2010 IEEE International Conference on Robotics and Automation
Anchorage Convention District
May 3-8, 2010, Anchorage, Alaska, USA

# The Search for Survivors: Cooperative Human-Robot Interaction in Search and Rescue Environments using Semi-Autonomous Robots

Barzin Doroodgar, Maurizio Ficocelli, *Member, IEEE*, Babak Mobedi, and Goldie Nejat, *Member, IEEE*

*Abstract*—**Current applications of mobile robots in urban search and rescue (USAR) environments require a human operator in the loop to help guide the robot remotely. Although human operation can be effective, the unknown cluttered nature of the environments make robot navigation and victim identification highly challenging. Operators can become stressed and fatigued very quickly due to a loss of situational awareness, leading to the robots getting stuck and not being able to find victims in the scene during this time-sensitive operation. In addition, current autonomous robots are not capable of traversing these complex unpredictable environments. To address this challenge, a balance between the level of autonomy of the robot and the amount of human control over the robot needs to be addressed. In this paper, we present a unique control architecture for semi-autonomous navigation of a robotic platform utilizing sensory information provided by a novel real-time 3D mapping sensor. The control system provides the robot with the ability to learn and make decisions regarding which rescue tasks should be carried out at a given time and whether an autonomous robot or a human controlled robot can perform these tasks more efficiently without compromising the safety of the victims, rescue workers and the rescue robot. Preliminary experiments were conducted to evaluate the performance of the proposed collaborative control approach for a USAR robot in an unknown cluttered environment.**

## I. Introduction

With the advancement of robotic research in recent years, mobile robotic systems are being developed to aid rescue workers in urban search and rescue (USAR) operations. In both human-caused and natural disasters, the fundamental tasks at hand are: (i) to find and rescue victims in the rubble or debris as efficiently and safely as possible, and (ii) to ensure that human rescue workers' lives are not put at great risk. Generally, USAR environments are highly cluttered and all robots that operate in these environments do not have a priori information about landmarks in the scene. These conditions make it extremely difficult for robots to autonomously navigate the scenes and identify victims. Therefore, current applications of mobile robots in USAR operations require a human operator in the loop to help guide a robot remotely.

Most robots' relationship to their environments is limited by sensor technologies and cost, where their location in the environment, the layout of the environment, and the presence of victims are usually extracted from a 2D video camera [1]. A human operator in USAR environments faces the important tasks of remembering, recognizing and diagnosing a scene and how the robot is positioned and oriented within the scene merely from a camera. This often leads to disorientation, the robot getting stuck and not being able to identify victims that are present in the scene.

Studies have shown that situational awareness (SA) is essential to the effectiveness of the use of a mobile robot in a USAR operation [2]. SA is defined to be the perception of the objects in an environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future [3]. Having good SA is critical, in fact it has been noted that operators will stop everything they are doing and spend an average of 30% of their time trying to acquire/re-acquire SA, even when they are performing a time-sensitive search and rescue task [2]. The need for high levels of SA in USAR situations can make it difficult for operators to safely navigate the robot and identify victims. To address the challenge of SA in cluttered USAR environments, we have developed a novel real-time 3D mapping sensory system capable of providing 2D and 3D images in real-time as well as identifying landmarks and performing 3D Visual Simultaneous Localization and Mapping (SLAM) [4,5]. The images as well as the 3D map of the environment can be utilized to situate the robot in cluttered USAR environments. The sensory system can be used by operators with minimal training in robotic search and rescue applications. However, in order to be able to utilize our 3D mapping sensor as effective sensory feedback in USAR environments, a robot control architecture for human in the loop operation must be developed.

Although human teleoperation can be effective, the unknown cluttered nature of the environments makes the tasks of robot navigation and victim identification highly challenging. Issues such as latency or a loss of communication can arise. Additionally, operators can become stressed and fatigued very quickly in USAR environments, causing crucial errors in control and victim identification [6]. Semi-autonomous control schemes have been proposed that allow control to be divided between a robot and a human operator [7,8]. In many of these schemes robot controllers are used to perform routine tasks so that the operator can focus on high level control and supervisory tasks. Although these schemes simplify the task of the human operator, the level of autonomy of the robot is fixed.

Manuscript received September 15, 2009. This work was supported in part by the National Science Foundation (U.S.) grant CNS-0721701 and the Natural Sciences and Engineering Research Council of Canada (NSERC).

Barzin Doroodgar, Babak Mobedi, and Goldie Nejat are with the Autonomous Systems and Biomechatronics Laboratory, Department of Mechanical and Industrial Engineering, University of Toronto, Toronto, Ontario, Canada (e-mails: barzin.doroodgar@utoronto.ca, mobedi@ mie.utoronto.ca, nejat@ mie.utoronto.ca).

Maurizio Ficocelli is with the Department of Mechanical Engineering, State University of New York at Stony Brook, Stony Brook, NY, U.S.A (e-mail: maurizio.ficocelli@stonybrook.edu).

Alternatively, some control schemes allow the operator to set the level of autonomy of the robot [9,10]. This allows for the level of autonomy of the robot to be changed continuously by the operator according to need. By allowing the operator to set the level of autonomy, the robot is relying on the experience and knowledge of the human operator to determine when the robot can operate independently. In this paper, we present a unique control architecture for semi-autonomous navigation of a robotic platform utilizing sensory information provided by the real-time 3D mapping sensor. We propose the development of a unique hierarchical reinforcement learning (HRL) control algorithm to allow the robot to learn and make decisions regarding which tasks should be carried out at a given time and whether the human or the robot should perform these tasks for optimum results. By giving the robot the decision making ability to decide when human intervention is required the human operator can take advantage of the robot's ability to continuously learn from its surrounding environment.

## II. System Architecture

The overall system architecture consists of two main components: (i) a real-time 3D mapping sensory system, and (ii) a novel control architecture that allows the robot to change the amount of autonomy it has during search and rescue operations in cluttered environments.

### A. 3D Mapping Sensory System

The recent literature has proposed the use of a combination of sensory systems including variations of video and thermal cameras, IR sensors, range finders, time-of-flight sensors, gyroscopes and accelerometers to improve a robot operator's SA in cluttered USAR environments [i.e., 11-13]. In particular, the use of stereovision, laser range finders and time-of-flight sensors demonstrate the need for 3D information to be obtained from a robot's environment for the two crucial tasks of robot navigation and victim identification.

Our work consists of utilizing a real-time structured light sensory system based on a digital fringe projection and phase shifting technique for 3D mapping in USAR environments [4,5]. The sensor can directly map rubble in 3D and in real-time at a frame rate of up to 40 fps and at a resolution of 640x480 pixels. The performance of the sensor is independent of three main limiting factors of current sensors: (i) the use of a scanning mechanism, which is time-consuming in real-time applications, (ii) slow scanning speed; the sensor can provide 3D mapping in real-time, and (iii) the illumination conditions of the environment; the sensor will successfully work in dim lit and dark environments. In addition to providing real-time 2D and 3D images to the human operator, sensory information is utilized to create a 3D virtualized map of the disaster environment with respect to a world frame in which victims can be found. In order to generate a 3D map of the environment, we have developed a unique and robust 3D Iterative Closest Point (ICP) –based SLAM technique in which landmarks in the environment are identified based on the Scale Invariant Feature Transform (SIFT) technique. The novelty of the method is in the utilization of both 3D and 2D images. Distinguishable landmarks are determined effectively within the images utilizing a combination of SIFT invariant features, and depth and geometric clustering techniques. The landmarks are then matched in consecutive frames and utilized as inputs for 3D ICP-based SLAM. For more details on the sensory system, the reader is referred to [4,5]. In addition to the sensory hardware, an enclosure and vibration isolation system has been developed for the sensory system to protect the system as the robot traverses the uneven and dusty terrain of USAR scenes. We have integrated our sensory system into a rugged robotic platform capable of traversing cluttered environments, Fig. 1.


Fig. 1. USAR sensory system on a rugged robotic

### B. Control Architecture

In general, robot control architectures can be defined as deliberative, reactive or hybrid [14]. Deliberative control consists of high-level planning, whereas reactive control is based on directly utilizing sensory information for low-level commands. Traditionally, the reactive architectures have been considered as behavior-based control. In behavior-based robotic control, the overall control of the robot is shared between a set of perception-action units known as behaviors [14]. Our proposed control architecture is based on a multi-layer hybrid behavior-based approach, Fig. 2.

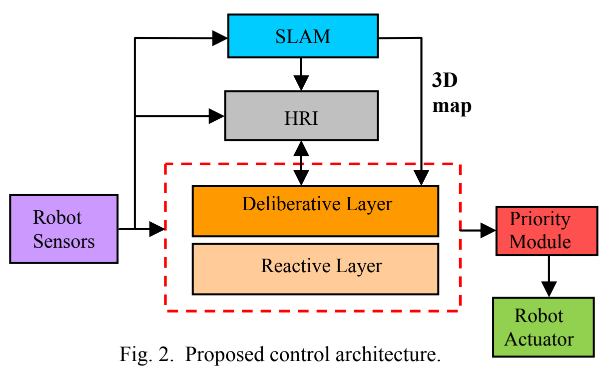


Fig. 2. Proposed control architecture.

The robot control architecture proposed in this work contains the following modules:

Robot Sensors: The inputs to the control system include the robot's internal/external sensory information. In particular, the 3D mapping sensor is utilized herein to gather 2D and 3D data from the environment. In addition, five infrared sensors distributed around the robot are also utilized.

SLAM Module: The 2D and 3D images provided by the mapping sensor are utilized in the SLAM module to identify and match 3D distinguishable landmarks, as the robot moves within the scene, in order to create a 3D global map. In order to build the 3D map in world coordinates, the robot must be able to localize itself utilizing these landmarks. This is achieved by stitching consecutive 3D range information corresponding to the landmarks provided by the sensor via the ICP-based SLAM method [4,5].

Deliberative Layer: The deliberative layer contains the decision making capabilities required to analyze real world scenarios. Real-time sensory data from the mapping sensor as well as the generated 3D map of the explored environment are utilized as inputs into the deliberative layer. Since the robot is designed to be semi-autonomous, it is within the deliberative layer where the level of autonomy is primarily decided. In particular, it is in this layer where the balance between human control and autonomous control is decided. If human control is prominent, then the decision making within this module is made by the human operator.

HRI Interface: The HRI interface module consists of the user interface for the operator. The interface allows the operator to obtain sensory information from the environment and the robot in order to control the robot's motion.

Reactive Layer: The reactive layer is used mainly for interaction situations that require an immediate response, i.e., the robot is in a dangerous situation that can cause it to be damaged.

Priority Module: The priority module decides the final behavior of the robot based on the precedence of information regarding robot health and safety during interaction.

Robot Actuators: The robot actuators module consists of the robot's motors and motor control boards. Herein, the appropriate motor signals are applied based on robot behavior information.

In the remainder of this paper we present the design of the deliberative layer since it is the main decision making module of the control architecture. Namely, within this architecture, the deliberative layer is the module responsible for providing the robot with the ability to learn and make decisions regarding which rescue tasks should be carried out at a given time and whether the robot or human can perform these tasks more quickly and efficiently without compromising the safety of the victims, rescue workers and the rescue robot. The remaining subsection presents the detailed design of the deliberative layer.

*1) Deliberative Layer:* The deliberative layer will allow the level of autonomy of the robot to vary depending on the robot's ability to function in the given environment. A learning algorithm will be used to allow the robot to make decisions regarding which tasks should be carried out at a given time and who (the robot or human) should perform these tasks for optimal results. For our proposed control architecture, a HRL algorithm is used for the robot intelligence and as the decision making scheme for the deliberative layer. The advantages of HRL are that it does not require information about the environment to be provided a priori, and the learning process is on-line.

***Hierarchical Reinforcement Learning***

Three main HRL approaches have been explored in detail in the literature: (i) the Options approach [15], (ii) the Hierarchical Abstract Machines (HAMs) approach [16], and (iii) the MAXQ approach [17]. Of these three HRL approaches, MAXQ requires less knowledge of the system when designing the learning policy, which is advantageous when dealing with unknown USAR environments. Hence, in this paper, we propose a MAXQ approach to be utilized for semi-autonomous control of a search and rescue robot. The task graph for the robot is presented in Fig. 3. The MAXQ approach is able to support temporal abstraction, state abstraction, and subtask abstraction, each of which is important in search and rescue applications. The need for temporal abstraction exists in this application since actions may take varying amounts of time to execute depending on the complexity of the scene and the location of the robot within the scene. State abstraction is important since when the robot is navigating to a particular location only the target location is important, the reason why it is navigating to that location is irrelevant and should not affect the robot's actions. Subtask abstraction is necessary because it allows subtasks to be learned only once; the solution can then be shared by other subtasks. For example, the *Navigate* subtask is used by both the *Navigate to Unvisited Regions* and the *Victim Identification* subtask, Fig. 3.

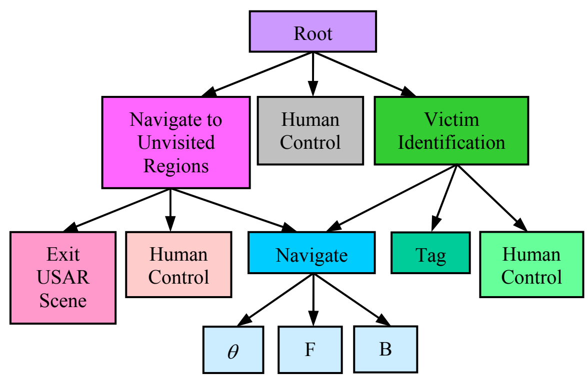

Fig. 3. Task graph for the USAR semi-autonomous robot platform.

In order to construct a MAXQ decomposition for the semi-autonomous robot problem, we must first identify the individual subtasks that will be used to solve the overall problem of finding a victim in a cluttered unknown environment. The robot will also have the ability to give control to a human operator whenever an operator can perform a function more efficiently. For this application the following task and subtasks are defined:

- *Root* – This is the overall USAR task where the goal is to find victims within a cluttered USAR environment.
- *Navigate to Unvisited Regions* – In this subtask the goal is to have the robot explore unvisited regions within the search and rescue environments.
- *Victim Identification* – In this subtask the goal is to scan the viewable scene to identify victims in the environment.
- *Human Control* – In this subtask the goal is to pass control over to a human operator if the robot cannot complete the tasks at hand while in the environment.
- *Navigate* – The goal of this subtask is to move the robot from its current location to a target location while also performing obstacle avoidance.

In addition to subtasks, a number of primitive actions can be implemented. These primitive actions include $\theta$, $F$ and $B$, which are used by the *Navigate* subtask to move the robot around the environment by rotating the robot by angle $\theta$ and moving the robot forward ($F$) or backward ($B$). The *Victim Identification* subtask uses the *Tag* action to mark a victim that is found in the USAR scene for future reference. The action *Exit USAR Scene* is used by the *Navigate to Unvisited Regions* subtask to guide the robot out of the USAR scene once it is determined that the environment has been fully explored. All tasks, subtasks and primitives are summarized in the task graph in Fig. 3.

***MAXQ State Definitions***

A set of states, $S$, have been determined for the aforementioned tasks and subtasks to be utilized within the MAXQ framework. In particular, the state functions for the robot's overall task and each of the robot's subtasks are defined as follows: (i) *Root*: $S(V, L_R, M_{xyz})$, *(ii) Navigate to Unvisited Regions:* $S(L_R,M_{xy})$, (iii) *Victim Identification:* $S(L_{V/R},M_{xy})$, and *(iv) Navigate:* $S(C_i)$. $V$ represents a potential victim. To help identify victims that are trapped within the rubble, image processing techniques using human features and shape analysis can be applied directly to the 2D images provided by the 3D mapping sensor and/or an additional thermal camera can be utilized. $L_R$ represents the robot's location with respect to a global coordinate frame (determined from the SLAM module in the control architecture) and $L_{V/R}$. represents the location of a potential victim relative to the robot as determined by the 3D mapping sensor. $M_{xyz}$ represents the 3D map created by the SLAM module and $M_{xy}$ represents the 2D grid map of the robot's explored environment at a certain time. $C_i$, where $i$=1 to $n$, represents the information of the $n$ neighboring cells of a 2D grid map of the environment surrounding the robot at a particular location. The significance of the 2D grid map is that it allows for a simple representation of the environment that the robot can utilize effectively for decision making. The grid is composed of an array of cells. Each cell in the grid defines the status of an area (i.e., in front, behind or to the sides of the robot) in the real environment. The size of the cells is defined by the range of the 3D mapping sensor. $C_i$ represents the status of cell $i$ in the grid and can be defined as follows with respect to cell information: an obstacle is present in the cell as detected by the robot's sensors, the cell has been visited before by the robot, the cell is unvisited by the robot but has been detected as an obstacle-free cell, and the cell information is unknown in which case the cell has not been explored and there is no sensory information available. As the robot explores the scene unknown cells are updated into either obstacle, visited, or unvisited cells in the grid. The information regarding the status of each cell is simply obtained from the depth information from the 3D map and the infrared sensors around the robot. Note that in the state definition of *Navigate* only the surrounding cells of the robot are used, not the entire 2D map. Based on the status of the robot's surrounding cells, $Q$-learning determines the optimal primitive action to take during navigation. The rewarding function rewards desirable behavior such as obstacle avoidance and global exploration of the scene with positive rewards and discourages undesirable behavior such as revisiting previously explored regions and colliding with obstacles with negative rewards.

## III. EXPERIMENTS

Preliminary proof-of-concept experiments were conducted to evaluate the performance of the proposed HRL control approach for our rugged USAR robot (55cm by 65cm) in an unknown cluttered environment. A cluttered $12m^2$ USAR-like scene was developed consisting of different types of objects that may be found in a disaster scene, Fig. 4. The objects included wood, metal, plastic, brick, ceramic, concrete, paper, cardboard, plaster, rubber-like polymers and rocks. In addition, the environment consisted of 8 victims represented by dolls and mannequins that were distributed within the scene. Rubble in the scene was strategically placed such that the robot would have to explore around corners and barriers to search for victims. A number of the victims were partially obstructed in which case only a portion of their body was visible, i.e., limbs or head. Five operators ranging in age from 18-35 years participated in the experiments. None of the operators had any experience in remote navigation or exploration.

Fig. 4. USAR-like scene.

### *A. Sensory System*

The structured light sensory system utilized in these experiments consists of a DLP projector with a native resolution of 800x600 pixels and a Prosilica GE680C CCD camera with a resolution of 640x480 pixels and a frame rate of 200Hz. The effective measurement range of the sensory system is 0.2–1.6 m and the current viewing angle for our experiments is approximately $11^o$. The accuracy of the system is 0.1 mm. The sensory system was able to obtain corresponding 2D texture and 3D depth images at 40 fps. We also placed a 2D video camera at the front of the robot that provided the operator with continuous 2D video feed of the scene for human teleoperation. For the purpose of these experiments, a simple skin-color blob tracking technique was sufficient to test the semi-autonomous control architecture. We are currently developing more robust techniques for victim detection to incorporate within the architecture as a part of our future work. The robot was also equipped with five infrared sensors distributed along the sides and back of the robot (i.e. two sensors on each side and one in the back) with a range of 0.1 to 0.8 m. Sensory information is communicated to the robot's deliberative layer utilizing Bluetooth communication via an ATMEL microcontroller. The microcontroller is also utilized to send control commands to the robot actuators.

### B. Robot Control

The operators were asked to navigate the robot within the scene in two different trial sets: (i) having full teleoperated control over the robot, and (ii) while the robot was in semi-autonomous mode, in which case the human and the robot shared decision making tasks as needed. For the latter trials, the MAXQ algorithm and task graph presented in Section II were utilized. An initial training stage for the MAXQ was implemented outside of the scene. For each set of experiments, the level of difficulty traversing the scene remained the same; however, the objects and the victims were placed in different locations to minimize human expert knowledge of the scene.

### C. HRI User Interface

During human operation of the robot, the interface presented in Fig. 5 was utilized. The control of the robot is achieved via a gamepad with two mini joysticks and a set of buttons. The 3D map is always available to the operator to allow for situations where control is passed from the robot to the operator.

(a) (b) (c)

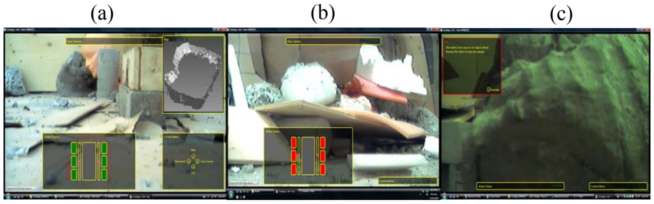

Fig. 5. (a) 2D image, 3D map, robot status (green indicates the wheels of the robot are moving) and control menu display, (b) robot status when wheels are not moving, and (c) alert window using sensory information from infrared sensors to warn the operator if the robot is in danger.

### *Experimental Results and Discussions*

Figs. 6, 7 and 8 present the performance comparisons of teleoperated control versus semi-autonomous control of the search and rescue robot. It is evident from the results shown in Fig. 6 that the total number of collisions was significantly reduced in the semi-autonomous search and rescue operation for the majority of the trials. These results suggest that the autonomous navigation can be more effective in using the sensory information for controlling the robot in cluttered scenes. It is important to note that the collisions that were detected in the semi-autonomous mode occurred on rare occasions where the robot was navigating through obstacle free sections of the map and colliding with parts of the rubble that were not within the field of view of the sensory system: Namely, the wheels of the robot were hitting small objects that were below the line of sight of the infrared sensors. On the other hand, during human teleoperation, even though these types of collisions occurred as well, the majority of collisions took place while the sensory system had detected the obstacles. In particular, many of these collisions occurred in parts of the scene where the robot was navigating through small openings in rubble piles or trying to make a sharp turn within a rubble filled area. In many cases for the teleoperated trials, the operators used a brute force approach to try to fit within narrow passages. In one trial, the robot became stuck in the scene for approximately 120 seconds during teleoperation before the robot could continue navigating again. It is also interesting to note that in addition to object-based collisions, three out of the five operators accidentally collided with a victim during teleoperation.

On average approximately 5 victims were found in the teleoperated robot mode within the trials versus 7 in the semi-autonomous robot mode, Fig. 7. False victim identification was also made by one operator during teleoperation who defined a small red cushion to be a victim.

Fig. 8 presents the percentage of the overall search and rescue scene that was traversed by the robot in all five trials. Only one operator traversed the full scene during teleoperation, whereas utilizing the semi-autonomous mode, the robot was able to traverse the total area of the overall scene for each of the five trials utilizing the MAXQ approach. The total time for each trial is presented in Table I and is defined herein as the time it took to identify victims and exit the scene. The average total time for completion of each trial was determined to be: 337 seconds for teleoperated mode and 167 seconds for semi-autonomous mode. Hence, the total operation time was decreased on average by 50% while the robot was in semi-autonomous mode.

In general, semi-autonomous control of the robot appeared to have better performance results. In addition to the experiments, the human operators were also asked to complete a survey reflecting their experiences. Within the survey, they were asked questions regarding their stress level during robot operation and in their opinion the impact semi-autonomous control had on their decision making abilities. In regards to stress levels, all 5 participants mentioned that they felt stress during full teleoperation of the robot. These stress levels varied from 3 participants feeling little stress to 2 participants (Participants 2 and 3) feeling a lot of stress. For semi-autonomous control, 4 participants stated they felt no stress and 1 participant (Participant 2) mentioned that he felt a lot of stress. Since Participant 2 felt high levels of stress in both scenarios, he clarified that in comparison, he was more stressed during robot teleoperation. All 5 participants agreed that they found that the robot having autonomous capabilities improved their own decision making abilities in the scene due to the fact that they did not have to continuously multi-task and it was easier for them to understand the overall scene better. A number of them mentioned that they also did not worry as much about the robot getting stuck or hitting obstacles which seemed to be one of their main concerns. The preliminary experimental and survey results validate that a semi-autonomous control architecture utilizing a MAXQ approach for rescue robots has the potential of improving the success rate of a search and rescue mission.

## IV. Conclusion

In this paper, we propose a unique control architecture for semi-autonomous navigation of a robotic platform. The control architecture utilizes a HRL algorithm to provide the robot with the ability to learn and make decisions regarding which rescue tasks should be carried out at a given time and whether autonomous control or human control should be utilized to perform these tasks more quickly and efficiently without compromising the safety of the victims, rescue workers and the rescue robot. The preliminary experiments show the potential of further exploring the integration of a HRL based approach for semi-autonomous robotic control in unknown cluttered environments. Future work will consist of performing extensive experiments with a larger pool of participants and USAR scenes. Furthermore, we will look into the design of guidelines to help balance the delicate partnership between human and robot control in these complex environments. We envision that this approach will allow the robot to learn from different situations and scenarios it will be placed in and assist the human operator in a number of tasks that need to be completed during search and rescue. Thus, minimizing the stress and burden placed on the operator to solely complete the tasks.

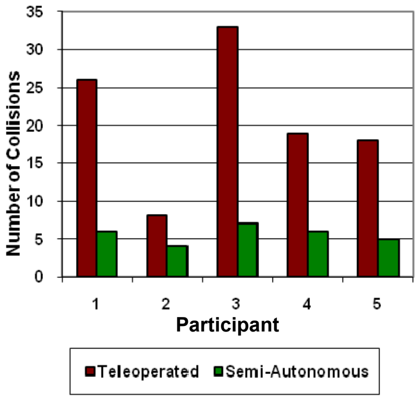


Fig. 6. Number of collisions.

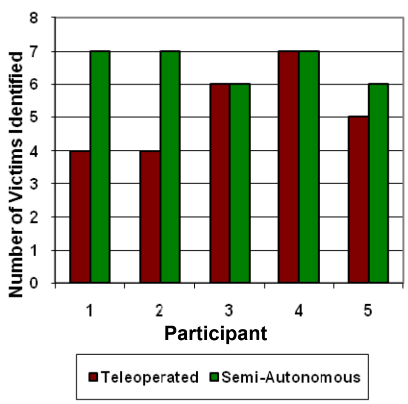


Fig. 7. Number of victims identified.

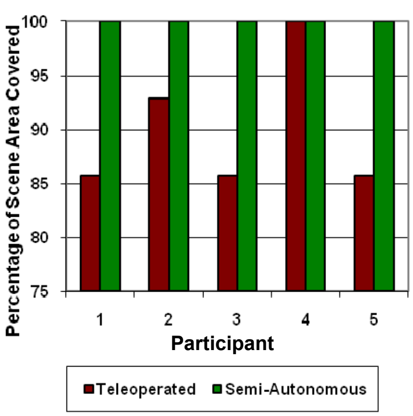


Fig. 8. Percentage of scene explored.

TABLE I: SUMMARY OF TRIAL TIMES.

| **Participant** | **1** | **2** | **3** | **4** | **5** |
|---|---|---|---|---|---|
| **Teleoperated Trial Times (s)** | 418 | 159 | 505 | 271 | 334 |
| **Semi-Autonomous Trial Times (s)** | 215 | 148 | 194 | 130 | 148 |

## ACKNOWLEDGMENT


The authors would like to thank Z. Zhang, H. Guo and Professor P. Huang for all their contributions to the development of the real-time 3D mapping sensor.


## REFERENCES


[1] R. R. Murphy, "*Human-Robot Interaction in Rescue Robotics*," IEEE Trans. on Systems, Man, and Cybernetics-Part C: Applications and Reviews, Vol. 34, No. 2, pp. 138-153, 2004.

[2] J.L. Drury, J.L., B. Keyes, H.A. Yanco, "*LASSOing HRI: Analyzing Situation Awareness in Map-Centric and Video-Centric Interfaces*", ACM/IEEE Int. Conf. on Human-Robot Interaction, pp. 279-286, 2007.

[3] M.R. Endsley, "Design and evaluation for situation awareness enhancement", Human Factors Society 32nd Annual Meeting, Santa Monica, CA, pp. 97–101, 1988.

[4] Z. Zhang, H. Guo, G. Nejat, and P. Huang, "*Finding Disaster Victims: A Sensory System for Robot Assisted 3D Mapping of Urban Search and Rescue Environments*," IEEE, Int. Conf. on Robotics and Automation, Rome, pp. 3889-3894, 2007.

[5] Z. Zhang and G. Nejat, "Robot-Assisted 3D Mapping of Unknown Cluttered USAR Environments," the Disaster Robots Special Issue, Advanced Robotics Journal, vol. 23, pp. 1179–1198, 2009.

[6] J. Casper, R. Murphy, "Workflow study on human-robot interaction in USAR", IEEE Int. Conf. on Robotics and Automation, pp. 1997-2003, 2002.

[7] R. L. Murphy and J. J. Sprouse, "Strategies for searching an area with semi-autonomous mobile robots," Robotics for Challenging Environments, pp. 22-28, 1996.

[8] D. Doroftie, E. Colon and Y. Baudoin, "A modular control architecture for semi-autonomous navigation," Int. Conference on Climbing and Walking Robots, pp. 712-715, 2006.

[9] R. Wegner and J. Anderson, "B*alancing robotic tel*eoperation and autonomy for urban search and rescue environments," Lectures in Computer Science, Vol. 3060, pp. 16-30, 2004.

[10] D. J. Bruemmer, D. D. Dudenhoeffer, J. L. Marble, "*Dynamic-Autonomy for Urban Search and Rescue*," AAAI Mobile Robot Workshop, 2002.

[11] B. McKinnon, J. Baltes, and J. Anderson, "*Region-Based Approach to Stereo Matching for USAR*," RoboCup 2005, LNAI 4020, pp.452-463, 2006.

[12] B. Keyes, R. Casey, H. A. Yanco, B. A. Maxwell, Y. Georgiev, "*Camera Placement and Multi-Camera Fusion for Remote Robot Operation*," IEEE Int. Workshop Safety, Security, Rescue, 2006.

[13] M. W. Kadous, R. K.-M. Sheh, C. Sammut, "*CASTER: A robot for Urban Search and Rescue*," Australasian Conference on Robotics and Automation, pp. 1-10, pp. 1-7, 2005.

[14] P. Pirjanian et al., "*CAMPOUT: A control architecture for multi-robot planetary outposts*," SPIE – The International Society for Optical Engineering, pp. 221-230, 2000.

[15] R. S. Sutton, S. Singh and D. Precup, "*Between MDPs and semi-MDPs: a framework for temporal abstraction in reinforcement learning*," Artificial Intelligence, Vol, 112, No. 1, pp. 181-211, 1999.

[16] P. Parr and S. Russell, "*Reinforcement learning with hierarchies of machines*," Conf. on Advances in Neural Information Processing Systems, MIT Press, 1997.

[17] T. G. Dierrerich, "*Hierarchical reinforcement learning with maxq value function decomposition*," Journal of Artificial Intelligence Research, Vol. 13, pp. 227-303, 2000.