

RoboCupRescue 2010 - Robot League Team RescueRobots Freiburg (Germany)

Christian Dornhege, Johannes Bendler, Roxana Bersan, Philipp Blohm, Martin Gloderer, Andreas Hertle, Thomas Liebetaut, Diego Cerdan Puyol, Alexander Kleiner, and Bernhard Nebel

Institute of Computer Science
Foundations of AI
Universität Freiburg
79110 Freiburg, Germany

<http://www.informatik.uni-freiburg.de/rescue/~robots>

Abstract. This paper describes the software and hardware system developed by the University of Freiburg team of search and rescue robots for the RoboCup Rescue 2010 competition. This system is an extension to the software that finished in first place the 2005 and 2006 autonomy challenge, focusing on two key areas: autonomous navigation and manipulation. Our team, consisting mainly of students, originates from the former CS Freiburg team (RoboCupSoccer), the ResQ Freiburg team (RoboCupRescue Simulation), and RescueRobots Freiburg teams '05 and '06.

1 Introduction

Our main contributions to the league are robust techniques for autonomous navigation and manipulation. We are implementing victim search with a manipulator carried by the mobile robot platform. This particularly allows to examine victim locations on different levels of height. Our goal is it to make this process autonomous and to execute it while the mobile robot platform explores the environment. The deployed mapping technique has been tested by First Responders during several events, and is currently integrated into a commercial robot.

For operator assisted exploration a flexible human robot interface (HRI) is developed, that displays all mission-relevant data to the operator. The GUI contains different views and allows for optimal control during autonomous and manual driving of the robot. To determine the possible location of victims, a computer vision system searches for holes in the surrounding walls. If they emit heat, the probability for a victim is increased and mission control then further analyzes the location. We use a 2D grid map to integrate data from the laser range finder (LRF) and possible victims as reported by the victim detection component.

To track the exploration progress of the robot, we update the sensor coverage in every cell for each sensor separately. The planned trajectories of the robot are then executed by a reactive driving behavior that follows way points along the trajectory path while avoiding obstacles. The *mission control* connects all on-line systems of the robot

together by monitoring their states and sending them commands. It is implemented as a set of timed automata.

RescueRobots Freiburg is a team of students from the University of Freiburg. The team's approach proposed in this paper is based on experiences gathered at RoboCup during the last ten years. The team originates from the former CS Freiburg team [Weigel *et al.*, 2002], which won the RoboCup world championship in the RoboCupSoccer F2000 league three times, and the ResQ Freiburg team [Kleiner *et al.*, 2005a], which won the RoboCup world championship in the RoboCupRescue Simulation league in 2004, and Rescue Robots 2005 [Kleiner *et al.*, 2005b], and 2006 [Kleiner *et al.*, 2006].

2 Team Members and Their Contributions

- Team Leader/Manipulation: Christian Dornhege
- SLAM: Andreas Hertle, Alexander Kleiner
- Controller Design and Behaviors: Martin Gloderer
- Victim Identification: Diego Cerdan Puyol, Thomas Liebetaut
- HRI: Roxana Bersan, Philipp Blohm
- Mission Control: Johannes Bendler
- Advisor: Alexander Kleiner, Bernhard Nebel

3 Operator Station Set-up and Break-Down (10 minutes)

Our robot is controlled by a lightweight laptop via a *Logitech Rumblepad*, which all can be transported together in a backpack. The *Matilda* robot can be transported by a moveable case with wheels and is easily two-man-portable. The whole setup and break-down procedure can be accomplished within less than 10 minutes, including booting the computers, checking the network connection, and checking whether all sensors work properly.

4 Communications

Autonomous as well as teleoperated vehicles are communicating via wireless LAN. We use an access point from Linksys, which is capable of operating in the *5 GHz* as well as in the *2.4 GHz* band. All communication is based on the Inter Process Communication (IPC) framework, which has been developed by Reid Simmons [Simmons, 1997]. The simultaneous transmission of multiple digital video streams is carried out by an error-tolerant protocol which we developed based on the IPC framework.

5 Control Method and Human-Robot Interface

We implemented a Human Robot Interface (HRI) for interacting with the robot. There are two situations to be considered. First we place emphasis on the autonomous control of the robot. Thus our main focus lies on the yellow arena and possibly the orange arena in autonomous operation.

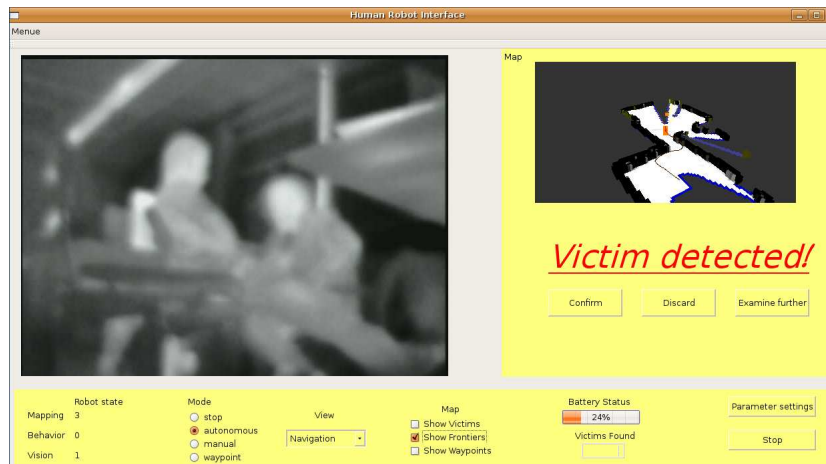


Fig. 1. This figure shows the Human Robot Interface when detecting a victim.

Apart from that teleoperation with a joystick linked to a portable laptop is possible. This mode of operation is primarily intended for the more complex arenas. In order to give the operator a comprehensive overview of the robot's state, different data is displayed. Besides images of the CCD and thermo camera, an elaborate map is shown. In addition to the 2D grid map of the surroundings, which is build based on laser scans, the robot position, victim locations, and additional information about the environment are displayed.

In order to highlight the relevant information in different situations our GUI can switch between certain views. For operator navigation a camera image is displayed large and centered, while additional information is shown smaller on the side. Additionally, there is a view centering the map and giving the operator the chance to define destinations the robot should travel towards, autonomously. Finally for the examination of victims there exists a view displaying the relevant information delivered by the vision systems. Here the operator can further examine the information gathered about the victim if unable to decide on whether the observation actually shows a victim on first sight.

Furthermore, the GUI has the functionality for the operator to switch between the four different robot modes, which are pause, autonomous operation, manual control and driving autonomously based on the waypoints set by the operator. Finally the states of the different modules of the robot as well as additional information about the robot state are displayed for supervision.

6 Map generation/printing

6.1 Simultaneous Localization And Mapping (SLAM)

We are generating 2D maps from laser-based scan matching. During the last decades a rich set of solutions for building maps from 2D laser range data has been proposed, such

as [Lu and Milios, 1997; Gutmann, 2000; Hähnel, 2005]. In contrast to scan matching methods, more sophisticated methods, such as *FastSlam* [Montemerlo *et al.*, 2002], and *GMapping* [Grisetti *et al.*, 2005], were introduced that are correcting the entire map at once when loop-closures, i.e., re-visits of places, are detected.

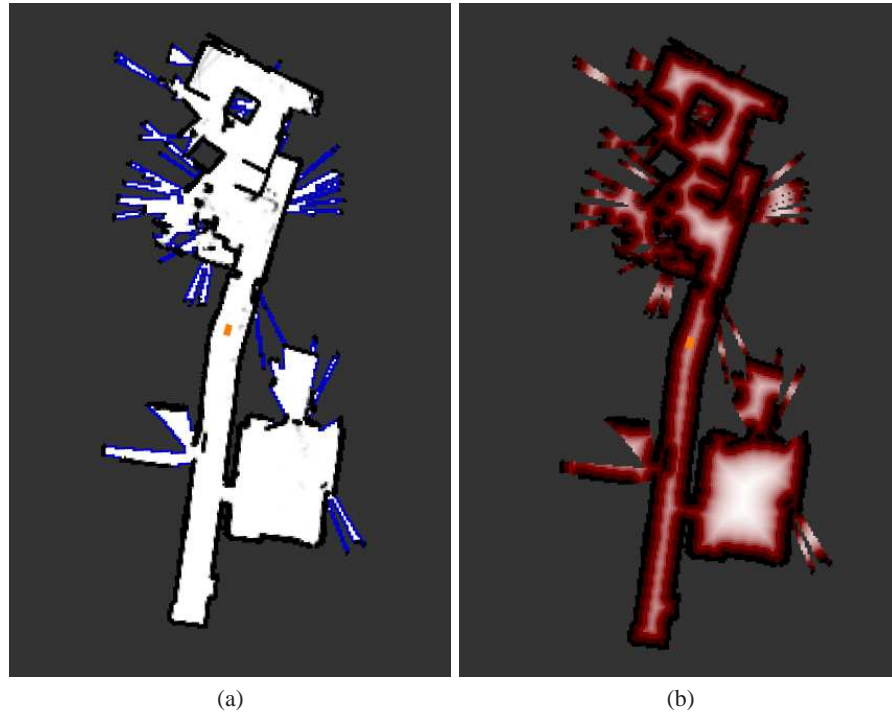


Fig. 2. These screenshots show the 2d grid map generated autonomously by the robot. (a) Occupancy map. Obstacles are colored black, traversable cells are colored white. Frontiers are colored blue. (b) Same map showing the discomfort values of cells in red. The darker the color, the bigger the discomfort cost for traversing that cell.

Although existing methods are capable of dealing with sensor noise, they do require reasonable pose estimates, e.g., from wheel odometry, as an initial guess for the mapping system. As commonly known wheel odometry tends to become unreliable given an unpredictable amount of wheel slip, which is frequently the case on rough terrain, such as found during USAR missions. Furthermore, methods performing loop-closures are mostly not applicable in real-time since their computational needs can unpredictably increase within unknown environments.

The mapping approach utilized for our robot team focuses on the application scenario of realistic teleoperation. Under certain constraints, such as low visibility and rough terrain, first responder teleoperation leads to very noisy and unusual data. For example, due to environmental make-up and failures in control, laser scans are frequently

taken under a varying roll and pitch angle, making it difficult to reliably find correspondences from successive measurements. In contrast to artificially generated data logs, logs from teleoperation seldom contain loops.

Most existing methods are following the principle of minimizing the squared sum of error distances between successive scans by searching over scan transformations, i.e., rotations and translations. Scan point correspondences are decided only once before the search starts based on the Euclidean distance. In contrast to other methods, our scan matching approach re-considers data associations during the search, which remarkably increases the robustness of scan matching on rough terrain. The algorithm processes data from laser range finder and gyroscope only, making it independent from odometry failures, which likely occur in such domains, e.g., due to slipping tracks.

The mapping approach has been extensively tested on robot platforms designed for teleoperation in critical situations, such as bomb disposal. Furthermore, the system was evaluated in a test maze by first responders during the Disaster City event in Texas 2008. Experiments conducted within different environments show that the system yields comparably accurate maps in real-time when compared to more highly sophisticated offline methods, such as Rao-Blackwellized SLAM. More details on the utilized mapping approach are found in [Kleiner and Dornhege, 2009].

The generated map integrates all sensor measurements: The laser data is used to find walls and obstacles like ramp, stairs and step fields. The visual camera and the thermal camera provide possible victim locations. In the grid map we also track the sensor coverage of the environment. Our thermal camera has a narrow field of view and a low range (compared to the laser range finder). With this information the robot can determine which parts of the environment were not yet explored with all available sensors. The same grid map data is transferred from the robot to the operator, so that the operator can monitor the progress of the mission. Figure 2 a) shows a map generated by the robot.

6.2 Exploration and Path Planning

The robot uses the map to plan a path to the next mission objective. Mission objectives can be frontiers or points of interests (e.g. heat sources). Our planning algorithm does not necessarily determine the shortest path to a mission objective, but rather the safest path, where the collision with obstacles and traversal over rough terrain can be avoided.

We employ a technique called Exploration Transform as proposed by Wirth and Pellenz [Wirth and Pellenz, 2007]. Based on the obstacle data acquired from the laser range finder we compute a distance map. With this information we assign the cells in the grid map additional discomfort cost. When we compute the length of a path, the discomfort cost is added to the traversal cost of the cell. So paths that bring the robot close to obstacles have a higher cost than paths with a safety distance to those obstacles. Figure 2 b) visualizes the discomfort cost.

With the Exploration Transform we can also select the best of multiple mission objectives. We initialize the grid cells of the objectives with a cost of zero. Then we use an efficient flood fill algorithm to build a gradient from the mission objectives to the robot's position in the map. To retrieve the safest path we follow the gradient until




Fig. 3. These screenshots of the 2d grid map show examples of the gradient computed during planning. (a) Planning to a nearby frontier. (b) Planning to an arbitrary chosen location.

we reach a mission objective. Figure 3 shows two planning gradients and the extracted paths.

7 Sensors for Navigation and Localization

7.1 Sensor Setup

Figure 4 shows the current sensor setup on our Matilda robot. Our main sensor is a Hokuyo UTM-30 laser range scanner that has a maximum range of 30 meters and delivers data at 40 Hz. The sensor can be mounted on a tilt unit enabling auto-levelling and 3D data acquisition. Additionally a vertically mounted Hokuyo URG-04LX supports the horizontal scanner. It is used to support SLAM on non-flat flooring and to detect prominent three dimensional structures as stairs during navigation. An XSens Mti inertial measurement unit completes the sensor setup for navigation.




Fig. 4. This figure shows the current sensor setup on our Matilda robot. The Hokuyo UTM-30 on a tilt unit (1), a vision camera (2) and a thermal camera (3) can be seen.

8 Sensors for Victim Identification

8.1 Manipulation

The sensor setup displayed in Figure 4 differs from the final setup that can be seen in Figure 5. Additionally, a Schunk 5-dof manipulator with more than one meter reach that is currently in production will be mounted in the center of the robot. The thermal camera, a vision camera, and a Hokuyo-URG-04LX laser range scanner will be mounted at the tool center point of the manipulator. During autonomous operation the manipulator is used to position these sensors directly in front of possible victim locations and thus enables us to find victims at any height.

8.2 Hole detection from Vision

Victim detection from vision uses a manipulator-mounted computer vision camera. The main target of the victim detection by vision is to identify the circular holes in walls, as those are spots where victims are possible. The real identification of the victim is then done by further analyzing a specific hole with the thermal camera.

The first step of finding such holes is shown in Figure 6. The camera image is converted to a gradient image using the Sobel edge detection operator. The brighter the




Fig. 5. This figure shows the final setup of the Matilda robot when the Schunk 5-dof manipulator is mounted.

edge, the more significant the edge. It can be clearly seen that the contours in the wall are detected as good edges, while the wood grain adds just little noise. As the holes as viewed from this angle are ellipses rather than circles, an efficient ellipse detection algorithm [Xie and Ji, 2002] is used. It derives the center position, the semimajor axis, and the rotation of the ellipse from two randomly chosen points. Then, a third point is chosen to determine the fourth parameter of the ellipsis. Using all remaining points for this third point then gives a voting table for different ellipse parameters, among which the best ellipse is used. This is similar to a normal Hough transformation for ellipses, but the dimension of the search space is reduced to one for each pair of points.

This algorithm can be speed-up by combining it with a RANSAC [Fischler and Bolles, 1981] approach. Not all pixels are evaluated but only randomly chosen ones until the ellipse is "good enough". To further reduce the input data, the gradients received from the Sobel operator can be used. As the holes are dark, only pixels with opposing gradient vectors should be used for the first two pixels in the algorithm. We are confident that these improvements are sufficiently accelerating the detection enabling real-time detection of hole structures in the environment. Finally, detected ellipses are utilized by the mapping system for marking potential victim locations in the map.

8.3 Victim Detection from Vision

While the robot is moving we perform heat detection with the thermal camera trying to find victims that are directly visible.

We use laser range information to detect walls and discard parts of the image that are higher than 1.2 meter and would cause problems like lights or spectators outside the arena boundaries. With this approach we avoid false positives and reduce the computation power of our algorithm getting a higher frame rate. After thresholding the heat




Fig. 6. Camera picture of three adjacent circles that are a possible location for victims. (a) Original camera image as input. (b) Same image after edge detection using the Sobel edge detection.

image we use a two pass connected component labeling algorithm to build the heat blobs. Small blobs are discarded and large blobs are sent to the mapping system to be positioned in the map and clustered to find victims.

Finally, the thermal camera is used to look inside holes and determine if a victim is inside or not.

9 Robot Locomotion

For locomotion we have implemented a reactive controller responsible for steering the robot around obstacles and towards a goal position. The controller works as follows: The space in front of the robot is searched for directions that are not obstructed by obstacles. To evaluate in which direction the robot can move freely, the controller tries to fit a driving channel in all possible directions using scans of the main laser range finder (see Figure 8). A driving channel always starts from the position of the robot and extends in the investigated direction. The channel has a fixed width, slightly bigger than the width of the robot, and a variable length equal to the distance to the nearest obstacle in the area covered by the channel. Channels that point in the direction of the goal position are considered better than channels leading away from it, and so are channels that minimize the amount the robot has to turn.

The channel can be parameterized for specifying the width directly in front of the robot and the width at the extent of the channel. Additionally, the minimum allowed length under which a channel is not considered useful can be set, as well as the maximum channel length at which all channels are capped.

To steer the robot through the best channel that was found, the velocity of the robot is set according to the length of the channel normalized by the maximum channel length, and the robot turns to face in the direction of the channel. Besides the driving channel the controller takes inertial information about the robot's pitch and roll angle into account to adapt to the current situation.




Fig. 7. Thermal camera picture of a corridor with three persons. The blue zone is discarded thanks to the information of the range laser.




Fig. 8. The figure displays the reactive controller. The driving channel is shown in orange and the endpoints of the laser scan are shown in red. As can be seen, the controller steers the robot through the gap between the wall and an obstacle.

10 Other Mechanisms

10.1 Mission Control System

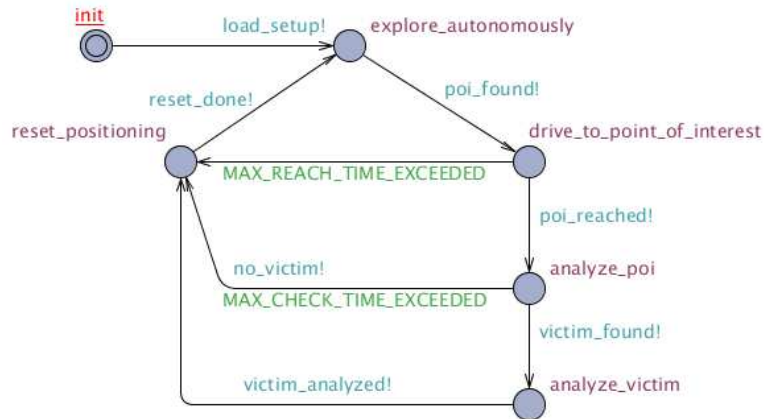


Fig. 9. An example of a top-level timed automaton for mission control.

The *mission control* connects all other on-line systems of the robot together by monitoring their states and sending new commands. It is implemented as a set of timed automata.

Figure 9 shows an exemplary top-level timed automaton. Referring to this as an example, the robot will start *autonomous_exploration* after the desired setup has been loaded. Whenever an on-line system (which may be e.g. thermal vision or the user interface) signals a new point of interest, the mission control checks whether it is important enough to be checked and acts accordingly. To reach the desired location, a sub-procedure sends commands to the software which controls the robot's motion. Once arrived, other sub-procedures also modelled as timed automata manage the analysis of the point of interest, e.g. by driving towards it or moving the manipulator, in order to search for a victim.

Using this architecture, the mission control system remains well-arranged and maintainable. It is used to implement various behaviors for different situations the robot may encounter and is based on system-wide available parameters (states and commands) served by a parameter daemon. Each on-line system delivers its own state to the daemon and listens for commands.

11 Team Training for Operation (Human Factors)

For the development of autonomous robots a sufficiently accurate physics simulation is absolutely necessary. Therefore, we utilized the USARSim simulation system [Wang *et*

al., 2005], which is based on the *Unreal2004* game engine (see figure 10 (a)) for simulating and developing the autonomous behavior of our robots. We have demonstrated the feasibility of this approach in the context of the proposed inter-league-challenge. Previous to Robocup German Open 2009 we developed SLAM and controller algorithms solely in the simulation without access to the Matilda robot. At the robocup site within three days of adaption we could reach stable autonomous driving in the yellow arena.

For real-world experiments we rebuild a rescue arena from standardized elements to test tele-operation and autonomous control.



Fig. 10. (a) shows the simulated Matilda in USARSim, based on the Unreal2004 game engine. (b) shows a part of our real-world test arena.

12 Possibilities for Practical Application to Real Disaster Site

Our team had no direct experience with any practical application in the context of real disaster response. However, we are confident that some of the techniques utilized by our team are very useful in the context of USAR. Within several efforts we are cooperating with NIST towards the goal of developing standards for benchmarking autonomous robots, particularly in the context of rescue and security.

Especially our mapping system has been demonstrated in the response robot exercise in Disaster City and is currently integrated into a commercial robot.

13 System Cost

The following table lists the approximate costs for our robot.

Name	Part	Price in USD	Number	Price Total in USD
Robot Base	Matilda Element	10000	1	10000
Manipulator	Schunk 5-dof	54000	1	54000
IMU	XSens Mti	2000	1	2000
Laser Range Finder	Hokuyu URG-04LX	1600	1	1600
Laser Range Finder	Hokuyu UTM 30	4500	1	4500
Thermo Camera	Thermal Eye	5000	1	5000
Laptop	Lenovo X61	3000	1	3000
USB Camera	PointGrey Chameleon	375	1	375
Sum Total:				80475

Table 1. Costs for the Matilda robot.

References

- [Fischler and Bolles, 1981] Martin A. Fischler and Robert C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6):381–395, 1981.
- [Grisetti *et al.*, 2005] G. Grisetti, Stachniss C., and Burgard W. Improving grid-based SLAM with rao-blackwellized particle filters by adaptive proposals and selective resampling. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, pages 667–672, Barcelona, Spain, 2005.
- [Gutmann, 2000] J.-S. Gutmann. *Robuste Navigation autonomer mobiler Systeme*. PhD thesis, Albert-Ludwigs-Universität Freiburg, 2000. ISBN 3-89838-241-9.
- [Hähnel, 2005] D. Hähnel. *Mapping with Mobile Robots*. PhD thesis, Universität Freiburg, Freiburg, Deutschland, 2005.
- [Kleiner and Dornhege, 2009] A. Kleiner and C. Dornhege. Operator-assistive mapping in harsh environments. In *Proc. of the IEEE Int. Workshop on Safety, Security and Rescue Robotics (SSRR)*, Denver, USA, 2009.
- [Kleiner *et al.*, 2005a] A. Kleiner, M. Brenner, T. Braeuer, C. Dornhege, M. Goebelbecker, M. Luber, J. Prediger, J. Stueckler, and B. Nebel. Successful search and rescue in simulated disaster areas. In *Proc. Int. RoboCup Symposium '05*. Osaka, Japan, 2005. submitted.
- [Kleiner *et al.*, 2005b] A. Kleiner, B. Steder, C. Dornhege, D. Höfer, D. Meyer-Delius, J. Prediger, J. Stücker, K. Glogowski, M. Thurner, M. Luber, M. Schnell, R. Kümmerle, T. Burk, T. Bräuer, and B. Nebel. RoboCupRescue - robot league team RescueRobots Freiburg (Germany). In *RoboCup 2005 (CDROM Proceedings), Team Description Paper, Rescue Robot League*, Osaka, Japan, 2005. (1st place in the autonomy competition, 4th place in the teleoperation competition).
- [Kleiner *et al.*, 2006] A. Kleiner, C. Dornhege, R. Kümmerle, M. Ruhnke, B. Steder, B. Nebel, P. Doherty, M. Wzorek, P. Rudol, G. Conte, S. Durante, and D. Lundstrom. RoboCupRescue - robot league team RescueRobots Freiburg (Germany). In *RoboCup 2006 (CDROM Proceedings), Team Description Paper, Rescue Robot League*, Bremen, Germany, 2006. (1st place in the autonomy competition).
- [Lu and Milios, 1997] Feng Lu and Evangelos Milios. Robot pose estimation in unknown environments by matching 2d range scans. *J. Intell. Robotics Syst.*, 18(3):249–275, 1997.
- [Montemerlo *et al.*, 2002] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit. FastSLAM: a factored solution to the simultaneous localization and mapping problem. In *Eighteenth national conference on Artificial intelligence*, pages 593–598, Menlo Park, CA, USA, 2002. American Association for Artificial Intelligence.
- [Simmons, 1997] Reid Simmons. <http://www-2.cs.cmu.edu/afs/cs/project/TCA/www/ipc/ipc.html>, 1997.
- [Wang *et al.*, 2005] J. Wang, M. Lewis, S. Hughes, M. Koes, and S. Carpin. Validating usarsim for use in hri research. In *Proceedings of the 49th Annual Meeting of the Human Factors and Ergonomics Society*, September 2005.

- [Weigel *et al.*, 2002] T. Weigel, J.-S. Gutmann, M. Dietl, A. Kleiner, and B. Nebel. CS-Freiburg: Coordinating robots for successful soccer playing. *IEEE Transactions on Robotics and Automation*, 18(5):685–699, 2002.
- [Wirth and Pellenz, 2007] Stephan Wirth and Johannes Pellenz. Exploration transform: A stable exploring algorithm for robots in rescue environments. In *Workshop on Safety, Security, and Rescue Robotics*, pages 1–5, 9 2007.
- [Xie and Ji, 2002] Yonghong Xie and Qiang Ji. A new efficient ellipse detection method. *Pattern Recognition, International Conference on*, 2:20957, 2002.