Amsterdam Oxford Joint Rescue Forces Team Description Paper Virtual Robot competition Rescue Simulation League Latin American RoboCup Open 2008

Arnoud Visser*, Tijn Schmits*, Steven Roebert*,
Gideon Maillette de Buy Wenniger* and Julian de Hoog†
*Universiteit van Amsterdam, 1098 SJ Amsterdam, the Netherlands
†Oxford University Computing Laboratory, Parks Road, Oxford OX1 3QD, UK
http://www.jointrescueforces.eu

***Abstract*— With the progress made in active exploration, the robots of the Joint Rescue Forces are capable of making deliberative decisions about the frontiers to be explored. The robots select the frontiers having maximum information gain. The robots incorporate the positions of their team mates into their decisions, to optimize the gain for the team as a whole. Active exploration is based on a shared occupancy map, which is generated online. The images of the omnidirectional camera can be used to generate bird-eye view maps and to visually estimate free space around the robot.**

INTRODUCTION

The RoboCup Rescue competitions provide benchmarks for evaluating robot platforms' usability in disaster mitigation. Research groups should demonstrate their ability to deploy a team of robots that explore a devastated area and locate victims. The Virtual Robots competition, part of the Rescue Simulation League, is a platform to experiment with multi-robot algorithms for robot systems with advanced sensory and mobility capabilities [1].

This year, shared interest in the application of machine learning techniques to multi-robot settings has led to a joint effort between the laboratories of Oxford and Amsterdam.

I. TEAM MEMBERS

UsarCommander was originally developed by Bayu Slamet and all other contributions have been built into his framework.

Arnoud Visser
: supervision [2], exploration & navigation algorithms [3], communication protocol [**?**]

Bayu Slamet
: user interface, real time visualization [4], several scan matching algorithms, manifold-SLAM [5], [6], communication protocol [**?**], exploration behaviors [3]

Max Pfingsthorn
: off-line rendering, several scan matching algorithms, manifold-SLAM, navigation behaviors [5], [6]

Tijn Schmits
: image processing, victim detection [4], sensor development [7], [8], user interface, communication protocol

Xingrui-Ji *et al.*
: occupancy grid map interpretation, beyond frontier exploration [9]

Aksel Ethembabaoglu
: image processing, active target tracking [10]

Steven Roebert
: map attribution, omnidirectonial camera usage [11], [12]

Gideon Maillette de Buy Wenniger
: image interpretation, learning to visually recognize free space

Julian de Hoog
: user interface, semi-autonomy, multi-robot exploration algorithms, communication roles [13], [1]

II. SCAN MATCHING

The possibilities for active exploration are heavily dependent on a correct estimation of a map of the environment. Many advanced techniques that aim to detect and correct error accumulation have been put forward by SLAM researchers. Although these SLAM techniques have proven very effective in achieving their objective, they are usually only effective once errors have already accumulated. With a robust scan matching algorithm the localization error is minimal, and the effort to detect and correct errors can be reduced to a minimum (see e.g. [14]).

Slamet and Pfingsthorn [5] performed an extensive survey of the performance of three scan matching algorithms in different environments. The survey demonstrated strong performance indoors, but less reliable results outdoors. Outdoor environments can contain large free spaces, where only sparsely obstacles are detected. Consequently, the scan matching algorithms were extensively tested in 2007 for outdoor environments and it was demonstrated that the robustness of the scan matching algorithms could be improved by matching against accumulated scans. With a storage technique like quad trees this accumulation can be done without losing the accuracy of the measurements.

One of the experiments was performed in the outdoor area of the 2006 Virtual Robot competition, which we call 'The Park' (see Fig. 1) . The experiment involved use of two implementations of the ICP algorithm [15]; IDC [16] and WSM [17]. The point-correlation procedures of the original implementations were replaced with a nearest neighbor-search in a quad tree. No additional modifications were made to the internal workings of these scan matchers, so we refer the interested reader to prior research [6], [5] and the original papers for further details. The experiments investigated the improvements that can be gained from using quad trees for both algorithms. The visualizations were created with the standard occupancy rendering techniques from [6]. All presented results are strictly based on scan matching.

Fig. 1. The outdoor area called 'The Park' of the 2006 competition world

For the experiment, an area of approximately 80 by 40 meters was used, with the robot starting in the bottom-right corner and traversing the park in a clockwise direction. The robot's path is shaded with gray for clarity and should describe a single closed loop from tip to tail. Both original scan matchers accumulate significant error; IDC 'overshoots' the end of the loop and WSM leaves a gap of several meters. Using the accumulated scans in the q-tree both IDC and WSM close the loop implicitly. Over the whole dataset the average correlation distance reduces from 9.83 mm to 4.83 mm for IDC and from 10.20 mm to 5.62 mm for WSM.

III. Localization and Mapping

The mapping algorithm of the Joint Rescue Forces is based on the manifold approach [14]. Globally, the manifold relies on a graph structure that grows with the amount of explored area. Nodes are added to the graph to represent local properties of newly explored areas. Links represent navigable paths from one node to the next.

The mapping algorithm is not dependent on information about the movement of the robot for the creation of links. In practice the displacement as reported by the inertial navigation sensor serves as an initial estimate for scan matching. Thereafter, displacement is estimated by comparing the current laser scan with laser scans recorded shortly before, stored in nearby nodes of the graph. As soon as the displacement becomes so large that the confidence in the match between the current scan and the previous scan drops, a new node is created to store the scan and a new link is created that corresponds to the displacement. A new part of the map is learned.

As long as the confidence is high enough, the information on the map is sufficient and no further learning is necessary. The map is just used to get an accurate estimate of the current location. The localization algorithm maintains a single hypothesis about where the robot currently is and does an iterative search around that location when new measurement data arrives. For each point the correspondence between the current measurement data and the previous measurement data is calculated. The point with the best correspondence is selected as the center of a new iterative search, until the search converges. Important here is the measure for the correspondence. For the Joint Rescue Forces, several scan matching algorithms are available (as introduced in the previous section) which can be used as correspondence measure.

The graph structure means that it is possible to maintain multiple disconnected maps. In the context of SLAM for multiple robots, this makes it possible to communicate the graphs and to have one disconnected map for each robot. Additionally, it is possible to start a new disconnected map when a robot loses track of its location, for example after falling down stairs.

The graph structure of the manifold can be easily converted into occupancy grids with standard rendering techniques, as demonstrated in Fig. 2 and [6].

IV. Multi-Robot Exploration

The approach of the UvA Rescue Team in previous years [18], [4] was to passively acquire the information to be stored in the map while the robot or operator was wandering around pursuing other objectives, like finding victims. This year however the focus will be on *active* exploration: to explicitly plan the next exploration action a which will increase the knowledge about the world the most. In this paradigm victim finding becomes the side-effect of efficient exploration.

A key aspect of this year's approach is that the *information gain* for areas of the environment not yet visited by the robot can be estimated with long-range laser range measurements. It is possible to generate two occupancy grids simultaneously [**?**]: one based on the maximum sensing range r_{max} of the range sensing device and another one based on a more conservative safety distance r_{safe}. Typical values for r_{max} and r_{safe} are 20 meters and 3 meters respectively.

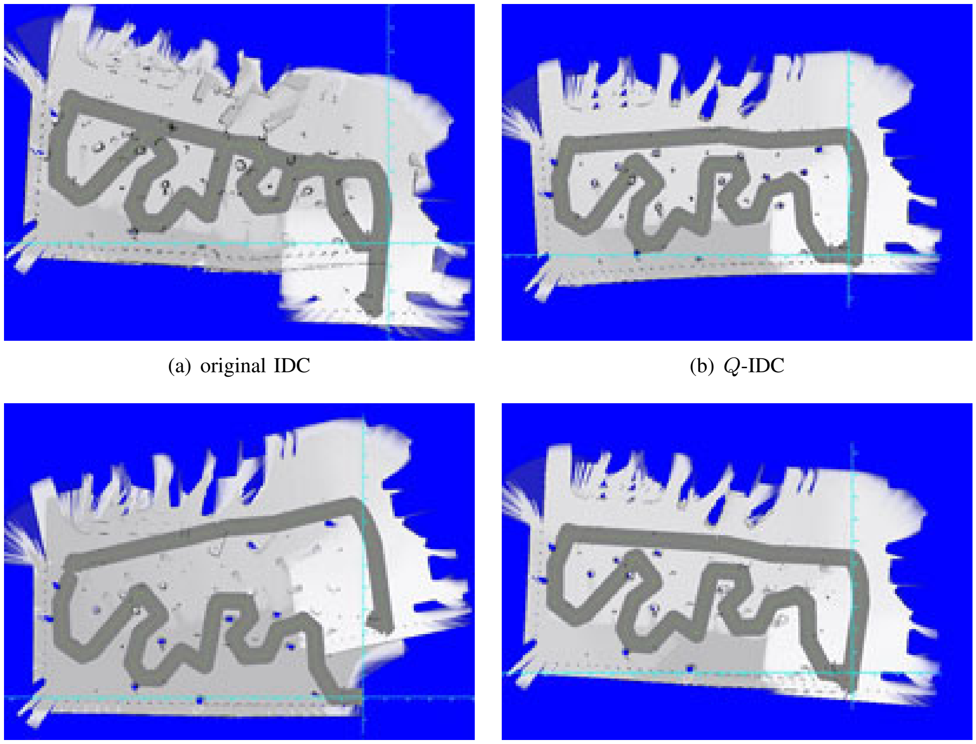

(a) original IDC (b) Q-IDC

(c) original WSM (d) Q-WSM

Fig. 2. Comparison of scan matching algorithms for a drive through a park, with poor odometry and sparse range scans.

The result is that the safe region determined by r_{safe} is a subset of the open area. Frontiers can then be extracted on the boundaries of the safe region where the robot can enter free space, and the area beyond each frontier (i.e. its associated *information gain*) can be estimated directly from the current map by calculating the amount of free space detected beyond it. For each frontier it is also straightforward to calculate how hard it is to reach (i.e. its associated *movement cost*) using a path planner.

Knowing both (i) the *information gain* and (ii) the *movement cost* for each frontier allows for active exploration. As discussed in [3], active exploration of the robots can be easily tuned by adjusting the balance between these two values. Shifting the balance in favor of information gain has the effect that robots explore mainly the corridors, while shifting the balance towards movement costs has the effect that the robots enter the rooms along the corridors.

V. Omnidirectional vision

Camera images can be used to automatically detect victims, independent from the Victim sensor provided by USARsim, as indicated in [4]. This independent information can be used to increase the robustness of the detection. This year the OmniCam sensor[1] is introduced in USARsim [7]. An omnidirectional catadioptric camera has some great advantages over conventional cameras, one of them being the fact that visual landmarks (such as victims) remain in the field of view much longer than with a conventional camera. This characteristic can be exploited during the competition.

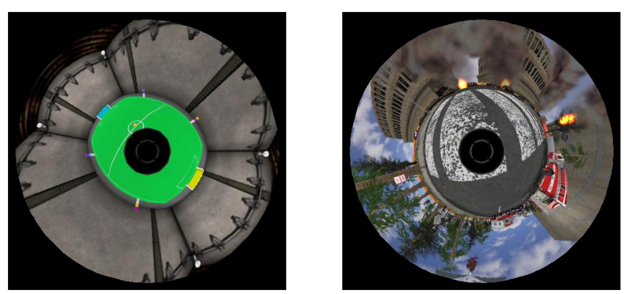

(a) Omnidirectional view of DM-spqrSoccer2006_250.utx (b) Omnidirectional view of DM-compWorldDay1_250.utx

Fig. 3. Images taken in the USARSim environment.

Omnidirectional pictures can be transformed to a Birds-Eye view. The correspondence between a pixel in the omnidirectional image $p_{\text{omni}} = (x_{\text{omni}}, y_{\text{omni}})$ and a pixel in the Birds-Eye view image $p_{\text{be}} = (x_{\text{be}}, y_{\text{be}})$ is defined by the following equations

$$\theta = \arccos \frac{z}{\sqrt{x_{\text{be}}^2 + y_{\text{be}}^2 + z^2}}, \quad \phi = \arctan \frac{y_{\text{be}}}{x_{\text{be}}} \tag{1}$$

$$\rho = \frac{h}{1 + \cos\theta} \tag{2}$$

$$x_{\text{omni}} = \rho \sin\theta \cos\phi, \quad y_{\text{omni}} = \rho \sin\theta \sin\phi \tag{3}$$

where h is the radius of the circle describing the 90 degree incidence angle on the omnidirectional camera effective

[1]The OmniCam package is available at `http://student.science.uva.nl/~tschmits/USARSimOmniCam/`

viewpoint. The variable z is defined by the distance between the effective viewpoint and the projection plane in pixels [19].

Using the previous equations, the following images were created. These resulting bird-eye view images portray a perspectively correct top-down view of the environment from directly above the robot.

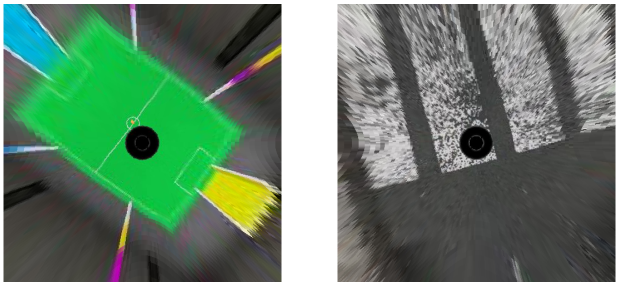

(a) Bird-eye view transformation of figure 3(a).

(b) Bird-eye transformation of figure 3(b).

Fig. 4. 500×500 pixel bird-eye transformations of Figures 3.

Ideally, all landmarks and soccer goals in Figure 4(a) would be depicted as if they were observed from far above as would be the case with a vertical orthographic projection of the environment. Unfortunately, orthographic projection cannot be performed on images produced by cameras which have a single effective viewpoint close to the ground. Perspective projection in combination with a relatively low position of the camera results in a depiction of landmarks and goals which is exceptionally stretched.

VI. Validation on Real World Data

To ensure the validity of our *scan matching* approach with data that suffers from real-world odometric errors and sensor noise, our algorithm is tested on a wide variety of datasets which are available thanks to the initiative of Andrew Howard and Nicholas Roy[2].

The occupancy grid maps illustrated in figure 5 were all created with the standard occupancy rendering techniques from [6]. All presented results are strictly based on scan matching, the SLAM algorithm was purely incremental. The occupancy grid maps can be compared with the original results (see for instance [20], [21]).

The data is collected in indoor environments, with many overlapping feature-rich, dense laser scans. For instance, the 'AP Hill' dataset (Fig. 5a) was collected for the DARPA/IPTO SDR project when four robots had to explore an unknown building at Fort AP Hill. The dataset is difficult because people were walking around the robots to check their progress. The 'CMU Newell Simon Hall' (Fig. 5b) is a relatively old and small. The difficulty in this dataset are the straight corridors without many features. The 'Intel Campus, Oregon' dataset (Fig. 5c) is collected by a P2DX robot during a tour of the part of the Intel Lab in Hillsboro, Oregon. The last dataset (Fig. 5d) is collected at our own location, a small tour around a staircase with a Nomad robot equiped with a Hokuyo laserscanner. These results illustrate the general applicability of our approach and more generally that developments in the Virtual Robot competition can be directly applied to fielded robotic systems.

[2]The Robotics Data Set Repository (Radish) available on `http://radish.sourceforge.net`

VII. Conclusion

This paper summarizes the approach of the Amsterdam Oxford Joint Rescue Forces as developed during the research of many enthousiastic students. At the Latin American RoboCup Open 2008 our autononomous exploration algorithm, as described in section IV, will be extensively tested with our remote participation. Although it became already clear at the RoboCup 2008 competition in Suzhou that it is difficult to outperform teams of teleoperated robots, it should be the goal of an Artificial Intellence researcher to take the challenge and try.

References

[1] A. Visser and J. de Hoog, "Amsterdam Oxford Joint Rescue Forces - Realistic Simulations to aid research and education in advanced Robot Control algorithms," in *Proceedings of the Scientific ICT Research Event Netherlands (SIREN 2008)*, September 2008, p. 22.

[2] S. Balakirsky *et al.*, "Towards heterogeneous robot teams for disaster mitigation: Results and performance metrics from robocup rescue," *Journal of Field Robotics*, vol. 24, no. 11-12, pp. 943–967, November 2007.

[3] A. Visser and B. A. Slamet, "Balancing the Information Gain Against the Movement Cost for Multi-robot Frontier Exploration," in *European Robotics Symposium 2008*, ser. Springer Tracts in Advanced Robotics. Springer-Verlag, February 2008, pp. 43–52.

[4] A. Visser *et al.*, "Design decisions of the UvA Rescue 2007 Team on the Challenges of the Virtual Robot competition," in *Proc. 4th International Workshop on Synthetic Simulation and Robotics to Mitigate Earthquake Disaster*, July 2007, pp. 20–26.

[5] B. A. Slamet and M. Pfingsthorn, "ManifoldSLAM: a Multi-Agent Simultaneous Localization and Mapping System for the RoboCup Rescue Virtual Robots Competition," Master's thesis, Universiteit van Amsterdam, December 2006.

[6] M. Pfingsthorn, B. A. Slamet, and A. Visser, *A Scalable Hybrid Multi-Robot SLAM Method for Highly Detailed Maps*, ser. Lecture Notes on Artificial Intelligence. Springer-Verlag, July 2008, vol. 5001, pp. 457–464.

[7] T. Schmits and A. Visser, "An Omnidirectional Camera Simulation for the USARSim World," in *Proceedings of the 12th RoboCup International Symposium*, July 2008, proceedings CD. To be published in the Lecture Notes on Artificial Intelligence series.

[8] T. Schmits, "Development of a Catadioptric Omnidirectional Camera for the USARSim Environment," Master's thesis, Universiteit van Amsterdam, June 2008.

[9] A. Visser, Xingrui-Ji, M. van Ittersum, L. A. González Jaime, and L. A. Stancu, *Beyond frontier exploration*, ser. Lecture Notes in Artificial Intelligence. Berlin Heidelberg: Springer-Verlag, July 2008, vol. 5001, pp. 113–123.

[10] A. Ethembabaoglu, "Active target tracking using a mobile robot in the USARSim," Bachelor's thesis, Universiteit van Amsterdam, June 2007.

[11] S. Roebert, "Creating a bird-eye view map using an omnidirectional camera," Bachelor's thesis, Universiteit van Amsterdam, June 2008.

[12] S. Roebert, T. Schmits, and A. Visser, "Creating a Bird-Eye View Map using an Omnidirectional Camera," in *Proceedings of the 20th Belgian-Netherlands Conference on Artificial Intelligence (BNAIC 2008)*, October 2008.

[13] A. Visser, T. Schmits, S. Roebert, and J. de Hoog, "Amsterdam Oxford Joint Rescue Forces - Team Description Paper - Virtual Robot competition - Rescue Simulation League - RoboCup 2008," in *Proceedings CD of the 12th RoboCup International Symposium*, July 2008.

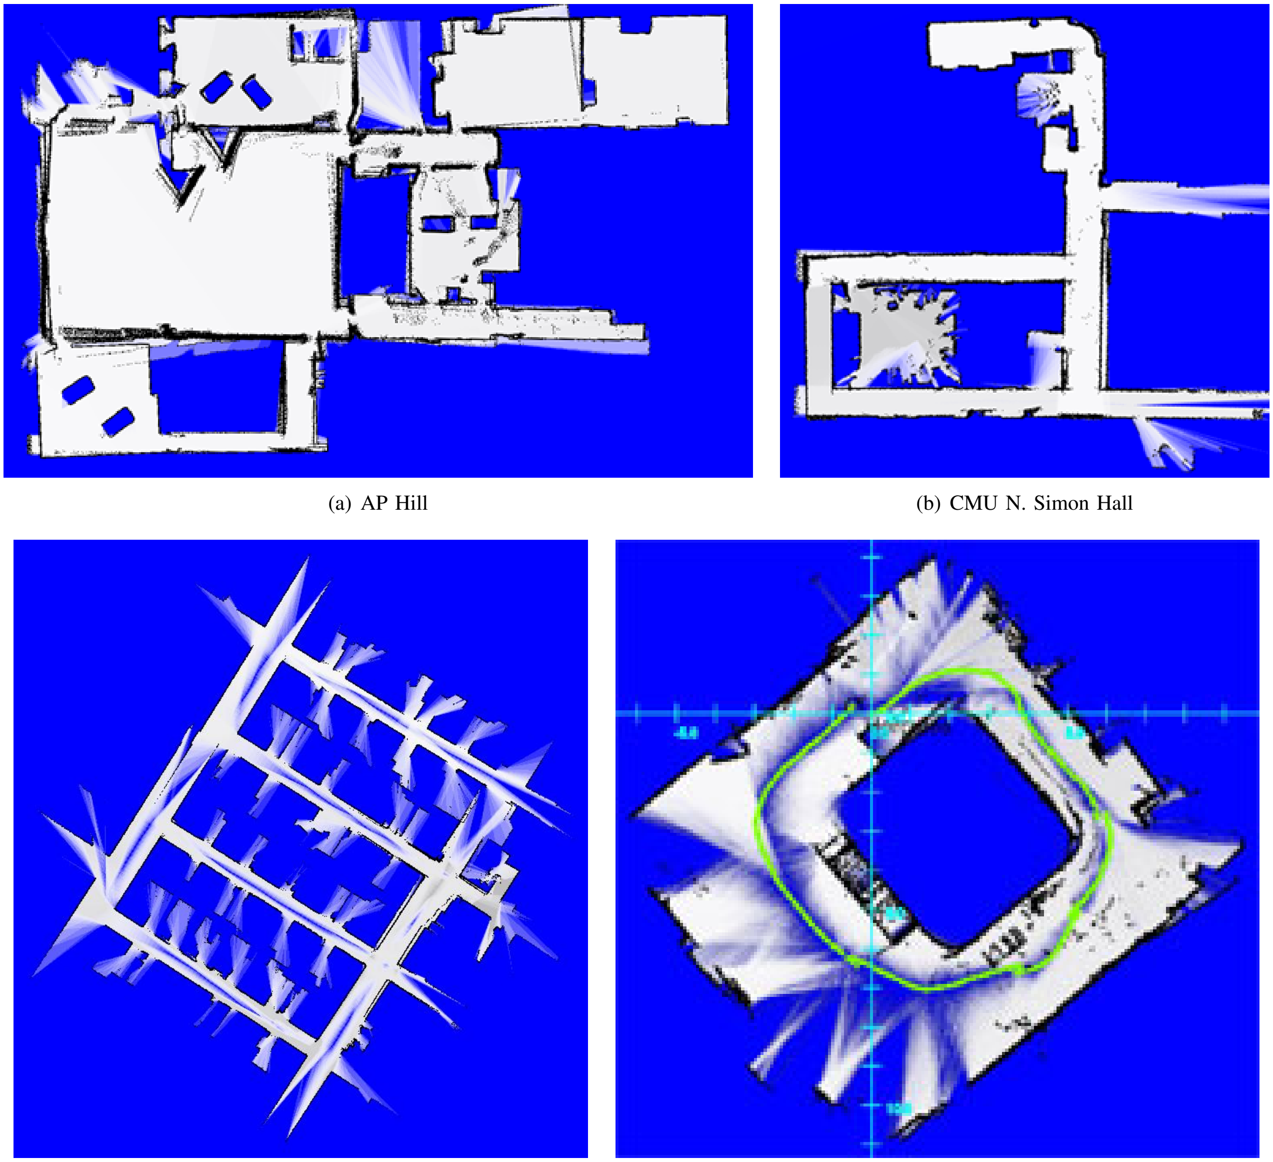

(a) AP Hill

(b) CMU N. Simon Hall

(c) Intel Campus, Oregon

(d) NIKHEF Amsterdam

Fig. 5. Occupancy maps generated for several datasets (mainly from Radish). The results that are acquired by running a version of our scan matcher, Q-WSM, as an iterative process, without any global optimization at the end of the process.

[14] A. Howard, G. S. Sukhatme, and M. J. Matarić, "Multi-robot mapping using manifold representations," *Proceedings of the IEEE*, vol. 94, no. 7, pp. 1360–1369, July 2006.

[15] S. Rusinkiewicz and M. Levoy, "Efficient variants of the ICP algorithm," in *Third International Conference on 3D Digital Imaging and Modeling (3DIM)*, Jun. 2001, pp. 145–153.

[16] F. Lu and E. Milios, "Robot Pose Estimation in Unknown Environments by Matching 2D Range Scans," *Journal of Intelligent and Robotic Systems*, vol. 18, pp. 249–275, 1997.

[17] "Weighted line fitting algorithms for mobile robot map building and efficient data representation," 2003, pp. 1667–1674.

[18] M. Pfingsthorn *et al.*, "UvA Rescue Team 2006; RoboCup Rescue - Simulation League," in *Proceedings CD of the 10th RoboCup International Symposium*, 2006.

[19] S. K. Nayar, "Omnidirectional vision," in *Proc. of Eighth International Symposium on Robotics Research ISRR'97*, Y. Shirai and S. Hirose, Eds. Springer-Verlag, 1998.

[20] D. Hähnel, D. Fox, W. Burgard, and S. Thrun, "A highly efficient fastslam algorithm for generating cyclic maps of large-scale environments from raw laser range measurements," in *Proceedings of the Conference on Intelligent Robots and Systems (IROS)*, 2003.

[21] T. Kollar and N. Roy, "Trajectory Optimization using Reinforcement Learning for Map Exploration," *The International Journal of Robotics Research*, vol. 27, no. 2, pp. 175–196, 2008. [Online]. Available: http://ijr.sagepub.com/cgi/content/abstract/27/2/175