victim finding performance comparable to that of participants in Wang et al. (2009a) who used paths generated by other participants.

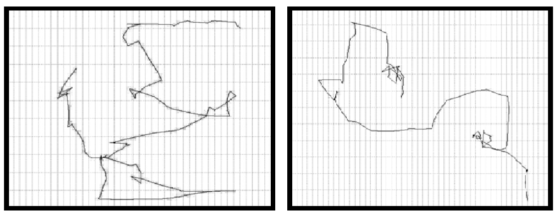

7a. autonomous path

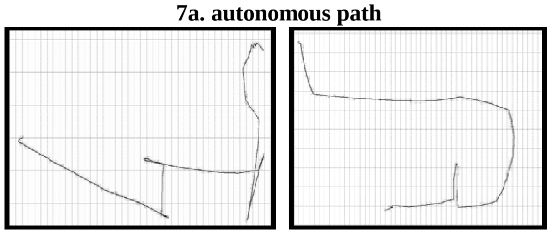

7b. manual path

Figure 7. Autonomous/Manual/Neglected Robot's Path

Despite these similarities, the actual performance of the distributed path planner was substantially different from that of human operators. As differences in fractal dimension indicate, the routes among autonomous, manual, and neglected robots are quite distinct (figure 7). The top paths (7a) which have higher FD were taken from robots in the autonomous condition. These paths are very convoluted, filled with abrupt reversals, crossings, and few straight line segments. This is due to the distributed path planning algorithm that seeks to maximize information gain. Because these decisions are made locally based solely on the exploration frontier and occupancy grid they do not reflect knowledge or expectations about the domain and are subject to revision at frequent intervals when, for example, another robot enters a region it had intended to explore. Differences in teleoperation frequency and duration between autonomous path planning and manual participants observed in Wang et al. (2009b) are the single difference in performance found between use of automated and human generated paths. Wang et al. (2009a) found no such difference in teleoperation for manual and human generated paths.
Because the comparison of path planning methods was not the primary purpose of these experiments some resolution is lost. The strong correspondence between the autonomous path planning condition and high FD in these data, for example, makes results relating FD to performance equally well explained by comparisons between the autonomous and manual conditions. While lowered workload was found for non-path planning operators in Wang et al. (2009a) but not Wang et al. (2009b) an explanation involving freeing of cognitive resources appears most plausible for gains in performance in victims found, victims found/area searched, and accuracy in marking victims.

The results conclude that human operators generated longer and straighter paths with fewer waypoints (figure 7b). Although lower FD represents less contorted paths, it could also reveal lower SA. FD is a well established measurement, however, whether higher or lower FD is better is a matter of interpretation.

While the increased level of teleoperation when working with an algorithmic path planner may suggest a slight penalty, the lack of effects on task performance in exploration and finding victims validates the use of distributed path planners without special adaptation for use in human-robot teams.

ACKNOWLEDGMENT

This research has been sponsored in part by AFOSR FA955008-10356 and ONR Grant N0001409-10680

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*, 24(11-12), 943-967.

Benhamou, S. (1989). An olfactory orientation model for mammals' movements in their home ranges. *Journal of Theoretical Biology*, 139: 379–388.

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

Craighead, J. (2009). *IEEE International Workshop on Safety, Security and Rescue Robotics.*

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.

Gotceitas, V. & Colgan, P.(1988). Individual variation in learning by foraging juvenile bluegill sunfish. *Journal of Comparative Psychology*, 102(3):294–299.

Latombe, J. C. (1991). *Robot motion planning*. New York, NY: Springer.

Lewis, M. (1990) Modeling individual differences at a process control task. *Proceedings of the Human Factors Society 34th Annual Meeting*, Orlando, FL, pp. 917-921.

Mandelbrot, B. (1967). How long is the coast of Britain? Statistical self-similarity and fractional dimension. *Science*, 156: 636–638.

Milne, B.T. (1997). Applications of fractal geometry in wildlife biology. In: Bissonette, J.A. (Ed). Wildlife and Landscape Ecology. *Effects of Pattern and Scale*. New York, NY: Springer, pp. 32–69.

Montemerlo, M., Roy, N. & Thrun, S. (2003). Perspectives on standardization in mobile robot programming: the carnegie mellon navigation (carmen) toolkit. *In Proc. of IROS*, Vol. 3, pp. 2436- 2441.

Nams, V. O. (2006). Improving accuracy and precision in estimating fractal dimension of animal movement paths. *Acta Biotheoretica*, vol. 54, no. 1, pp. 1–11.

Neumeister, H., Cellucci, C. J., Rapp, P. E., Korn, H. & Faber, D. S.(2004). Dynamical analysis reveals individuality of locomotion in goldfish. *The Journal of Experimental Biology*, 207(697-708).

Phillips, F. & Voshell, M. G. (2006). A Novel Metric for Evaluating Human-Robot Navigation Performance. *NATO/PFP UNCLASSIFIED.*

Steinfeld, A., Fong, T., Kaber, D., Lewis, M., Scholtz, J., Schultz, A. & Goodrich, M. (2006). Common metrics for human-robot interaction. *1st Annual Conference on Human-Robot Interaction*, Salt Lake City, Utah.

Voshell, M. G., Woods, D. D. & Phillips, F. (2005). Human-Robot Interaction: From Fieldwork to Simulation to Design. *Proceedings of the Human Factors and Ergonomics Society 49th Annual Meeting*. 26-28 September, Orlando FL.

Wang, H., Chien, S., Lewis, M., Velagapudi, P., Scerri, P. & Sycara, K. (2009b). Human teams for large scale multirobot control, *Proceedings of the 2009 IEEE International Conference on Systems, Man, and Cybernetics*, 1306-1311.

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