Managing Risk in Disaster Scenarios with Autonomous Robots

Daniel P. Stormont and Vicki H. Allan
Department of Computer Science
Utah State University
Logan, Utah 84322–4205
Email: daniel.stormont@aggiemail.usu.edu; vicki.allan@usu.edu

***Abstract*— Disaster areas are one of the most challenging environments faced by mankind. Uncertainty, hazards, and limited availability of rescuers all impact the ability to save lives. Prepositioned autonomous rescue robots offer promise in assisting the first responders to a disaster site, but there is a challenge to using robots in hazardous environments: numerous studies have shown that human rescuers lack trust in fully autonomous systems.**

This paper introduces the aspects of disaster areas that make them so challenging. The use of robots as a risk management tool for human rescuers is introduced. Then some of the factors that limit human trust in robots are addressed – including one of the key factors: reliability. The design of a computer model used to investigate issues of trust and the impact of reliability in a firefighting scenario is discussed and the results are analyzed. Finally, some preliminary conclusions and plans for further work in this area are presented.

***Index Terms*— disaster management, autonomous robot, trust, reliability, simulation**

I. Introduction

Regardless of their cause, human or natural, disaster areas all share a number of characteristics:

- There is uncertainty about the extent and degree of damage.
- The initial response to the disaster is limited to only those local rescue assets that have survived the incident.
- There are many hazards in the area, whose location and nature are unknown.
- There is a high likelihood of trapped victims in the area, whose location and condition are unknown.
- As a counterpart to the limited number of first responders available, studies have shown that the first 72 hours are essential for rescuing victims. The survival rate drops geometrically with time, to nearly zero after 72 hours [1].

One tool becoming available to first responders in a disaster area is robotics. Robots have had some limited use in disaster response, most notably the responses led by the Center for Robot-Assisted Search and Rescue (CRASAR) at the University of South Florida [2]. Fig. 1 shows the robots that responded to the World Trade Center disaster in September 2001. However, rescue robots have still not been widely accepted by human rescuers.

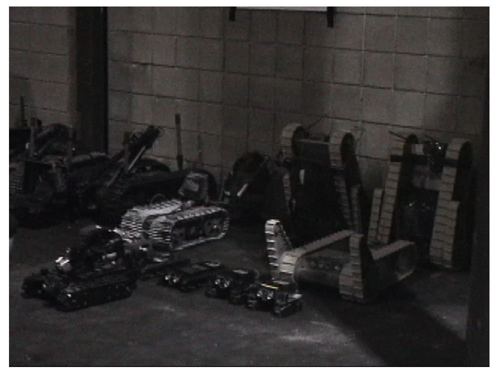

Fig. 1. Robots deployed to the World Trade Center disaster site in 2001. *Photograph by the Center for Robot-Assisted Search and Rescue*

Fig. 2. Robots in use in Iraq: A bomb disposal robot (left) and a UAV (right).*Photographs from U.S. Air Force.*

The situation has been somewhat different in combat areas, where bomb disposal robots and unmanned aerial vehicles (UAVs) have been well accepted, as shown in Fig. 2.

However, even in areas where robots have been accepted as tools to augment the capabilities of humans or to protect humans from unnecessary exposure to dangerous situations, autonomous robots have been rejected. An example is the Special Weapons Observation Reconnaissance Detection System (SWORDS) manufactured by Foster-Miller, pictured in Fig. 3. Even though this robot is operated under human control, the Army has never authorized the use of the weapon on the three SWORDS robots deployed to Iraq [3] and the robot was the center of a widespread controversy when it was claimed that the gun mounted on the robot slewed without any input from the operator [4].

Fig. 3. Foster-Miller SWORDS robots. The middle robot is the type that was deployed to Iraq. *Photograph by US Army.*

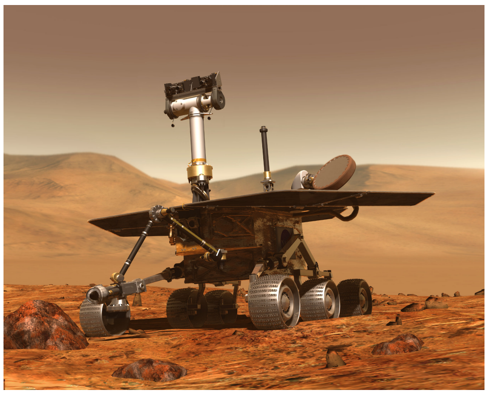

Fig. 4. Mars Exploration Rover – a semi-autonomous exploration robot. *Image courtesy of NASA/JPL-Caltech*

II. Autonomous Robots and Risk Management

Robotics researchers have long pursued the goal of building autonomous robots that could assist humans. Manufacturers were quick to incorporate robots in industrial settings, where they could replace humans on the assembly line performing dull, repetitive, and dangerous tasks – often with lower error and rework rates than the humans they replaced [5]. These systems were automated, but not truly autonomous, since they could only carry out preprogrammed sequences of motions, were not mobile, and reacted to unexpected environmental conditions by shutting down until a human operator could clear the error condition. Autonomy is necessary if a robot is going to interact with a realistic environment and be able to assist humans where they live. A good definition of the characteristics of an *autonomous robot* was provided by Arkin: "An intelligent robot is a machine able to extract information from its environment and use knowledge about its world to move safely in a meaningful and purposeful manner [6]." Autonomous robots have been confined mostly to research laboratories, with a few exceptions: limited functionality consumer robots, like the iRobot® Roomba®, and space probes, like the Mars Exploration Rover shown in Fig. 4, that navigate autonomously with human direction to establish goals for exploration.

While autonomous robots for the most part have not broken free of the laboratory, there has been a great deal of interest in using autonomous robots to manage the risk of disasters. The Japanese Central Disaster Prevention Committee and a number of local Japanese government agencies have been collecting data on historical disasters (mostly earthquakes and the resulting tsunamis) to establish plans for risk mitigation of future disaster events [1] [7] [8]. One of the research projects of the Central Disaster Prevention Committee is the Special Project on Development of Advanced Robots for Disaster Response (DDT Project). The project has nine task forces researching technologies to assist with disaster risk management; including the development of a number of different types of rescue robots, human-robot interfaces, communications aids, infrastructure improvements, and field evaluations. Japan is not alone in performing evaluations of robotics technologies in disaster situations. In the United States, the Center for Robot-Assisted Search and Rescue has been deploying to disaster sites since 2001, as mentioned in the introduction. The U.S. National Institute of Standards and Technology (NIST) has a very active program to evaluate and promote the development of robots for managing disasters; including key roles in annual disaster response exercises and the annual RoboCup Rescue competitions [9]. Obviously, disaster managers see potential for robotic systems to save lives.

III. Trust and Reliability

The primary impediments to the adoption of autonomous robots systems for managing the risk of disaster scenarios are the related elements of human trust in autonomous robots and the reliability of autonomous robots. This section will consider these two elements independently, though it will be apparent that they are interrelated.

A. *Trust*

Trust has been an implicit element of human relations for as long as people (or animals, for that matter) have interacted with each other. Trust is an integral part of cooperative games and negotiations [10]. However, trust as an area of research in its own right can be traced back to the classic collection of papers by Gambetta in 1988 [11]. Gambetta's definition of trust was:

> trust (or, symmetrically, distrust) is a particular level of the subjective probability with which an agent assesses that another agent or group of agents will perform a particular action, both before he can monitor such action (or independently of his capacity ever to be able to monitor it) and in a context in which it affects his own action. When we say we trust someone or that someone is trustworthy, we implicitly mean that the probability that he will perform an action that is beneficial or at least not detrimental to us is high enough for us to consider engaging in some form of cooperation with him. Correspondingly, when we say that someone is untrustworthy, we imply that that probability is low enough for us to refrain from doing so [12].

Since that time, trust relationships and their influence on teamwork (e.g., coalition formation) have been an active area of research. *Trust* is defined by Griffiths and Luck as "an agent's estimate of how likely another agent is to fulfill its cooperative commitments. The risk of whether to cooperate, and with whom, may be determined by, among other things, the degree of trust [13]."

Castelfranchi and Falcone identified two parts to the trust of one agent in another agent: *core trust* and *reliance*. Core trust has two components: *competence* (Is the agent I want to perform a task for me capable of performing that task?) and *disposition* (Is the agent I want to perform a task for me willing to perform that task?). If the core trust components are satisfied, then the other part of trust comes into play: reliance. Once I have made the decision that an agent is able and willing to perform a task for me and I have asked the agent to do so, I am relying on that agent to complete the task [14].

Ramchurn, et al., considered two ways of assessing trustworthiness in agents: "(i) confidence derived (mainly) from analyzing the result of previous interactions with that agent and (ii) reputation acquired from the experiences of other agents in the community through gossip or by analyzing signals sent by an agent. Both measure the same property; that is, the agents believed reliability in doing what it says it will regarding particular issues of a contract [15]."

The problem with trust as it relates to autonomous robots at disaster sites is that human rescuers do not trust robots. This issue was noted by Murphy, who stated: "One impact of the human side is that rescue workers today refuse to consider fully autonomous systems designed to act as 'yes/no theres something down there' search devices [16]." The trust rescuers have in robots is related in large measure to their reliability.

B. Reliability

Reliability is a well defined term in engineering, with a long history of statistical measures and methods for probabilistic estimation. The simple definition for *reliability* is "the probability that a system will operate without failure under given conditions for a given time interval. We express reliability on a scale from 0 to 1: A system that is highly reliable will have a reliability measure close to 1, and an unreliable system will have a measure close to 0 [17]." Informally, human rescuers would consider a robot to be reliable when it was available for use when needed, it could carry out its task without human assistance or intervention, and it would not pose a hazard to either rescuers or the people being rescued.

As Murphy noted in [18] and [19], even after years of research and development, the reliability of search and rescue robots is abysmal. Murphy and her team from CRASAR deployed to the La Conchita, California mudslides in 2005. La Conchita is a coastal area of California, north of Los Angeles, that has a history of mudslides damaging or destroying houses. On January 10, 2005, a mudslide buried 15 houses. The CRASAR team responded two days later with two robots that had been specifically built for the search and rescue mission. These were not autonomous robots, they were teleoperated robots controlled by an operator via a tether. The first deployment was into a partially crushed home where an odor had been detected in the crushed area. A small robot was inserted into a void to seek the source of the odor. Before the robot even left the direct line of sight of the operator, it threw a track on a lump of clay and had to be retrieved manually by reeling in the tether. The second deployment was in a house that had already been searched, to provide a more benign environment for the robot to operate in with the goal of allowing the rescuers on site the opportunity to evaluate the utility of the robot. The robot was inserted through a window to provide a more realistic entry into the void – even though it was large enough for a person to enter. Shortly after being inserted into the house, the robot's tracks got tangled in the shag carpeting and the robot threw a track. Again, the robot was manually retrieved. The evaluation by the rescuers who observed the robot trials was that the robots were not reliable enough to be useful to the rescuers.

In 2007, Murphy's team responded to the Crandall Canyon coal mine collapse in Utah. The objective was to deploy a tethered robot down a six inch wide bore hole that had been drilled into a cavity where it was hoped some surviving miners could be located. Murphy's team managed to get the robot to the bottom of the hole, where they were able to look around, but they were not able to explore the void. When they attempted to retrieve the robot, the bore hole collapsed and the robot was trapped. They finally abandoned the robot. Losing the robot is not really a concern, as it is expected that a robot won't always be capable of being retrieved; however, it was disappointing that the robot could not explore the void in the nearly liquid coal slurry they encountered. Murphy's conclusion is that, as of 2007, search and rescue robot systems are not reliable enough for use by rescuers in the field.

C. Impact of Trust and Reliability on Use

The interrelation between these two factors should be apparent. Human rescuers will not use robotic systems if they do not trust them (especially true for autonomous systems). In order to trust robotic systems, they need to be demonstrated to be reliable. Right now, even teleoperated robotic systems are not reliable enough to gain the trust of human rescuers. This cycle tends to perpetuate the mistrust of robotic systems and will make the eventual use of robots in risk mitigation for disaster scenarios very difficult to achieve.

IV. Agent Model of Robotic Firefighters

To explore the interaction between the elements of trust and reliability, and to try to establish a utility metric for the use of autonomous robots in a disaster scenario, a firefighting model was created in the NetLogo agent modeling language [20]. NetLogo was selected as the language for the model because it has a simple, agent-oriented syntax, and it is highly extensible. In fact, the only real limitation on the size of the model is the resources available on the computer it is executing on. While the agent programming is done in a highly specialized and extended version of the Logo programming language (the famous "turtle" language created to introduce children to computer programming, algorithms, and logic), the development environment is written in Java; so the models created in NetLogo can be run on any platform with a Java Runtime Environment (JRE) or can be saved as an applet

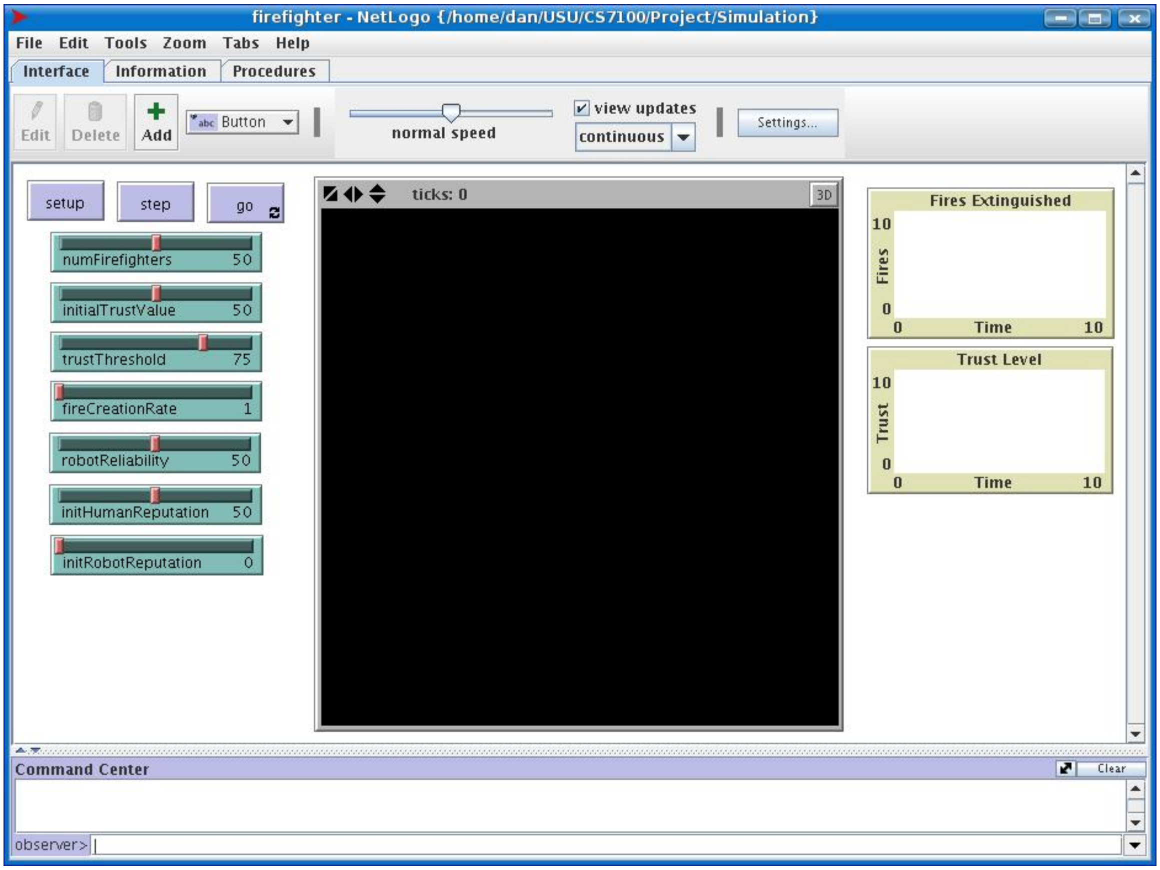

Fig. 5. The firefighting simulation in NetLogo.

embedded in a web page. A screenshot of the simulation is shown in Fig. 5.

The simulation interface is laid out with the simulation controls on the left-hand side of the window, the display of the firefighting scenario in the center pane, and the data collection plots on the right-hand side of the window. As can be seen in Fig. 5, there is still plenty of room on the interface for expansion of the model.

For the current iteration of the model, there are three control buttons and seven slider bars for setting the parameters of the simulation. (A slider bar allows the user to vary the value of a variable within program-established boundaries, prior to executing the simulation.) The *setup* button clears any remaining data or displays from a previous run of the simulation and initializes the simulation display. The *step* button causes the simulation to advance by one time step. The *go* button will run the simulation continually, until the button is clicked again to stop the simulation. If the go or step buttons are clicked after stopping the simulation, it will continue from the current state – unless the setup button has been clicked in the meantime. The *numFirefighters* slider selects the number of human firefighters to be created when the setup button is clicked. A number of robotic firefighters equal to the number of human firefighters will also be created. The slider selects between 0 and 100 firefighters, with a default value of 50 (the zero value was included for testing purposes during development of the model - obviously, the results of a simulation run with no firefighters will not be very interesting). The *initialTrustValue* slider allows a human trust value of 0 to 100 to be selected, with a default value of 50. In the current implementation, all of the human agents trust equally. The *trustThreshold* slider selects a value of 0 to 100, with a default value of 75. The trust threshold establishes the point at which the human firefighters start to trust the robotic firefighters and will call on them for assistance. In the current implementation, the primary influence on the human firefighters' trust level is workload – as the humans have increasing numbers of fires to put out, they will trust the robots more. The *fireCreationRate* slider has values of 1 to 10, with a default value of 1. This is the number of fires randomly started in each time step of the simulation. Values greater than one tend not to be very useful in practice (the fires spread too fast for the firefighters to react), so this control may be changed to the number of time steps between a new fire starting, instead of fires per time step. The *robotReliability* slider can take values from 0-100 and indicates the percentage reliability of the robots, where 0 would be the case where they always fail and 100 would mean they never fail. The default value is in the middle at 50. The next two sliders affect reputations exhanged between human agents. The *initHumanReputation* is the initial value for the reputation of other human agents held by each individual human agent. The range is 0-100 for this slider, with 0 indicating a totally negative reputation and 100 being a perfect reputation. The *initRobotReputation* is the same thing for initial human reputations for robotic agents.

The center pane of the simulation is the display of the firefighting scenario. Green areas are unburned grassland, the red “fire department logo” icons are human firefighters, the blue icons are robotic firefighters, the flame icons are active fires, and the black areas are burned grass. When the simulation is initialized, the firefighters (both human and robotic) are randomly distributed throughout the map. Also, the terrain is all identical – there is no rough terrain, no variation in elevation, and no difference in plant cover. There also is no difference in the firefighting agents – they are all equally effective and equally mobile.

The right side of the simulation window contains two data plots at present. One shows the number of fires active (orange line) and the number of fires extinguished by each agent type (red for human firefighters and blue for robotic firefighters). The other shows the value of the human trust value with relation to time.

Fig. 6 shows screen shots from the various phases of a typical simulation run. The first image shows the simulation after it has been initialized, the second shows the point the robotic firefighters are called upon by the human firefighters, and the third image shows the end of the simulation. Note that the firefighters have not been very successful in controlling the fires.

V. Preliminary Results

Fig. 7 shows the results of running the simulation. These results are typical of the current iteration of the simulation software. The human firefighters work alone, until they are not making enough progress putting out the fires. At that point, they call upon the robotic firefighters. The robots assist in putting out the fires until the situation is back under control. At that point, the humans decide they no longer need the robots and they stop requesting assistance. This usually continues to just about the end of the run, when the human firefighters can not reach the last of the fires and call on the robotic firefighters one last time. The predictability of these results are proof that the model needs to be extended to consider more factors and to be made more realistic.

VI. Conclusion and Future Work

Autonomous robots would appear to offer some real benefits in the management of the risk of disaster events. However, the very real social and technological barriers to the acceptance of robotic assistance at disaster sites need to be overcome. The modeling effort described in this paper is a step towards understanding the dynamics of these barriers, but the model needs to be extended to provide better data collection and more trust parameters that can be varied.

Some examples of changes that will be made in a future revision of the simulation are: A human firefighters only scenario is necessary so that a baseline performance metric can be established – this provides a point of comparison for the efficiency of fighting fires with and without the robotic firefighters. Individual trust values for the human firefighters, instead of the homogeneous global value, would be more

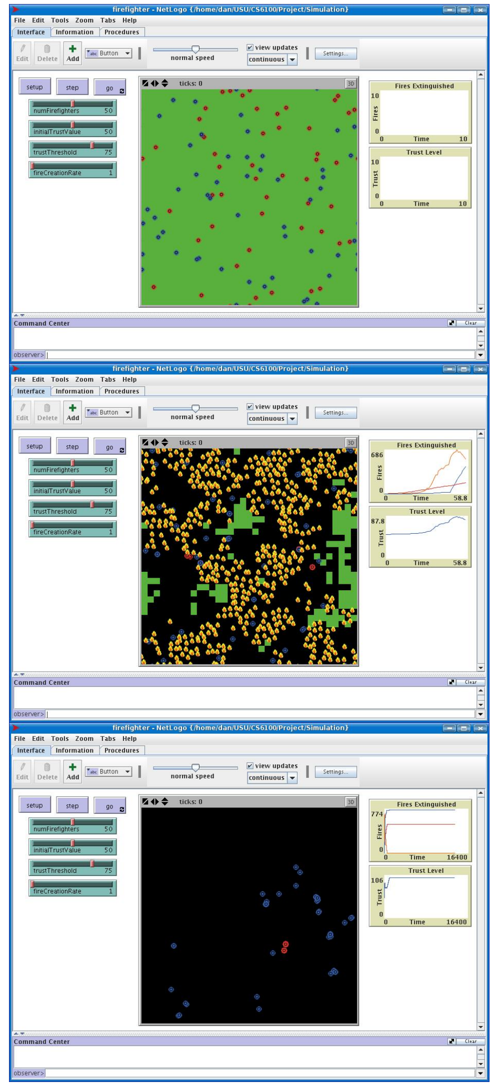

Fig. 6. The phases in a typical simulation run: Initialization, robot firefighters activated, and end of simulation.

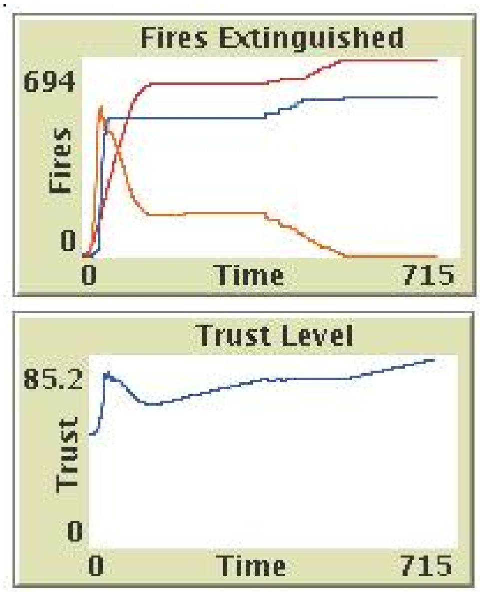

Fig. 7. The plots from a typical run of the firefighting simulation.

realistic. A switch to select a single starting trust value or a Gaussian distribution of trust values around the value set on the slider would also provide more opportunity for exploring the problem space. Future plans for the simulated environment are to make the terrain three dimensional, with varying terrain types, and different plant cover types. There also will be structures added to the terrain that will have a higher priority for the firefighters than open terrain. The structures will include both residential and commercial structures, as well as fire stations, where the firefighters will be based. Another future enhancement planned is to have heterogeneous agent types for both humans and robots; i.e., there will be human firefighters with different fire fighting and mobility capabilities. The same will be true for the robot firefighters. To improve data visualization, plots for future iterations of the model will include robot failure and human fatigue levels, as well as including the current average trust threshold on the plot with the trust level.

Once the extended functionality of the simulation has been implemented, the next logical step is to move to a higher fidelity simulation. While the NetLogo language is easy to develop in and it does have language support for large-scale agent simulations, there are limits to how much fidelity is possible in NetLogo.

The RoboCup Rescue Simulation system is a city-scale client/server environment that is being actively supported and extended by the RoboCup Federation [21]. The simulation consists of a number of servers responsible for different agent types: police, fire, and ambulance; operating in a high fidelity model of a city that has been struck by an earthquake. The simulation is extensible through agent, communication, and environment classes. Agents can be programmed in Java or C++ and the simulation can be viewed using a 2D or a 3D viewer, as shown in Fig. 8. There are also a number of data collection and analysis tools that can connect to the simulation. The tools allow analysis of agent performance, as well as comparison of agents over multiple simulation runs. This tool is the next logical step for a simulation that can investigate trust and reliability for a wide range of human and robotic agent types interacting in a realistic disaster scenario.

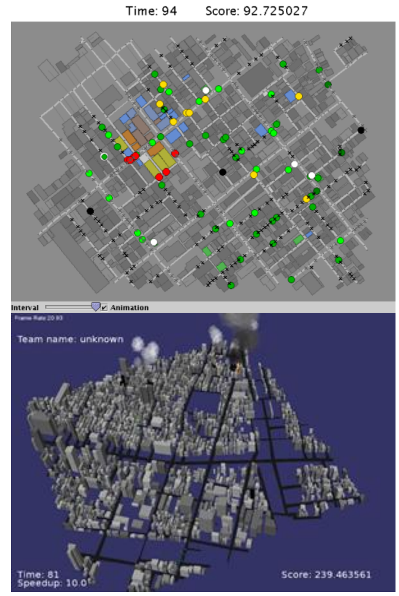

Fig. 8. The RoboCup Rescue Simulation as seen through the 2D viewer (top) and the 3D viewer (bottom).

REFERENCES

[1] S. Tadokoro, "Special project on development of advanced robots for disaster response (DDT Project)," *Proceedings of IEEE Workshop on Advanced Robotics and its Social Impacts (ARSO'05)*, Nagoya, Japan, June 12-15, 2005.

[2] University of South Florida, *Center for Robot-Assisted Search and Rescue*, Available [Online:] http://crasar.csee.usf.edu/MainFiles/index.asp.

[3] "SWORDS," *Wikipedia*, Available [Online:] http://en.wikipedia.org/wiki/Foster-Miller_TALON.

[4] "The Inside Story of the SWORDS Armed Robot 'Pullout' in Iraq: Update," *Popular Mechanics*, April 15, 2008, Available [Online:] http://www.popularmechanics.com/blogs/technology_news/4258963.html

[5] R. Murray, Z. Li, and S. Sastry, *A Mathematical Introduction to Robotic Manipulation*, Boca Raton, FL: CRC Press LLC, 1994, pp. 1-8.

[6] R. Arkin, *Behavior-Based Robotics*, Cambridge, MA: The MIT Press, 1998, p. 2.

[7] H. Takahashi, T. Suzuki, and T. Takahashi, "Hinomiyagura as a pacemaker in evaluating GIS environments," *RoboCup Rescue Simulation League 2005 Team Description Paper*, Osaka, Japan, July 11-19, 2005. Available [Online:] http://sakura.meijo-u.ac.jp/ttakaHP/MeijoRescue/HinomiyaguraAgent05.pdf.

[8] T. Takahashi, "Hinomiyagura: Comparison Simulation Results to Data of Local Government," *Robocup Rescue Simulation League 2006 Team Description Paper*, Bremen, Germany, June 14-20, 2006. Available [Online:] http://sakura.meijo-u.ac.jp/ttakaHP/MeijoRescue/HinomiyaguraAgent06.pdf

[9] National Institute of Standards and Technology, "Performance Metrics and Test Arenas forAutonomous Mobile Robots." Available [Online:] http://www.isd.mel.nist.gov/projects/USAR/.

[10] J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior*, Princeton, NJ: Princeton University Press, 2004.

[11] D. Gambetta, ed., *Trust: Making and Breaking Cooperative Relations*, New York, NY: Basil Blackwell Inc, 1988. Available [Online:] http://www.nuffield.ox.ac.uk/users/gambetta/gambetta_trust%20book.pdf.

[12] D. Gambetta, "Can We Trust Trust?," in *Trust: Making and Breaking Cooperative Relations*, electronic edition, Department of Sociology, University of Oxford, chapter 13, pp. 213-237, Available [Online:] http://www.sociology.ox.ac.uk/papers/gambetta213-237.pdf.

[13] N. Griffiths and M. Luck, "Coalition Formation Through Motivation and Trust," in *Proceedings of the 2nd International Conference on Autonomous Agents and Multi-agent Systems (AAMAS 2003)*, Melbourne, Australia, July 2003. Available [Online:] http://users.ecs.soton.ac.uk/mml/papers/aamas03.pdf.

[14] R. Falcone and C. Castelfranchi, "Social trust: a cognitive approach," in *Trust and Deception in Virtual Societies*, Norwell, MA: Kluwer Academic Publishers, 2001, pp. 55-90.

[15] S. Ramchurn, et al., "Devising a Trust Model for Multi-Agent Interactions using Confidence and Reputation," *Applied Artificial Intelligence*, issue 18, 2004, pp. 833-852.

[16] R. Murphy, "Rescue Robotics for Homeland Security," *Communications of the ACM*, special issue on Homeland Security, vol. 27, no. 3, March 2004, pp. 66-69. Available [Online:] http://crasar.csee.usf.edu/research/Publications/CRASAR-TR2004-1.pdf.

[17] S. Pfleeger, *Software Engineering:Theory and Practice*, Second edition, Upper Saddle River, NJ: Prentice Hall, 2001, p. 404.

[18] R. Murphy and S. Stover, "Rescue Robot Performance at the 2005 La Conchita Mudslides," *Proceedings of 1st Joint Emergency Preparedness and Response/Robotic and Remote Systems Topical Meeting (SSEHE 2006)*, Salt Lake City, Utah, February 11-16, 2006.

[19] R. Murphy, Presentation at 2007 IEEE International Workshop on Safety, Security, and Rescue Robotics, Rome, Italy, September 27-29, 2007.

[20] U. Wilensky, *NetLogo*, Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL. Available [Online:] http://ccl.northwestern.edu/netlogo/.

[21] RoboCup, *Rescue Simulation Leagues*, Available [Online:] http://www.robocuprescue.org/simleagues.html