

Submitted to IEEE Transactions on Robotics and Automation as a regular paper

- Title: Mixed-Initiative Control of Multiple Heterogeneous Robots for Urban Search and Rescue
- Authors: Robin R. Murphy, Jennifer Casper, Mark Micire, Jeff Hyams
- Corresponding Author: Dr. Robin R. Murphy
- Address: Computer Science and Engineering  
University of South Florida  
4202 East Fowler Ave ENB 118  
Tampa FL 33620, USA
- email: murphy@csee.usf.edu
- phone: (813) 974-4756
- fax: (813) 974-5456

**abstract:** This article describes a mixed-initiative system for Urban Search and Rescue, its implementation on a team of heterogeneous robots, and data collected with the NIST USAR Standard Test Course. It contributes to the justification and understanding of heterogeneous robot teams in perceptually and navigationally challenging domains, and to the theory of architectures for mixed-initiative systems. The mixed-initiative control scheme used a novel *Three-Agent Society* organization which places an Intelligent Assistant Agent as the middleware between the Physically-Situated (remote robots) and Cognitive (human) Agents. The Intelligent Assistant Agent provided perceptual assistance, cueing the teleoperator to the possible presence of victims using a fusion of four affordances: heat, motion, skin color, and color difference. The mixed-initiative system was novel in that the robot took the initiative for perception rather than for navigation. Navigation through the confined space was done cooperatively using *collaborative teleoperation*. Data collected at the 2000 AAI Mobile Robot Competition showed that the robot team found four times as many victims under more demanding conditions as any competitor. The robots cooperatively identified 209% (23) more victims than they would have individually (11). In addition, an analysis of the recordings showed that three types of navigation errors were introduced by the teleoperators: *iterative course corrections* (2 per minute per robot), *collisions* with the environment (0.13 per minute), and *pose errors* (0.3 per minute).

**keywords:** artificial intelligence, multi-agents, heterogeneous robots, mixed-initiative control, urban search and rescue, distributed sensing, sensor fusion, vision-based robotics



# Mixed-Initiative Control of Multiple Heterogeneous Robots for Urban Search and Rescue

R.R. Murphy, J. L. Casper, M. J. Micire, J. Hyams

*Abstract*— This article describes a mixed-initiative system for Urban Search and Rescue, its implementation on a team of heterogeneous robots, and data collected with the NIST USAR Standard Test Course. It contributes to the justification and understanding of heterogeneous robot teams in perceptually and navigationally challenging domains, and to the theory of architectures for mixed-initiative systems. The mixed-initiative control scheme used a novel *Three-Agent Society* organization which places an Intelligent Assistant Agent as the middleware between the Physically-Situated (remote robots) and Cognitive (human) Agents. The Intelligent Assistant Agent provided perceptual assistance, cueing the teleoperator to the possible presence of victims using a fusion of four affordances: heat, motion, skin color, and color difference. The mixed-initiative system was novel in that the robot took the initiative for perception rather than for navigation. Navigation through the confined space was done cooperatively using *collaborative teleoperation*. Data collected at the 2000 AAI Mobile Robot Competition showed that the robot team found four times as many victims under more demanding conditions as any competitor. The robots cooperatively identified 209% (23) more victims than they would have individually (11). In addition, an analysis of the recordings showed that three types of navigation errors were introduced by the teleoperators: *iterative course corrections* (2 per minute per robot), *collisions with the environment* (0.13 per minute), and *pose errors* (0.3 per minute).

## I. INTRODUCTION

Urban Search and Rescue (USAR) involves victim detection and extrication from urban structures such as collapsed buildings after an earthquake, hurricane, or terrorist attack. These structures are characterized as confined spaces where an agent cannot turn around in place, and consist of rubble and destroyed interiors. Due to the shortage of certified workers and dogs that can enter and work in a collapsed site, plus the extreme time pressure on survivors, mobile robots have been proposed as a third “leg” of a human-dog-robot triad.[1], [2] As detailed in [1], [2], [3], small robots can perform tasks that neither humans, dogs, nor existing tools can do.

Previous work [4] has established the utility of teams of heterogeneous mobile robots and polymorphic, or shape-shifting, robots for the victim detection task in urban search and rescue. At first, the obvious divide-and-conquer nature of search of physical spaces suggests a homogeneous, not a heterogeneous, team. However, heterogeneity is desirable because small robots capable of operating in confined spaces (voids of less than 2 meters cross-section) are unlikely to carry all the sensing modalities used in vic-

tim identification, e.g. thermal imaging, acoustic, video. According to FEMA statistics taken from disasters in the United States, only 50% of the victims will be surface victims, easy to detect. The remainder will be lightly covered (30%) or trapped in voids spaces (15%) or entombed (5%).[3] The non-surface victims pose the most potential gains for robotic detection. Video cameras can be used to detect those victims who are conscious and moving as long as they are not entombed behind rubble. Thermal cameras can detect whether a visible motionless victim is alive but unconscious using body heat, and can often perceive body heat through light coverings of rubble. Note that the incorporation of a human in the loop does not overcome the limitations of video in distinguishing death or unconsciousness. The information as to the state of the victim is very important to the Incident Commander who must decide which areas of the building to excavate.

While the above example shows that miniature thermal FLIR cameras are favored, they cost on the order of \$14,000 USD versus the \$500 USD cost of CCD cameras. It is unlikely that in the near future all robots will be equipped with FLIRs. A solution to these constraints is to distribute the sensing modalities across numerous robots, then have the robots cooperate. For example, a robot with a video camera may identify possible victims but be unable to determine if they are unconscious or dead. The robot with the FLIR can then visit the locations and make the final disambiguation.

The need for heterogeneity also emerges from the diversity in mechanisms needed to navigate through different voids and rubble types. Polymorphic robots capable of changing shapes, such as the RWI Urban, are now available and being tested in USAR situations. Polymorphic robots have a navigational advantage in being able to enter and navigate through more challenging voids spaces. However, controlling the pose of the robot is often difficult even with teleoperation, and as will be shown in this article, leads to a significant number of poses where the robot cannot easily recover and resume navigation.

This article describes a mixed-initiative system for controlling a team of heterogeneous robots, with both physical heterogeneity and distributed sensing capabilities, and its application to USAR. The mixed-initiative system is based on a novel *three-agent society architecture*, following [5]. In this architecture, a middleware Intelligent Assistant Agent complements the Physically-Situated Agents (remote robots) and Cognitive Agents (human). The mixed-

initiative approach reported here is also unique because the robot takes the initiative in victim detection, not navigation. The Intelligent Assistant Agent uses an intelligent vision sub-agent to *cue* the operator’s attention to possible victims using *behavioral sensor fusion* of low order computational complexity vision algorithms. This design decision to focus on perception is an outcome of our direct experience with rescue workers, detailed in Sec. III-C. The mixed-initiative system also supported cooperative navigation and problem solving through the use of *collaborative teleoperation*.

The article presents data collected from the 2000 AAAI Mobile Robot Competition USAR event held in August, 2000, in Austin, Texas, at the National Conference of the American Association for Artificial Intelligence (AAAI). The University of South Florida (USF) fielded a heterogeneous team consisting of two different outdoor robots, a RWI ATRV wheeled vehicle and a RWI polymorphic (shape shifting) Urban, using the mixed-initiative system. The data collected shows that the robots were able to cooperatively i) correctly identify four times as many robots as other teams during the final competition, ii) identify 133% more victims cooperatively than independently, and iii) navigate two challenging arenas that no other team was able to enter and that one of the USF robots would have been unable to navigate without collaborative teleoperation. The data also shows the error rate for navigation in confined spaces by both the polymorphic and fixed platforms and the time spent in collaborative teleoperation mode.

The article is laid out as follows. Sec. II reviews the related work in robotic teams for USAR and mixed-initiative systems. The approach taken by USF is detailed in Sec. III. It first describes the Three-Agent Society organization, the Physically-Situated Agents (robots) used for the competition, and Intelligent Assistant Agent. The section concentrates on the vision sub-agent which aids the operator in detecting victims. The use of low computational complexity algorithms exploiting directly perceivable affordances gives the USAR community a set of algorithms to use for their systems. On a theoretical level, the vision sub-agent serves as an example of how Gibsonian direct perception is viable for complex domains. Sec. IV presents the data collected at the AAAI Competition in terms of victim detection rates and time spent in collaborative teleoperation mode. The ramifications of the results are discussed in Sec. V, presenting additional data on the types and frequency of operator navigation errors (iterative course corrections, collisions, pose errors). This taxonomy and error rate data is expected to be useful to any mobile robot mixed-initiative system, especially those using a polymorphic robot. Overall, the article contributes to the justification and understanding of heterogeneous robot teams in perceptually and navigationally challenging domains, and to the theory of architectures for mixed-initiative systems.

## II. RELATED WORK

The robotics work most closely related to the work reported in this article falls into two categories: intelligent robots designed explicitly for USAR applications and mixed-initiative control systems.

### A. Intelligent Robots for USAR

Urban search and rescue tasks have been considered as a motivation for robotics since at least 1983 [6]. Tangible robotic work has been conducted primarily in constructing, or proposing, platforms suitable for operating in confined space and rubble[7], [8], [9], [10]; these efforts generally ignore the issues of semi- and full-autonomy in victim detection. Blitch [1] created an intelligent assistant for aiding the USAR Incident Commander in choosing which resources (people, dogs, robots) to deploy given the characteristics of a search site and has produced a series of overviews of the opportunities for artificial intelligence to assist with USAR.[2] Work by Murphy et al has been the source of many of the algorithms explicitly aimed at victim detection [11], [4], [12], though that should change with the advent of the RoboCup Rescue Competition Physical Agent League.[13] Other groups are looking at perception for search and rescue, but are not considering confined space situations, for example [14].

Our ongoing work in marsupial and shape-shifting robot teams explores using cooperative robot teams for USAR [4] and a preliminary investigation won one of three 1998 Robotics Industries Association paper awards.[15] The only other efforts in using cooperative robotic teams explicitly intended for rescue operations that we are aware of to date is Jennings et al [16]. However, that work concentrated on developing distributed control algorithms with no communication to allow multiple robots to autonomously swarm around an object and “rescue” it by manipulating it. This is a different problem set than those in urban search and rescue, even in terms of victim extrication, due to the importance of extricating victims in a concerted, probably single, movement.[3] The work also used RWI indoor research robots with blocks-world objects, limiting its applicability to more perceptually and navigationally challenging domains.

### B. Mixed-Initiative Efforts

Our approach assumes a mixed-initiative style of human-robot interaction to combat the extreme perceptual and physical demands of the USAR domain. Mixed-initiative systems tend to focus on the blending of deliberation (planning) by the human and reaction (perception-action) by the robot. See Kortenkamp [17] for a tutorial on mixed-initiative systems, also known as adjustable autonomy.

The mixed-initiative approach taken in this article is most similar to that of Horiguchi, Sawaragi and Akashi[18] which attempts to divide tasks based on perception, rather than on planning, using Kirlik’s Generalized Lens Model. According to [18], the generalized lens model classifies tasks into four categories:

1. *proximal perception, proximal action*. This is essentially a purely reactive behavior, where the direct perception is sufficient for a direct action by the agent.
2. *proximal perception, distal action*. Here the action requires some type of inference, rather than a stimulus response action.
3. *distal perception, proximal action*. In this case, perception is incomplete and requires inference, but the appropriate action is clear once perception is complete.
4. *distal perception, distal action*. Both the perception and action require inference.

[18] used the generalized lens model to describe the interface characteristics for displaying information to the user. Our work takes a different approach. It assumes that victim detection of non-surface victims requires distal perception (e.g., human performs explicit recognition) and distal action (e.g., human decides what to do next), but proximal cues (e.g., direct perception by the robot, or Gibsonian affordances [19]) can trigger assistive proximal actions (e.g., responses by the robot). In this case, victim cues such as heat, motion, and color, can trigger the interface to display these percepts and actively alert the operator to the possible presence of a victim. The operator can then infer whether a victim is present or if active sensing is required to disambiguate the scene.

### III. APPROACH

Our approach to mixed-initiative systems for USAR is novel in i) its organization of intelligence, ii) its focus on automating the perceptual tasks rather than the navigational tasks, and iii) its use of collaborative teleoperation to improve navigation. Traditional mixed-initiative work has viewed teleoperation as a two-agent society consisting of the Cognitive Agent and the Physically-Situated remote robot. This work considers mixed-initiative as a three-agent society. Traditional mixed-initiative systems have concentrated on relieving the operator of navigation tasks, thereby freeing the operator to provide *mission sensing*, e.g., look for landmarks, targets, etc. This focus on navigation is perhaps motivated by the availability of a large body of work in autonomous navigation. However, our experience with USAR suggests the opposite: that it is preferable for the physically situated agent to provide mission sensor and the cognitive agent control navigation.

There are two supporting arguments for having the robot take the initiative in perceptual search. First, perceptual search for victims appears even more demanding than navigation in confined spaces. This is because many factors prevent an operator from actually “seeing” a victim in rubble. The perceptual context is a factor, since the environment has uncontrolled, extreme lighting conditions, the operator is given a reduced field of view and unnaturally low viewpoint from most small rovers, and expected to use exotic sensors such as thermal imagers which often produce hard to interpret results. Perceptual search is also challenging because of the need to pay attention and search for very small cues of survivors, while under significant time and emotional pressure. This means that searching for vic-




Fig. 1. Three-Agent Society.

tims is similar to other cognitively fatiguing tasks such as air-traffic control. It is likely that a human operator will miss victims due to a failure of attention. During field work with the Hillsborough County Fire Rescue and Tampa Fire Rescue departments, the members of the technical rescue teams specifically said they would prefer to “joystick” the robot rather than examine the image for signs of survivors. The rescue workers themselves believe that perception is the hardest, most fatiguing part of their job. Fortunately, the same perceptual cues used by rescue workers are directly perceivable as Gibsonian affordances [19]: heat, motion, color, and sound.

The second justification for having the robot take the initiative in sensing rather than navigation is that most robot navigation is currently performed using range data extracted from sonars. Indeed, many micro-rovers only have a camera and sonars; they do not have any other method to extract range to obstacles except sonar. However, since the micro-rover is operating in confined spaces, the proximity of walls, obstacles, etc., are within the sonar “dead zone.” It is our experience that sonars are largely blind in realistic USAR spaces. Furthermore, the sonar transducers are exposed and vulnerable to mud, water, or other damage. Even if the robot was capable of semi-autonomous navigation initially, it is reasonable to expect that the sensors, and thereby navigational autonomy, would degrade over time.

This section describes the three-agent society approach, the robots used for the competition, and details the vision agent used to support perceptual search.

#### A. Three Agent Society

The mixed-initiative approach taken in this paper is based on the three-agent society work of Rogers and Murphy[5], which considered human-robot interactions. That work extends the classic two-agent approach to teleoperation, where the human supervisor at the local workstation formed the *Cognitive Agent* and the remote robots were the *Physically-Situated Agents*. Rogers and Murphy suggested that a third, middleware agent was needed to facilitate transactions between the Cognitive and Situated Agents, given the differences in cognitive ability, representation, and contextual knowledge. This third software agent was



Fig. 2. The operator station setup during the AAI conference. Mark Micire and Brian Minten are shown operating Taylor and Fontana, respectively.

dubbed the *Intelligent Assistant Agent* and can be thought of residing with the local workstation or where ever the user interface is generated.

The Intelligent Assistant Agent consists of a number of sub-agents, each a specialist in some aspect of the robot control. This work has produced a basic sub-agent for the Intelligent Assistant Agent, called the *vision agent*, which takes the initiative in examining the perceptual data and displaying the results to the operator. It fuses streams of concurrent sensor data in order to cue the operator to possible victims.

The overall layout of the ideal three agent society is shown in Fig. 1. In that society, the operator communicates with one or more of the remote, physically-situated agents as needed. The Intelligent Assistant Agent facilitates communication as well as handles the representation of data coming from the robots.

The agent society for the AAI Mobile Robot Competition consisted of two remote heterogeneous robots (described in more detail below), two operators, and an intelligent assistant assisting the larger, more challenging robots. Each robot was controlled by an individual operator from dedicated terminals. The two operators were located side-by-side so they could communicate directly (see Fig. 2) and look over and share views from their robots. The two Cognitive Agents performed classical deliberative AI; they cooperated and shared their grasp of the context and did allocation and problem solving. This allowed one Cognitive Agent to ask that a Physically-Situated Agent be re-tasked to confirm a possible detected victim or to aid with a navigation problem. This also allowed the Cognitive Agents to fuse the results of distributed sensing.

### B. Physically-Situated Agents

Physically-Situated Agents for the USAR domain are remote robots. It is assumed that due to size and power constraints that they will be limited in on-board computational power and carry only a few sensors. As such, these agents will be able to function reactively, under a Reactive



Fig. 3. USF USAR team: Fontana and Taylor.

Paradigm, with low computational complexity perceptual processing.

The Physically-situated agents for the competition consisted of two outdoor robots, as seen in Figure 3:

1. *Fontana*, a RWI ATRV wheeled vehicle with a footprint of 102 by 80 cm. Fontana's sensor suite consisted of twelve sonar sensors, color NTSC video, and Indigo Alpha miniature FLIR video camera. Inside, a dual Pentium 300 motherboard, 64 MB of Ram, a 6 GB hard drive, and two matrox meteor framgrabbers provided computing power for the platform. Redhat Linux 6.1 provided the OS support along with an installation of the RWI CORBA-based software interface, Mobility.
2. *Taylor*, a customized RWI Urban with a footprint of 62 by 50 cm. Urbans are polymorphic, or shape-shifting, robots capable of using its front flippers to climb stairs or rubble and to lift the main compartment and sensor suite up (similar to a small animal standing on its hind legs). The Urban had been customized with the addition of a second camera positioned upward toward the rear of the body. This allows video navigation while standing upright. Sensory input consisted of thirteen sonar sensors, NTSC video, magnetic compass, and an accelerometer. The computing power consists of a 240Mhz Intel based motherboard with 64MB of RAM and build in Ethernet. Video from the on-board cameras was digitized with an bttv848 based PCI framegrabber. A Redhat 6.1 Linux system was installed on a 160Mb Silicon hard disk along with RWI's Mobility software.

Fontana and Taylor are part of a tactical mobile robot marsupial team [20], where Fontana carries three Urbans, but at the time of data collection the transport mechanism for the team was still under construction at the time of the competition.

Each operator controlled one robot through the use of dedicated operator control units (OCU). The OCU used for Taylor was designed and implemented by RWI and consisted of a proprietary interface using Mobility. It consists

of a laptop, battery pack, and a wireless Ethernet in a carrying case. The OCU provided the user with a camera view, battery power levels, speed, inclinometer, compass, pose, and sonar. Taylor was controlled by moving the cross-hairs on the camera view using the touch pad. Her flippers were controlled similarly by cross-hairs on the power level display.

The OCU used for Fontana was an in-house design using X-windows display widgets. It attempted to simplify the teleoperation control challenges imposed by the RWI OCU as well as permit the display of results from the Intelligent Assistant Agent. It consists of a laptop and wireless Ethernet. The operator was provided with sonar, an intelligent operator system, and keyboard control. The arrow keys were used to drive Fontana.


The robots communicated with the OCUs for the operators through a pair of Breezecom wireless Ethernet nodes and a base unit connected to a hub. This wireless network operated in the 2.4Ghz band and provided up to 2Mbps data throughput to the robots. All the Intelligent Assistant Agent processing resided on Fontana, not on the OCU.

### C. Intelligent Assistant Agent

The Intelligent Assistant Agent for this work consisted of one sub-agent, a vision agent to aid with victim detection. The vision agent performed *cueing* and *behavioral fusion* based on the output of four concurrent detection algorithms: *motion*, *skin color*, *color difference*, and *thermal region*. The user interface displayed the extraction results from each applicable algorithm, highlighted in color. Thus, the operator had visual cues for victims. In addition, the operator was notified by beeping when the agent had sufficient confidence in the presence of a victim based on fusion of multiple detectors. Fig. 4 shows two of the detectors. All active detectors were tiled side by side along with an un-augmented view from the video camera. No screen-dumps of the actual display were taken.

Each of the four detection algorithms extracted a single feature that afforded the presence of a victim. The detection algorithms were written in C. Sufficient size of a feature or co-location of cues was interpreted by the vision agent as a possible victim. The operator was notified of the results of each algorithm through a continuously updating display, written in QT. The operator was also given the ability to switch on and off any of the algorithms. Figure III-C shows the composition of detectors for the vision agent, with the order complexity.  $m$ ,  $n$  are the dimensions of the image array, in this case 320 by 240. Each detector exhibited a low order computation for runtime execution, though initialization might take longer.

Three algorithms were used in conjunction with the color camera: Skin Color detection, Motion detection, and a Color Difference algorithm. Two algorithms were used with the FLIR: Heat detection and Motion detection. Each ran independently of the other, and each returned a list of possible candidate victims in the image. The vision agent then took these lists and matched the location of the candidate in the images to inform the operator if a possible



a.



b.

Fig. 4. Output of two of victim detectors used by Intelligent Assistant Agent: a) heat (output of Indigo Alpha FLIR) and b) motion (output of either FLIR or video camera).

| Robot   | Vision Sensors                  | Detectors               | Complexity       |
|---------|---------------------------------|-------------------------|------------------|
| Fontana | color camera                    | motion                  | $O(4m \times n)$ |
|         |                                 | skin color              | $O(m \times n)$  |
|         |                                 | color difference        | $O(m \times n)$  |
|         | IR                              | heat                    | $O(m \times n)$  |
|         |                                 | motion                  | $O(4m \times n)$ |
| Taylor  | <i>b&amp;w camera (forward)</i> | <i>motion</i>           | $O(m \times n)$  |
|         | <i>color camera (back)</i>      | <i>motion</i>           | $O(4m \times n)$ |
|         |                                 | <i>skin color</i>       | $O(m \times n)$  |
|         |                                 | <i>color difference</i> | $O(m \times n)$  |

Fig. 5. Two instances of the Intelligent Assistant Vision Agent. The agent on Taylor was not implemented at the time of the competition.

victim is found.

Ideally, the robot team would have had two instances of the vision agent: one for Fontana and one for Taylor. However, the agent was not migrated to Taylor for the competition (though it was added later) due to difficulties with the proprietary RWI OCU. Also, the black and white camera was the primary sensor for the competition; the short duration (20-30 minutes) of a run made it impractical to shape shift Taylor to use the rear color camera. The image compression from Taylor rendered the motion detection from the black and white camera useless. Therefore, it was not practical to use an intelligent vision agent for Taylor.

### C.1 Motivation for Choice of Detectors

The vision agent can support any number of vision algorithms. The choice of these four algorithms was based on our understanding of the cues used by USAR technical rescue personnel and on the use of low computational *affordances*. A living victim will always radiate heat, which can be measured by a thermal imager. Heat is a primary affordance of victims in USAR. The only heat sources in the rubble should be people or animals or indicators of a fire. The region properties of the heat source (e.g., does it look like the profile of a human) is not as critical. The rubble may insulate large portions of the victim, causing a more sophisticated (and computationally expensive) segmentation and matching algorithm to report false negatives (nothing is present).

Another affordance of a victim is color. In a collapsed structure, there is a grey dust from sheet rock or ceiling tiles which covers everything, and as such, most color is washed from the image. In a structure that has experienced a fire, the debris is similarly covered in black soot. Therefore rescuers look for is anything that is distinguishable from that grey dust. For instance, a victim may move and shake of some of the dust, or there may be distinguishable colors such as blood. Previous work[12] used a color detector tied to a specific color; in that case, international orange since that was the color of safety vests and helmets. Rather than use a large number of color segmentation algorithms, each targeted at a different color, this work used an algorithm that attempted to detected colors that were different and distinct from the background debris. A skin detector algorithm, which was also included

Notice that the two previous victim detection affordance apply regardless if the victim were conscious or not. If the victim is conscious, an easily detectable affordance is motion, e.g., moving arms and waving for attention as the robot comes near. The specifics of each detector is given below, along with the fusion mechanism.

### D. Heat and Skin Color Detectors

The heat and skin color detectors relied on region segmentation. A skin color detector was included because it was available. It used the spherical coordinate transform(SCT) to operate under the unconstrained USAR conditions. SCT has been used for mobile robots and has been

shown to perform well in unstructured lighting conditions such as outdoors.[21] Various skin colors segment very similarly over SCT space, thus multiple definitions of skin color are not needed[22].

The skin color detector required a range of values for the color. These values were trained using a simple two dimensional Gaussian histogram over SCT space for 10 sample images of skin color. Once the values were trained, when the algorithm ran in real time, the five largest connected areas of the skin color were then displayed to the operator as well as sent to the vision agent.

Heat detection was performed similarly for the FLIR camera. Since the image reports heat as a value between 0 and 255, the algorithm would find the five largest connected blobs of heat in the image above a threshold. The threshold was the value corresponding to the ambient temperature at the site.

### E. Motion Detector

Motion detection was done in the same manner for both cameras. First, the images are assumed to be coming in as grey scale. Second, it is assumed that the robot is not going to be moving during the motion detection routine. Two consecutive images are taken, and a difference operator is run on these images. Next, this new binary image is taken and applied to an evidence grid using a growth operator. This works as follows: if a pixel in the binary image is determined to be motion, then that pixel in the evidence grid is increased by 3(to a maximum of 255), and if a pixel in the binary image was not determined to be motion, 1 is subtracted from that location in the evidence grid(to a minimum of 0). Once this has been done on ten consecutive sets of images, the evidence grid is thresholded with to create a new binary image representing where motion has been continuously occurring. The five largest blobs of motion from this image are then sent to the vision agent for processing as well as being displayed to the user. This is continued for every new set of images once ten have gone by, until the motion detection is turned off by the operator. This method helps to remove noise as well as motion occurring across the visual field, such as the daughter robots moving around the area. Since victims will tend to be moving in the same areas, this algorithm should find only victims moving, and not other objects.

### F. Color Difference Detector

The last algorithm used for victim detection was a novel approach created after discussion with actual rescuers how they visually recognized victims in a disaster.

To this end, a color image is taken upon initial entering of the structure. It is converted to the SCT space, and is then quantized into 49 bins. The number of pixels in the image belonging to each bin is counted up, and any bin that is over 10% of the image size is added to the list of bins to not segment, in effect, these bins form the image *background* and the presence of pixels in other bins form the *foreground*. This initialization was controlled by the operator, and can be run any number of times, such as

when the robot enters a new search area. Once initialized, an image is segmented in SCT space, and any pixels not belonging to marked bins are set as possible victim pixels. Then the five largest connected regions of marked pixels are displayed and sent to the vision agent.

### F.1 Behavioral Sensor Fusion

While the operator could identify a victim from any of the cues displayed by the vision, the vision agent itself can more aggressively alert the operator to presence of a victim using behavioral sensor fusion[23]. The behavioral sensor fusion algorithm assumed that if two or more detectors were operating on the robot and two or more reported cues in the same location, then that was a sufficient positive.

After the detectors have run and updated their displays, they send the  $(x, y)$  centers of the connected regions to the vision agent. The vision agent translates the image location from the detectors operating on the FLIR thermal images to the video image coordinates.

These centers are then compared and if any two from different algorithms are within a certain threshold distance of each other, then a signal (an audible beep) is sent to the operator telling them to check more closely in this area.

In the case of Fontana, the video camera and FLIR were mounted approximately at the same height alignment, so only the relative horizontal distance had to be considered. To accomplish this, the focal length of each camera was set, and knowing the field of view and the actual distance apart, a constant value was computed for the horizontal distance offset in pixels, and thus the two cameras can match up. While not exact, this value is a good approximation of the actual location and appeared to be sufficient since victim cues covered a sizable portion of the image and exact location was not needed.

### G. Collaborative Teleoperation

A third characteristic of our approach to controlling robots in USAR environments is an expectation of collaborative teleoperation. Collaboration teleoperation is the use of a secondary robot to provide the operator with an external view of the primary robot. The external view is intended to help performance (e.g., navigate through a tight spot), diagnose and recover from a problem (e.g., high-centering of the robot, loss of a track, etc.) or cooperate (e.g., make an opening for the target robot to go through).

In the instance of USAR, a second robot can help in a variety of ways. In terms of victim detection, the second robot can provide confirmation of a victim through adding an additional viewpoint and/or sensor modality. The secondary robot can assist the primary robot by providing external lighting of a scene.

Our experience with fielding USAR robots at the Hillsborough County Fire Rescue (Tampa, Florida), SRDR (Miami, Florida), and the Montgomery County Fire Training Academy (Rockville, Maryland) training facilities is that it is very easy to flip or high center a robot in rubble. While it is easy to get the robot into trouble, it is difficult to diagnose the problem or construct a solution using only limited




Fig. 6. Layout of the NIST Standard Test Course for USAR. Upper left box is the red section. Upper right box is the orange section. And the bottom rectangle is the yellow section.

proprioceptive sensors (tilt, pose) or an exteroceptive view of the environment. What is needed is an exproprioceptive view, where the robot’s pose is known relative to the external environment. Since the robot cannot see itself, the viewpoint of another agent is helpful.

Collaborative teleoperation is a function that can be supported by an Intelligent Assistant Agent. The agent can display appropriate video feeds from both the primary and secondary robots. It could even request and negotiate for the re-tasking of a secondary robot. The agent used reported in this article did not provide these services, but these directions are being explored.

## IV. DATA AND RESULTS

The mixed-initiative system with the vision Intelligent Assistant Agent were tested at the 2000 AAI Mobile Robot Competition USAR event.

### A. 2000 AAI Mobile Robot Competition

The NIST Standard Test Course for USAR consisted of three sections, each providing a different level of difficulty in order to accommodate a range of robot platforms (see Figure 6). The victims were mannequins and dolls. “Live” victims were adjacent to a heating pad to generate heat. The course also contained parts of mannequins, which represented dismembered body parts. These body parts were to act as a source of false positives (that a victim was there when in fact it wasn’t). Some victims had a metronome type of device to create a motion signature and a tape recorder playing simulated victims’ cries for help.

The easiest section, Yellow, contained mainly hallways, blinds, and covey holes to search through. This was a little more difficult than the Office Delivery events held by the AAI Mobile Robot Competition in past years. The course could be traversed by any indoor research type of robot. Furthermore the victims in the Yellow Section were

all surface victims, that is, fully visible and not trapped. As noted in [24], the Yellow Section was not a realistic USAR setting.

The USF robots did not search the Yellow Section, concentrating on the more realistic Orange and Red Sections instead. The Orange and Red Sections required robot platforms capable of negotiating rubble. The intermediate Orange Section provided more challenge with the addition of a maze, negative obstacles, and a second level that was reachable by stairs or ramp. The victims were also all surface victims.

The Red Section was intended to be the most difficult. It contained piles of rubble and dropped floorboards that represented a pancake-like structure. The pancake surfaces were lightly supported, causing a secondary collapse if a robot platform traveled on it. Victims in the Red Section were in either pancake type of voids or entombed.

The USF team participated in three exhibition runs, a preliminary competition event, and a final event. The two competition events and one of the exhibition runs were completely recorded on video by USF. A second exhibition run was partially recorded. The placement of victims and the layout of the arena was changed by the judges between runs. The results fall into two main areas: victim detection and collaborative teleoperation.

The previously published scoring scheme for the AAAI competition was changed on site from a combination of quantitative performance and qualitative software metrics to a purely qualitative rating scheme in order to reward full autonomy approaches. Under the new scheme, USF received third place out of four teams participating. The other three teams operated only in the Yellow Section, while USF operated only in the Orange and Red Sections. The USF team found the most victims in both the preliminary and final rounds, 4 and 6 respectively. The data reported here is based on analysis of tape recordings made by USF, since the judges did not record the number of victims identified by each team. The exact number of victims in each section is also not known; the judges did not record that information.

### B. Narrative

The same exploration strategy was used in each run. The USF robots were placed on the boundary of the NIST Standard USAR Test Course. Each robot was controlled by a single operator. The agile Urban robot, Taylor, served as “point-man”, entering a new area first. Taylor would first enter the Orange Section maze and explore, then come back to provide a collaborative teleoperation view for Fontana to enter the maze and reach the open ramp area. Structural cross members and a door prevented Fontana from entering the main Orange area, but Taylor could crawl under the cross members and push open the door for Fontana. At that point, Fontana would be able to navigate unassisted and search the ground portion while Taylor would climb up the ramp, search the second story, then exit down the stairs. Together the two robots would exit on the far side of the Orange Section and cooperatively enter the Red Sec-

tion. The Red Section terrain was more challenging for Taylor, and Taylor would require Fontana to provide external views.

### C. Victim Detection Results

| Run    | Time (min)         | Total Vict Found | Robot   | Vict Found Solo | Vict Found Joint |
|--------|--------------------|------------------|---------|-----------------|------------------|
| Prelim | 32                 | 4                | Fontana | 1               | 2                |
|        |                    |                  | Taylor  | 1               | 2                |
| Final  | 30                 | 7                | Fontana | 1               | 3                |
|        |                    |                  | Taylor  | 3               | 3                |
| Exh #3 | partially recorded | 6                | Fontana | 0               | 3                |
|        |                    |                  | Taylor  | 3               | 3                |
| Exh #4 | 28                 | 6                | Fontana | 1               | 4                |
|        |                    |                  | Taylor  | 1               | 4                |

TABLE I

RESULTS OF VICTIM DETECTION IN THE FOUR RECORDED RUNS.

Table I shows the performance of the USF heterogeneous team for victim detection. The data shows that the majority of victims were detected jointly, or *cooperatively*. This was due to two reasons. First, the more agile Urban robot, Taylor, could quickly enter an area and find a sign of survivor, but often could not confirm that the sign was a victim using on a black and white camera with a narrow field of view. Second, Taylor could not confirm whether a victim was “living” or “dead” if the victim was not moving; this had to be disambiguated with the FLIR thermal imager on Fontana. Fontana found fewer victims by herself, or solo, because she was larger and less agile and had to slowly navigate the more restricted spaces of the Orange and Red Sections.

The detection did not report any false positives (label a non-victim as a victim). It did have at least one false negative (miss a victim) on every run. This mannequin was on a raised platform and was out of the view of the cameras; neither robot had the ability to pan cameras up.

The performance of the vision agent was difficult to quantitatively assess. The audio portion of the cueing had to be turned off. This was because the NIST Standard Test Course has Plexiglas panels to facilitate judge and spectator viewing. AAAI permitted spectators to ring the sections during the competition. Between the low height of walls and the Plexiglas, these spectators were visible and produced color, motion, and IR signatures. These spurious signals were present even when the USF robots were facing interior walls due to the ability to see through the exterior walls in other sections. As a result, USF had to turn off automatic victim notification through audio and rely strictly on color highlighting in the Intelligent Assistant Agent display windows.

### D. Collaborative Teleoperation Results

Table II shows the amount of time each robot spent acting as the secondary robot, providing an external view for the primary robot. The table indicates that both robots

| Run           | Time (minutes) | Robot   | Seconds Spotting |
|---------------|----------------|---------|------------------|
| Preliminary   | 32             | Fontana | 25               |
|               |                | Taylor  | 222              |
| Final         | 30             | Fontana | 45               |
|               |                | Taylor  | 160              |
| Exhibition #4 | 28             | Fontana | 10               |
|               |                | Taylor  | 250              |

TABLE II

RESULTS OF COLLABORATIVE TELEOPERATION DURING THE FOUR RECORDED RUNS.

needed collaborative views at some point during each of the recorded runs. As expected, the agile Urban, Taylor, needed much less help navigating, while the larger ATRV, Fontana, had to frequently rely on viewpoints from the Urban to assist the operator in guiding it through confined space.

The major areas of collaboration were:

- *Navigating confined space.* The majority of the time spent by Taylor spotting was to help Fontana navigate through confined space. Fig. 7 shows one such instance, where the tight clearance from the narrow hall was inside the minimum operating range of Fontana’s sonar. This meant that the sonars could not reliably report the robot’s position relative to the walls. The camera faced forward, so Fontana’s operator could not see when the robot was about to collide with a wall or a supporting bracket.
- *Cooperatively working together.* The two robots often had to watch each in order to cooperate. As shown in Fig. IV-D, Fontana could not enter one side of the Orange course unless it could open a door. Taylor was able to enter the area by going under a set of cross braces, then push open the door. Taylor then had to back up to get out of Fontana’s way as she entered, but if she backed up too quickly, the door would jam Fontana.
- *Recovery.* The majority of time when Taylor was being spotted by Fontana was due to Taylor getting trapped. This most often happened in the Red Section, where Taylor was vulnerable to getting tangled in the chicken wire and to falling off the pancaked platforms. See Fig. 9.

## V. DISCUSSION

While the data from the AAI Competition is incomplete, the results support the efficacy of heterogeneous teams for USAR. The efficacy of Three-Agent Society organization is more subjective. The results also show that although the conventional wisdom of rescue workers that perception is harder than navigation, teleoperated navigation is very hard.

### A. Efficacy of Heterogeneous Teams

The case study and data collected show the efficacy of heterogeneous teams for USAR. In this case, the robots



Fig. 7. Collaboratively navigating the maze in the Orange Section.



Fig. 8. Taylor opening a door for Fontana.




Fig. 9. Taylor successfully climbing on a narrow ledge using Fontana as the secondary robot. (Fontana is not shown.)

were heterogeneous both in navigation abilities and sensor suites. Heterogeneity led to a 133% improvement on the major robot task of victim detection, and facilitated the implicit task of robot survival.

Only one robot (Fontana) was capable of identifying a live victim. However, that robot was incapable of navigating in confined spaces without assistance, or would have navigated very slowly. If the robots had operated independently only 9 victims total over all recorded runs would have been detected. But by cooperating on identification, an additional 12 victims were confirmed, increasing the rate of detection by 209%. It should be emphasized that Fontana, the large robot, would have been unable to do more than circle the perimeter of the Orange Section if she hadn't been able to use collaborative teleoperation.

The cooperation was handled by the Cognitive Agents (human operators) who were sitting next to each other and could request help. We are currently exploring an Intelligent Assistant sub-agent for re-tasking robots for cooperation based on proximity and sensor suite. This activity is suitable for a software agent for several reasons. In the foreseeable future, USAR technical rescue will have more robots and operators, so it is unlikely that a human will be able to maintain an understanding of which robots are nearby, have appropriate complementary resources, and can be interrupted from their current activity. This line of research is both a resource allocation problem and a distributed sensing problem.

The results also showed a number of cases where the robots would have been trapped and considered to have "died in place." This is an important practical issue. A robot which has died in place is one less resource that can be used. More importantly, it is a resource that may be considered valuable enough to rescue (Taylor cost close to \$40,000 USD), thereby diverting rescue efforts from victim detection and extrication to robot rescue.

### B. Efficacy of the Three-Agent Society

The data gathered at the AAI Competition does not directly shed any light on the efficacy of a Three-Agent Society. There was no competing mixed-initiative system available to compare with, and it was beyond the scope of the competition effort to collect human-robot interface data.

The supporting side for the three agent society is anecdotal. The three robot operators were convinced it was much easier to control and identify victims than purely teleoperated versions. They also stated that the intermediate, non-fused results from the detector functions were not visually overwhelming.

One disappointing aspect of the Intelligent Assistant Agent was that the audible alert had to be turned off. This prevented us from determining how helpful it would be in a more realistic environment where large numbers of spectators would not be visible. As noted in [24], access to the area of collapse, the hot zone, and the surrounding buffer zone, the warm zone, is strictly limited at a rescue site. Therefore, it is expected that the vision agent, including

| Run    | Robot   | Corr. | Coll. | Pose |
|--------|---------|-------|-------|------|
| Prelim | Fontana | 90    | 2     |      |
|        | Taylor  | 70    | 3     | 1    |
| Final  | Fontana | 45    | 3     |      |
|        | Taylor  | 75    | 1     | 1    |
| Exh #2 | Fontana | n/a   | n/a   | n/a  |
|        | Taylor  | n/a   | n/a   | 2    |
| Exh #3 | Fontana | 64    | 0     |      |
|        | Taylor  | 38    | 3     | 1    |

TABLE III  
SUMMARY OF OPERATOR ERRORS DURING NAVIGATION: CORRECTIONS (CORR), COLLISIONS (COLL), AND POSE ERRORS (POSE). ONLY TAYLOR WAS POLYMORPHIC AND CAPABLE OF ENCOUNTERING POSE ERRORS.

behavioral sensor fusion, will be useful in practice.

### C. Teleoperated Navigation and Error Taxonomy

While a basic tenet of the USF mixed-initiative approach is that perceptual search is more important to automate than navigation, data collected indicates that teleoperated navigation is not easy for an operator either. Table V-C reports on the analysis done by examining the recorded runs. The analysis identified three major categories of navigational errors: *course corrections*, *collisions*, and *pose errors*.

The largest source of problems were course corrections. These were situations where the operator clearly overshot a desired location and had to be iteratively repositioned. A common case was repositioning after a collision with a wall or other obstacle. Due to human error, the robot would often go through a repetitive cycle of over-corrections before resuming more smooth control. Course corrections waste valuable time and add to the cognitive fatigue and frustration of the operator.

Another problem was collisions. These were often the result of attempting to navigate quickly (high velocity) and occasionally exceeding the communications time lag. One solution that is now being implemented is *guarded motion*, which uses sonar readings to prevent the robot from hitting obstacles. However, this is not a total solution because as noted previously the sonars often do not function in confined space and there are significant blind spots in the coverage on an Urban.

Fig. 10 shows an example of pose errors. Pose errors were when the polymorphic Urban robot, Taylor, were in a non-optimal position. The two most common non-optimal positions were being in the wrong configuration for rapid movement and being in the wrong configuration for climbing or traversing rubble. Pose errors are interesting because the RWI OCU does report pose information. However, we believe that the cognitive attention required to look at the video and navigate the robot prevent the operator from paying much attention to the pose information. Current work is exploring an intelligent assistant agent sub-agent to maintain the best pose for the speed and terrain.



Fig. 10. Taylor trying to rapidly translate while in an unfavorable pose.

#### D. Relevance to Other Domains

The data and lessons learned reported in this article can be applied to other domains, most especially military operations in urban terrain (MOUT). These types of operations, described in [25], will often involve teams of heterogeneous robots operating in confined space under a mixed-initiative control scheme.

The results suggest that no matter how agile a robot platform is, it will get trapped or stuck due to the intrinsic sensor limitations. Even if the robot is operating autonomously, the lack of sensors will preclude it from autonomous recovery in most cases. Therefore, an operator may have to assist. However, the operator is unlikely to do much better in diagnosis and recovery using only the on-board sensors. The results suggest that the viewpoint from a secondary robot enables the operator to more quickly identify the problem and solution.

The experience with the NIST test bed also suggests that robots will often have to negotiate risky terrain, such as ledges or structures where the support cannot be verified. The use of a secondary robot can minimize this risk of falling off and allow the operator to be alert to signs of an incipient collapse.

## VI. SUMMARY

This article describes a mixed-initiative system for Urban Search and Rescue, its implementation on a team of heterogeneous robots, and the data collected on the team's performance at the 2000 AAI Mobile Robot Competition. The objective of the competition was to find the most victims in a fixed period of time. The robot team consisted of a RWI ATRV vehicle with video and FLIR and a RWI shape-shifting Urban with video. Each robot had one teleoperator. The two robots cooperated in both victim detection and navigation. Cooperation in victim detection permitted the team to classify victims as dead or living with sufficient certainty, while cooperation in navigation permitted the robots to enter areas that it could not enter without the additional external viewpoint from the secondary robot or without manipulation of the environment

by the secondary robot.

The mixed-initiative control scheme used a three-agent society architecture, where the remote robots were Physically-Situated agents, teleoperators were Cognitive Agents, and an Intelligent Assistant Agent provided perceptual processing. The Intelligent Assistant Agent aided with perception rather than navigation based on results of work with the Hillsborough County Fire Rescue Department on user preferences. The agent ran four victim detector algorithms on two sensors to identify affordances of victims (heat, motion, skin color, and color difference). Each of these algorithms exploited directly perceivable affordances for victims, resulting in low ( $O(m \times n)$ ) order complexity implementation of perception. The agent displayed the output of each detector to the teleoperator with augmentation showing the regions of possible victims, serving the cue the operator's attention. The vision agent also fused data from multiple detectors and modalities to generate an audible alert when confidence in the presence of a victim was high.

The USF team found more than four times as many surface and non-surface victims in confined spaces as did any other competitor. All other competitors were all fully-autonomous systems limited to searching for surface victims in the open office areas of the NIST arena. The data shows that by cooperating heterogeneously the USF team found 209% more victims than by the individual robots working independently.

The article also provides a taxonomy and data on the types of errors that occur during teleoperated navigation in confined spaces. There are three distinct types of operator errors, *iterative course corrections* and *collisions* which apply to any robot platform, and *pose errors*, which apply to polymorphic platforms, such as an Urban. Of the three, iterative course corrections were the most prevalent, averaging 2.21 per minute for the ATRV and 2.03 per minute for the Urban, suggesting that a large block of the operator's time is spent in frustrating fine positioning actions. Although pose errors occurred less frequently, between 1 and 2 per run, they were serious and this suggests that this will be an important consideration for fielding polymorphic robots. An Intelligent Assistant sub-agent for navigation to complement the vision sub-agent is clearly called for.

## ACKNOWLEDGMENTS

Portions of the work presented in this article were supported by DARPA Grant F04701-99-C-0317, donations from SAIC, and a travel scholarship from the American Association for Artificial Intelligence (AAAI). The authors would like to thank Ron Rogers and Jeff Hewitt of the Hillsborough County Fire Rescue Department and LTC John Blitch for arranging for field testing with fire rescue teams, and Mike Larson of IRcameras.com for the loan of the Indigo Alpha FLIR.

## REFERENCES

- [1] J. G. Blitch and R. Mauer, "Knobsar: a knowledge based system prototype for robot assisted urban search and rescue," *Simulation*, vol. 66, pp. 375-391, 2000.

- [2] J. G. Blicht, "Artificial intelligence technologies for robot assisted urban search and rescue," *Expert Systems with Applications*, vol. 11, no. 2, pp. 109–124, 1996.
- [3] J. Casper, R.R. Murphy, and M. Micire, "Issues in intelligent robots for search and rescue," in *SPIE Ground Vehicle Technology II*, 2000.
- [4] R.R. Murphy, "Marsupial and shape-shifting robots for urban search and rescue," *IEEE Intelligent Systems*, vol. 15, no. 3, pp. 14–19, 2000.
- [5] R.R. Murphy and E. Rogers, "Cooperative assistance for remote robot supervision," *Presence*, vol. 5, no. 2, pp. 224–240.
- [6] A. Kobayashi and K. Nakamura, "Rescue robot for fire hazards," in *Proceedings of the 1983 International Conference on Advanced Robotics*, 1983, pp. 91–98.
- [7] A. Castano, W.M. Shen, and P. Will, "Conro: towards deployable robots with inter-robot metamorphic capabilities," *Autonomous Robots*, vol. 8, no. 3, pp. 309–324, 2000.
- [8] S. Hirose, "Snake, walking and group robots for super mechano-system," in *IEEE SMC'99 Conference Proceedings*, 1999, pp. 129–133.
- [9] R. Masuda, T. Oinuma, and A. Muramatsu, "Multi-sensor control system for rescue robot," in *1996 IEEE/SICE/RSJ International Conference on Multisensor Fusion and Integration for Intelligent Systems*, 1996, pp. 381–387.
- [10] R.M. Voyles, "Terminatorbot: a robot with dual-use arms for manipulation and locomotion," in *Proceedings 2000 ICRA*, 2000, pp. 61–66.
- [11] D. Diaz, T. Flowers, and R. Murphy, "1997 aaai find life on mars competition," in *Robotics 98. Proceedings of the Third ASCE Speciality Conference on Robotics for Challenging Environments*, 1998, pp. 237–243.
- [12] R. Murphy and J. Sprouse, "Strategies for searching an area with semi-autonomous mobile robots," in *2nd Conference on Robotics for Challenging Environments*, 1996, pp. 15–21.
- [13] S. Tadokoro, H. Kitano, T. Takahashi, I. Noda, H. Matsubara, A. Hinjoh, T. Koto, I. Takeuchi, H. Takahashi, F. Matsuno, M. Hatayama, J. Nobe, and S. Shimada, "The robocup-rescue project: a robotic approach to the disaster mitigation problem," in *IEEE International Conference on Robotics and Automation*, 2000, vol. 4, pp. 4089–4094.
- [14] Z. Zhu, K.D. Rajasekar, E.M. Riseman, and A.R. Hanson, "Panoramic virtual stereo vision of cooperative mobile robots for localizing 3d moving objects," in *Proceedings IEEE Workshop on Omnidirectional Vision*, 2000, pp. 29–36.
- [15] R. Murphy, M. Ausmus, M. Bugajska, T. Ellis, T. Johnson, N. Kelley, J. Kiefer, and L. Pollock, "Marsupial-like mobile robot societies," in *Agents 99*, 1999, pp. 364–365.
- [16] J.S. Jennings, G. Whelan, and W.F. Evans, "Cooperative search and rescue with a team of mobile robots," in *1997 8th International Conference on Advanced Robotics*, 1997, pp. 193–200.
- [17] D. Kortenkamp and G. Dorais, *Tutorial: Designing Human Centered Autonomous Agents*, 2000 PRICAI Pacific Rim International Conference on Artificial Intelligence, <http://www.traclabs.com/korten/tutorial.html>, 2000.
- [18] Y. Horiguchi, T. Sawaragi, and G. Akashi, "Naturalistic human-robot collaboration based upon mixed-initiative interactions in teleoperating environment," in *IEEE International Conference on Systems, Man and Cybernetics*, 2000, vol. 2, pp. 876–881.
- [19] J.J. Gibson, *The Ecological Approach to Visual Perception*, Houghton Mifflin, Boston MA, 1979.
- [20] J. Spofford, J. Blicht, W. Klarquist, and R. Murphy, "Vision-guided heterogeneous mobile robot docking," in *PIE Conference on Sensor Fusion and Decentralized Control in Robotic Systems II*, 1999, pp. 112–121.
- [21] J. Hyams, M. Powell, and R. Murphy, "Cooperative navigation of micro-rovers using color segmentation," *Autonomous Robots*, vol. 9, no. 1, pp. 7–16, 2000.
- [22] Scott E. C. Umbaugh, *Computer Vision and Image Processing*, Prentice-Hall, Englewood Cliffs, NJ, 1998.
- [23] R.R. Murphy, *Introduction to AI Robotics*, MIT Press, Cambridge, MA, 2000.
- [24] R.R. Murphy, J. Casper, M. Micire, and J. Hyams, "Assessment of the nist standard test bed for urban search and rescue," in *NIST Workshop on Performance Metrics for Intelligent Systems*, 2000.
- [25] E. Krotkov and J. Blicht, "The defense advanced research project agency (darpa) tactical mobile robotics program," *International Journal of Robotics Research*, vol. 18, no. 7, pp. 769–776, 1999.