UNIVERSITY OF AMSTERDAM

UvA-DARE (Digital Academic Repository)

Design decisions of the UvA Rescue 2007 Team on the Challenges of the Virtual Robot competition

Visser, A.; Slamet, B.; Schmits, T.; González Jaime, L.A.; Ethembabaoglu, A.

Publication date
2007

Published in
Proc. 4th International Workshop on Synthetic Simulation and Robotics to Mitigate EarthquakeDisaster

Link to publication

Citation for published version (APA):
Visser, A., Slamet, B., Schmits, T., González Jaime, L. A., & Ethembabaoglu, A. (2007). 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 EarthquakeDisaster* (pp. 20-26)
http://www.science.uva.nl/research/isla/pub/Visser07smred.pdf

General rights
It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations
If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: https://uba.uva.nl/en/contact, or a letter to: Library of the University of Amsterdam, Secretariat, P.O. Box 19185, 1000 GD Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (https://dare.uva.nl)

Download date:15 Oct 2026

Design decisions of the UvA Rescue 2007 Team
on the Challenges of the Virtual Robot competition

Arnoud Visser, Bayu Slamet, Tijn Schmits,
Luis A. González Jaime, and Aksel Ethembabaoglu

Universiteit van Amsterdam (UvA), Intelligent Systems Laboratorium Amsterdam,
Kruislaan 403, 1098SJ, Amsterdam, Netherlands
www.science.uva.nl/~arnoud/research/roboresc/

Abstract. This discussion paper highlights the design decisions of the UvA Rescue Team on the challenges imposed by the 2007 competition. Consistent with the approach of the previous year, the UvA Rescue Team research focus is dedicated to perceptional issues, to allow later extensions towards complex multi-agent exploration. The Simultaneous Localization and Mapping (SLAM) approach is extended with additional scan matching techniques. Victim detection is made possible with Histogram based Skin Detection.

Key words:
SLAM, Scan Matching, Skin Detection, Frontier Exploration

1 Introduction

Urban Search And Rescue is a challenging area of robotics research. In general the problem is not solvable by a single robot, and a heterogeneous team of robots that dynamically combine individual capabilities and cooperatively solve the task is needed [8]. Many research and development aspects still have to be worked out in order to be able to field efficient multi-robot solutions that jointly explore and map an environment while searching for victims. The Virtual Robot competition provides several elementary tests to demonstrate progress in the skills necessary for urban search and rescue. The skills tested are mapping, mobility, victim finding, communication and cooperation. This discussion paper highlights the design decisions of the UvA Rescue Team on the challenges imposed by the 2007 competition.

2 Virtual competition challenges

Based on the experiences acquired in the competition of previous year [1], the conclusion was drawn that it is difficult to discern the exact qualifications that support the claim "best in search and rescue". In 2006 a single equation was used that balanced the victim discovery, mapping and exploration skills with the autonomy of the team. In 2007 the focus will be on victim discovery with

small teams of robots with interactive support of a human operator. The scoring function of this year can be summarized with the following equation:

$$S = \frac{V_{LOC} * 10 + V_{PICT} * 5 + E * 50 - F_{NEG} * 5 - F_{POS} * 5 - C * 5}{(1 + N)^2} \quad (1)$$

This equation resembles the scoring function of 2006, but hides the increased difficulty of this year's competition in how the points on the different aspects can be acquired. A good example is the element V_{LOC}, which indicates the number of victims correctly localized (each worth 10 points). In the previous year all victims where equipped with a RFID-tag that broadcasted their ID, which could be detected with a specific victim sensor. This year the victim sensor returns unidentified body parts, which should be combined into a single victim with an accurate estimation of the location. /begincommment For the more challenging environments, where victims are trapped under, behind, or on top of objects, this is not trivial [6]. /endcomment

Another example is the element E. In the previous year this element represented the area that was explored (identified as traversable / not traversable). This year this element represents the area that is 'cleared' (traversable and victim free). Areas can only be cleared after close inspection with a camera, while traversable can be estimated with long range laser measurements. V_{PICT} represents the quality of a picture to asses the state of a victim (5 bonus points). F_{NEG} and F_{POS} represent respectively the false negative and false positives of the victim reports. A penalty of 5 points is given for each unreported victim in a cleared area and the same penalty is given for each victim reported while no victims are present in the vicinity. C indicates the penalty for a robot that collides with a victim.

In this equation no explicit reference is made towards the mapping quality. Still, the mapping quality is an important factor, because of its implicit influence on the elements V_{LOC}, F_{NEG} and F_{POS}. In the previous year maps could be shared between robots based on the IDs of the RFID-tags distributed throughout the map. This year these tags are no longer present, which means that the merging of maps has become a much harder task.

Last, but not least, is the discount based on the number of operators N. In 2006 the top three teams where fully autonomous, while the other teams had a single operator. This year every team should have at least one operator, to represent the USAR practice where a human operator is needed to setup and start the system. This means that the algorithms have to be modified to generate intermediate results that are understandable and to a human operator.

In the next sections the design decisions on the competition challenges are worked out.

3 Mapping

To demonstrate the mapping competency, an accurately georeferenced map should be delivered of an elementary-test world. An example of such an elementary-test

world is given in figure 1. The map must be annotated with the explored area, cleared area, victim locations and the robot's path. The accuracy of the map will be estimated by comparing it against the ground-truth using a georeferencing tool.

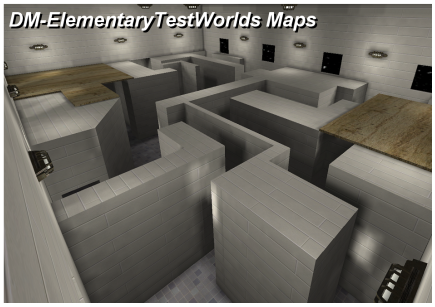

Fig. 1. A part of the elementary-test world used for mapping

The prizewinning mapping algorithm of the UvA Rescue Team is based on the manifold approach [3]. 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 while links represent navigable paths from one node to the next. The UvA Rescue Team takes no information about the actual movement of the robot into account while creating the links. All information about displacements is exclusively derived from the estimates obtained by scan matching. This displacement is estimated by comparing the current laser scan with laser scans recorded shortly before, stored in nearby nodes of the graph. In principle the scan matcher can also perform a comparison with measurements elsewhere in the graph, but such a comparison is only made under specific circumstances (as demonstrated in [10]). At the moment that the displacement becomes so large that the confidence in the match between the current scan and the previous scans drops, a new node is created to store the new scan and a new link is created with the displacement estimate. A new part of the map is learned.

Laser range scanners can deliver highly accurate measurements, and a position estimate based on scan matching outperforms position estimates based on more direct measurements of the displacement such as the measured odometry or the measurements yielded by inertial navigation sensors. The difference can be seen in the mapping results given in figure 2. At first glance both maps look equivalent, but a closer inspection reveals that obstacles as walls are much sharper represented in figure 2.b. Maps of this quality cannot be generated based on odometry measurements[1].

[1] See the logbook entry of 13 June 2007 at `http://www.slamet.nl/log.html`

a 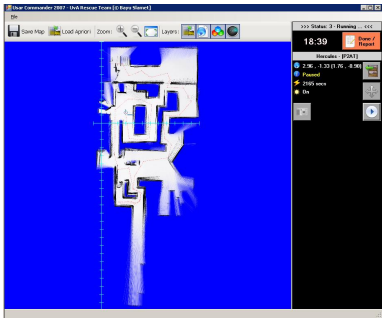 b 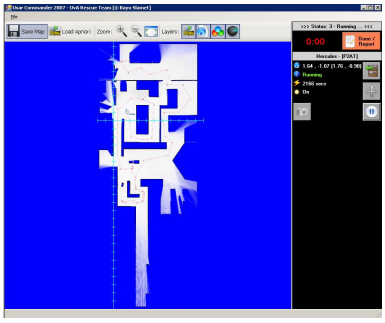

Fig. 2. Prelimenary results of a map created from the environment as displayed in figure 1. Figure 2.a is based on measurements of the inertial navigation sensor, figure 2.b is based on laser scan matching

Although the results are impressive for indoor environments, an evaluation study [12] indicated that the applied scan matching algorithm [11] performs less in outdoor environments. Current research concentrates on introducing alternative matching algorithms [5, 7], that perform better in outdoor environments. The ambition is to create a combination of those algorithms. Each algorithm gives an indication of its confidence of its match, and only the match with the highest confidence is needed for an accurate position estimate.

Another possible line of research is to globally optimize the map [9] in the case of loop-closure and map-merging. An important prerequisite for such an approach is an indepent criterion to start such an optimalization. Uniquely identifyable landmarks are needed to be sure that the same location is visited multiple times. In the previous year the unique RFID-tags of the victims could be used as landmarks. This year the RFID-tags should be released by the robot itself. An alternative approach is to use an independent sensor to start optimization, such as the inertial navigation sensor or the correspondence in appearance from the camera.

4 Victim detection

For the victim detection test, the robots must find, identify, and report the location of as many victims in the allotted time. The environment used for this test will not present mobility challenges but will present perception challenges.

To face this challenge, the reports (partly false) from the artificial VictSensor will be verified by automatically analyzing the images from the camera carried by the robot. An initial check will be based on a histogram based Skin Detection approach. A general 3D colour histogram model will be constructed in which discrete probability distributions are learned [4]. Given skin and non-skin histograms based on training sets we can compute the probability that a given colour value belongs to the skin and non-skin classes.

4.1 The Colour Model

We first construct a general colour model from the generic training set using a histogram with 32 bins of size 8 per channel in the RGB colour space. The histogram counts are converted into a discrete probability distribution $P(\cdot)$ in the usual manner: $P(rgb) = \frac{c[rgb]}{T_c}$ where $c[rgb]$ gives the count in the histogram bin associated with the RGB colour triple rgb and T_c is the total count obtained by summing the counts in all of the bins.

4.2 Skin Detection

We derive a skin pixel classifier through the standard likelihood ratio approach [2]. Given skin and non-skin histograms we can compute the probability that a given colour value belongs to the skin and non-skin classes: $P(rgb|skin) = \frac{s[rgb]}{T_s}$, $P(rgb|\neg skin) = \frac{n[rgb]}{T_n}$ where $s[rgb]$ is the pixel count contained in bin rgb of the skin histogram, $n[rgb]$ is the equivalent count from the non-skin histogram, and T_s and T_n are the total counts contained in the skin and non-skin histograms, respectively.

Given a certain threshold, Θ, based on the costs of false positives and false negatives, a skin pixel classifier is constructed:

$$\frac{P(rgb|skin)}{P(rgb|\neg skin)} \geq \Theta \tag{2}$$

An example of this classifier, preliminary trained in the small world 'DM-VictimTest' with only three victims, is given in figure 3. Because all three victims weared the same clothing, blue and white are still important components of this probability. Extending the training set with a wider variety of victims will reduce the influence of those colours, in favour of proper skin values.

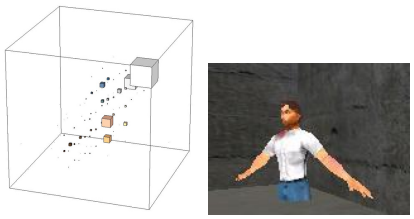

Fig. 3. A plot of $\frac{P(rgb|skin)}{P(rgb|\neg skin)}$ derived from an environment of which the image to the right is a camera-image during positive VictSensor readings.

This classifier can be used to verify the artificial VictSensor readings, and to detect victims on larger distances and behind glass. This classifier can also be used to initiate a tracking algorithm based on colour-histograms [14] to be able to cope with walking victims.

5 Mobility

For the mobility test, a team of robots will all start in a prescribed starting area. In order to pass this test, at least one robot must achieve a given goal location before time expires. It is the challenge to accomplish this goal by smart individual exploration behaviours. Cooperative exploration is tested in the next challenge.

Part of the current research [13] is frontier exploration. On a map several interesting locations can be present where the exploration can be continued, referred to as exploration frontiers. Typically, a greedy algorithm is used for the decision which frontier to explore next. Such a greedy algorithm only considers interesting locations locally, focused to reduce the movement costs. More sophisticated algorithms also take into account the information that can be gained along each frontier. This shifts the problem to estimate the amount of unexplored area behind the frontiers on the global map. Our algorithm exploits the long range of current laser scanners. Typically, during the previous exploration a small number of laser rays already passed the frontier, but this number is too low to have major impact on the generated map. Yet, the few rays through a frontier can be used to estimate the potential information gain from unexplored area beyond the frontier.

Extensions of this challenge include pitch and roll ramps, uneven flooring, and low overhangs. Currently no specific research is dedicated to navigate over such terrain, but initial test are performed with the Talon robot to address these challenges.

6 Cooperation

Cooperative exploration is possible, when frontier exploration is performed on a shared map. This requires the ability to merge individual maps, as indicated in section 3.

7 Communication

The prerequisite of this year is to perform all communications (operator-robot and robot-robot) via the Wireless Simulation Server. Due to the delays and speed-limit of this Server, this has a major impact on the performance of the team. Initial research has started to cope with this challenge, but in practice this prerequisite forces us to make all decisions autonomously on the robots.

8 Conclusion

The Technical Committee of the Virtual Robot competition has created an extensive set of challenges for the scientific community. No team will have the time or desire to become expert in every aspect of the competition. The UvA Rescue Team has decided to focus the research on perception, to allow later extensions towards complex multi-agent exploration.

References

1. S. Balakirsky, S. Carpin, A. Kleiner, M. Lewis, A. Visser, J. Wang and V. A. Ziparo, "Towards heterogeneous robot teams for disaster mitigation: Results and Performance Metrics from RoboCup Rescue", *Journal of Field Robotics*, to appear.
2. K. Fukunaga, *Introduction to statistical pattern recognition (2nd ed.)*, Academic Press Professional, Inc., San Diego, CA, USA, 1990, ISBN 0-12-269851-7.
3. A. Howard, G. S. Sukhatme and M. J. Matarić, "Multi-Robot Mapping using Manifold Representations", *Proceedings of the IEEE - Special Issue on Multi-robot Systems*, 2006.
4. M. J. Jones and J. M. Rehg, "Statistical Color Models with Application to Skin Detection.", *International Journal of Computer Vision*, volume 46(1):pp. 81–96, 2002.
5. F. Lu and E. Milios, "Robot Pose Estimation in Unknown Environments by Matching 2D Range Scans", in "IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 1994", pp. 935–938, 1994.
6. M. Micire, *Analysis of the Robotic-Assisted Seach and Rescue Response to the World Trade Center Disaster*, Master's thesis, University of South Florida, July 2002.
7. J. Minguez, L. Montesano and F. Lamiraux, "Metric-Based Iterative Closest Point Scan Matching for Sensor Displacement Estimation", *IEEE Transactions on Robotics*, volume 22(5):pp. 1047–1054, October 2006.
8. R. Murphy, J. Casper, M. Micire and J. Hyams, "Mixed-initiative Control of Multiple Heterogeneous Robots for USAR", Technical Report CRASAR-TR2000-11, Center for Robot Assisted Search & Rescue, University of South Florida, 2000.
9. E. Olson, J. Leonard and S. Teller, "Fast Iterative Optimization of Pose Graphs with Poor Initial Estimates", in "Proceedings of the IEEE International Conference on Robotics and Automation", pp. 2262–2269, 2006.
10. M. Pfingsthorn, B. Slamet and A. Visser, "A Scalable Hybrid Multi-Robot SLAM Method for Highly Detailed Maps", in "Proceedings of the 11th RoboCup International Symposium", July 2007.
11. S. T. Pfister, K. L. Kriechbaum, S. I. Roumeliotis and J. W. Burdick, "A Weighted Range Sensor Matching Algorithm for Mobile Robot Displacement Estimation", *IEEE Transactions on Robotics*, 2007, (to appear).
12. B. 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.
13. A. Visser, Xingrui-Ji, M. van Ittersum, L. A. González Jaime and L. A. Stancu, "Beyond frontier exploration", in "Proceedings of the 11th RoboCup International Symposium", July 2007.
14. Z. Zivkovic and B. Kröse, "An EM-like algorithm for color-histogram-based object tracking", in "IEEE Conference on Computer Vision and Pattern Recognition", volume 1, pp. 798–803, June 2004.