Rescue Robotics - a crucial milestone on the road to autonomous systems

Andreas Birk and Stefano Carpin
International University Bremen – Germany
{a.birk,s.carpin}@iu-bremen.de

Abstract

In this article we argue that rescue robotics is an important steppingstone in the scientific challenge to create autonomous systems. We motivate why we believe that there is a significant market for rescue robots, which has unique features that allow a fruitful combination of application oriented developments and basic research. Based on several application examples, for example the estimate that about 2000 road accidents per year in Germany alone involve hazardous goods, we conclude that there is a tremendous need for rescue robots. Unlike other markets for advanced robotics systems like service robots, the rescue robotics domain benefits from the fact that there is a human in the loop, which allows a stepwise transition from dumb teleoperated devices to truly autonomous systems. Current teleoperated devices are already very useful in this domain and they benefit from any bit of autonomy added as human rescue workers are a scarce resource at disaster scenarios and a single operator should ideally supervise a multitude of robots. We present results from the rescue robots and the International University Bremen (IUB) in a core area regarding autonomy, namely mapping.

keywords: autonomous systems, world modeling, mapping, cooperation, service robots

1 Introduction

The quest for autonomous systems is figuratively speaking the holy Grail of research in robotics and artificial intelligence. This does not only hold in respect to the basic research interests in autonomy, which can be traced down to the investigation of the core principles of the human mind, but it is also important from an application view point. For example, service robots that work in public places, offices and our homes will need a high degree of autonomy to be useful and safe. They need to operate in various dynamic environments where they encounter a multitude of unknown objects. They should be able to handle partially specified tasks. They have to adapt to unforeseen situations, and so on. Furthermore, service robots will need to be inexpensive devices produced for consumer markets to be affordable by everyone. It follows that for future autonomous systems, for which service robots can serve as a prime example, there are two development goals that are hard to combine. On one hand, there is the need to develop leading edge high-technology enabling the systems with autonomy, a challenge which includes still many unsolved research issues and which requires the according large investments. On the other hand, there is the need for simple, cost-effective systems to address mass markets. From this perspective, we see rescue robots as a crucial milestone on the road to autonomous systems. Unlike in other domains, rescue robotics already features a high end market from which the field can stepwise develop toward autonomous systems used in mass markets.

When we look at the current state-of-the-art of safety and rescue robots, we can see that it is dominated by special-purpose custom-made systems. There are for the example about thirty bomb squad robots in Germany. Each of these robots costs in the order of several hundred thousand euros and relies heavily on the human operator who activates motors in open loop, only guided by visual and acoustic feedback via a camera and microphones. The robots only allow to access normal buildings, i.e., they can negotiate stairs, and they provide very crude manipulation skills. To open doors for example, the onboard robot arms are not suited and explosive charges are used. So there is a large potential to improve this type of robot, both in respect to its capabilities as well as in respect to its cost.

Take for example the potential to use this type of robot in accidents, which involve hazardous material. In these scenarios, robots that allow an inspection from a safe distance are tremendously helpful tools. Note that about half of the accidents in Germany that lead to a severe pollution of water happen during transportation according to the German Ministry for the Environment. According to the German Ministry of Transportation, there are each year about 250 million transports via truck in Germany of which about 10 million involve highly hazardous material. About 50,000 accidents per year on German roads involve trucks. Based on the general percentage of hazardous transports compared to the overall number of road transports, we can estimate that about 2000 accidents per year in Germany include trucks carrying hazardous material. The 5.5 accidents per day are spread over the overall area of Germany with its 35 million hectares. It is therefore highly desirable to have a dense coverage of Germany with robots that can serve as tools to assess the state of hazardous material after road accidents. Supposing that a working solution could be supplied at a cost of approximately 10,000, we expect that there is sufficient interest to equip each of the 25,000 German fire brigade stations with an according system. Based on this example, which can be easily extended to other countries than Germany and to other domains than road accidents with hazardous material, it should be clear that there is a large economical potential for rescue robotics already right now. We dare to say that in the not too distant future we can even expect a diversification and simplification of rescue robots where the most basic systems are installed at every floor of every office building right next to each fire hose.

The rest of this article is structured as follows. The structure of the rescue robotics domain with its benefits and its challenges is discussed in section two. A contribution from the IUB rescue robots team in respect to the challenges discussed in the previous section, namely mapping as a core competency for world modeling, is presented in section three. Further results in respect to mapping are presented in section four where the important issue of cooperation is addressed. Section five concludes the article.

2 The challenges and the prospects of rescue robotics

As motivated in the introduction, the rescue robotics domain features very promising economical aspects. There is already a market where current systems, which are specialized high-end devices, find sufficient interest from customers who are willing to pay the necessary investments. Any further progress in technology and research in this domain will lead to a decrease in cost and larger production series. This in turn will open up additional markets, leading to further revenues and the possibility of additional investments. As already stated before, we expect a development where in the end every floor of every office building is equipped with a rescue robot as standard safety equipment much like a fire hose. In contrast to other advanced robotics domains like service robotics [16], rescue systems may be expected to face a high number of lost systems in disaster scenarios. The robots can more be considered as consumption material, even today where costly high end systems dominate the field.

One particular aspect of the rescue robotics domain eases the fruitful combination of highly challenging basic research and application oriented developments for large markets. This is the fact that rescue robots strongly benefit from autonomy while there is a human in the loop [2]. This speciality of this domain tremendously eases the transition from dumb tele-operated devices to intelligent autonomous systems. Note that for example in service robotics there is the pressure to provide autonomy right from start of the field. A vacuum cleaning robot is of no use if it has to be constantly supervised by humans. For rescue missions human rescue workers are and for a longtime still will be completely in charge of the operations. Teleoperated devices are already useful tools [17], which become more and more useful with any tiny bit of autonomy added [12]. This starts with onboard processing to enable sensor fusion and advanced perception much like in other RoboCup fields [3]. It goes over world model construction where for example learned maps can guide rescue workers to the found victims. Also, any intelligent manipulation and locomotion skill is an important added value over mere teleoperated devices. Last but not least there is a high pressure for autonomy in the rescue robotics domain due to the fact that human rescue workers are a scarce resource in disaster scenarios. Ideally, a single rescue worker should supervise a multitude of robots, which cooperate in their missions. So all the core aspect of autonomy, namely

- perception
- world modeling
- locomotion and manipulation
- cooperation

are of crucial interest in this domain. We can conclude that rescue robotics allows researchers interested in autonomous systems to start with teleoperated devices to which any contribution toward autonomy is a highly valued added benefit.

In the following sections we present results from the IUB rescue robots team, whose goal it is to develop fieldable systems within the next years. In doing so, the team has participated in several

Figure 1: The red arena at the RoboCup 2003 in Padova (left) and one of the IUB rescue robots in this arena (right).

RoboCup competitions to test their approach and to exchange their experiences with other teams in the field. The IUB rescue robot team ended on the fourth place at the RoboCup world championship in Fukuoka 2002 as well at the RoboCup world championship in Padova 2003. At the American open in New Orleans 2004, the team scored a second place. Furthermore, the IUB rescue robots received a technical award in Padova for their mapping capabilities. Based on the classification of core topics given above, the main scientific contributions of the IUB team are at the moment in the areas of world modeling and cooperation. Especially, they center around the topic of mapping, which we consider in this domain to be the foremost interest of world modeling.

3 Mapping, an example challenge for autonomy

As pointed out by others [18], mapping is one of the core problems in robotics. In the contest of rescue robotics this issue is even more important as maps are supposed to be used by rescue operators. Having a good map would enable humans to minimize the time needed to reach victims. This increases the chances of bringing them help promptly and at the same time decreases their own risk. However it has to be considered that while impressive progresses have been achieved in the field of mapping and localization [19][20][8], the problem is still far from being solved. In particular it is acknowledged that mapping deeply unstructured or dynamic environments is still open and challenging problem. The lack of structure is one they key aspects of rescue scenarios.

Subproblems like data association and feature detection become extremely difficult in these environments. In addition, the noise affecting the robot and the sensor models tends to increase. Think at the skidding experienced by a robot while trying to crawl over a pile of debris. It follows that most of the formerly developed methods cannot be applied. While developing the rescue robots we used for the RobCup Rescue competition [1] we then had to face a problem for which very few solutions have been proposed in the past. Among the different possible approaches, we have decided to adopt a system similar to the probabilistic grids introduced by Moravec [11]. The map produced by our system is called *occupancy grid*. An occupancy grid is a grid where each cell obtains *votes* indicating whether it is free or occupied. Votes are added over time but cannot grow arbitrarily. This to leave the possibility that cells previously considered free can turn to be considered occupied, or vice versa, as a consequence of a rearrangement of the environment. The mapping procedure uses the input coming from the robot odometry system and from the a cheap and not extremely precise range finder sensor. It is worth outlining that the odometry system does not just integrate the motors output, but also takes into consideration the input coming from a magnetic compass. This gives us the possibility to bound the orientation error that would arise from pure integration. Field experience showed that odometry works reasonably well as long as the robot moves along straight lines, but goes significantly bad when turning. This is no surprise, as in order to turn in place our platforms rely in the skidding of the wheels or tracks. Thus, when the robot turns, odometry data coming from the encoders are corrected with the orientation coming from the compass. This matter of localization is an important aspect, as occupancy grids based approaches rely on the assumption of precise pose estimation. We are aware this simple schema for bounding such errors is doomed to accumulate large errors in a long run. On the other hand, it turned out to work reasonably well in the short time horizon of a rescue operation.

The occupancy grid accounts for three different kinds of information, namely obstacle detected, free area detected, and no information available. This information is expressed in terms of *beliefs*. The robot starts with a completely empty grid. At every time step, the input from the range scanner is acquired. By combining this with the actual pose (x,y and orientation θ) coming from the odometry measurement subsystem, it is possible to update the beliefs of the covered grid cells. Technically, every grid cell holds an integer value, initially set to 0. This means that no information is available for that grid cell. When an obstacle is detected in the grid cell, the value is incremented. When the grid cell is determined to be free, such value is decremented. Both increments and decrements are bounded. This means that such beliefs cannot arbitrarily grow or decrease when the robot is standing at a fixed position (similar to what happens with [4]).

Figure 2 illustrates an occupancy grid produced by our robot while exploring the hall of a building. The IUB team has been the first one to introduce autonomous mapping capabilities in the Robocup Rescue 2003 competition held in Padova [7], and received for this reason a technical award from the organizers [9].

4 Multirobot mapping and the benefits of cooperation

The use of multiple robots while performing rescue operations is intriguing from many points of view. Increased system robustness, possibility to bring more and different sensors to the operation area, and decreased time to accomplish a mission are just some of them [13]. From the point of view of mapping the use of different robots is advantageous as a set of robots can visit and map different parts of the interesting area simultaneously, thus decreasing the mission time. On the other hand, as robots are likely to spread and explore different parts, the problem of combining different maps together arises. Multirobot mapping and localization gained quite some attention in the past [14][15][19]. But, as pointed

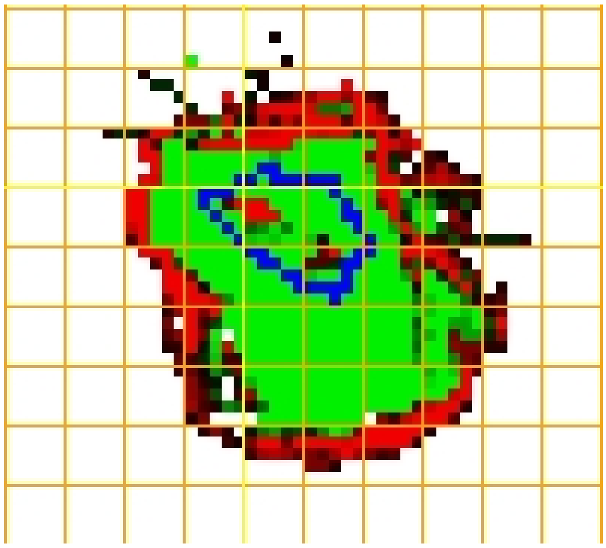

Figure 2: An occupancy grid produced by the robot while exploring the hall of a building. Green colored cells are cells for which the majority of votes indicates it is free. Red color is for cells where the voting system indicates it is occupied. White cells indicate lack of information, or an equal number of opposite votes. Color intensity is proportional to the difference of votes in one direction or the other.

out by others in [10], the problem of map merging has received much less attention.
In our system robots explore the environment and send sensed data back to the operator station where the different maps are produced and updated. Each robot is completely unaware of the presence of other robots, i.e. they do not know neither the initial positions nor how many robots are present. When the operator decides that enough information has been acquired all the different maps are to be merged into a single one to be provided to the rescue personnel who will enter the scene. Merging maps together can be seen as an optimization problem [5] [6]. While the mathematical details and proofs of are provided in the previous references, we here provide the an intuitive description of the algorithm we developed. Given two maps, say M_1 and M_2, the goal is to find a transformation, i.e. a translation and a rotation, that will lead to a good overlapping between the two maps. It is worth outlining that rotation is indeed needed as robots enter the operative scenario from different points and have different reference systems not only in terms of x and y but also in terms of orientation. *Good overlapping* is measured in terms of a *dissimilarity function*, i.e. a function that indicates how much two maps can be overlapped. Ideally, dissimilarity will be 0 for two maps that are completely overlapping and will grow more and more as one tries to superimpose increasingly different maps. It is then evident that the map merging problem can be seen as an optimization problem, i.e. it is required to determine the transformation that minimizes the overall dissimilarity. Multi-point hill climbing or simulated annealing are two approaches that can be used to solve the depicted problem. However we developed a new algorithm that turns out to be faster and more accurate. In addition, the previous two can be seen as specific cases of it. The algorithm performs a random walk in the pace of possible transformations. At each step a new candidate point is

generated according to a Gaussian transformation. The candidate transformation is accepted or refused according to the associated dissimilarity function. After a new candidate sample has been accpeted or refused, the algorithm updates the distribution's mean and variance, then the name adaptive random walk.

Algorithm 1 Random walk

Require: $numSteps \geq 0$

1: $k \leftarrow 0, \quad t_k \leftarrow$ Initial random transformation
2: $\Sigma \leftarrow \Sigma_{init}, \quad \mu \leftarrow \mu_{init}$
3: $c_0 \leftarrow$Initial dissimilarity between M_1 and M_2
4: **while** $k < numSteps$ **do**
5:     Generate a new sample $s \leftarrow t_k + v_k$
6:     $c_s \leftarrow$ Dissimilarity induced by the new transformation s
7:     **if** $c_s < c_k$ OR $RS(t_k, s) = s$ **then**
8:         $k \leftarrow k+1, \quad t_k \leftarrow s, \quad c_k = c_s$
9:         Update μ and Σ
10:     **else**
11:         discard the sample s
12:     **end if**
13: **end while**

Algorithm 1 provides an algorithmic sketch of the described procedure. The element v_k at line 5 indicates a random sample generated according to a Gaussian distribution with mean μ and covariance matrix Σ. t_k is the current point reached by the random walk while and c_k is its own transformation. Finally, RS (line 7) is a random selector that allows to randomly accept a given transformation even if it does not bring an advantage in terms of dissimilarity function.

It can be formally proved that the proposed algorithm will approximate the optimal solution with arbitrary precision under rather mild conditions.

Figure 4 illustrates an example of maps gathered by two different robots mapping the same hallway starting from two different points and the obtained result. Subfigure d shows the trend of the dissimilarity function while exploring the space of possible transformation. The characteristic pattern of the dissimilarity function can be observed, with repeated descent sequences alternated with jumps to new staring points.

5 Conclusion

Rescue robotics is an important domain from which we can expect significant scientific contributions toward the development of autonomous systems. It features a tremendous application potential while facilitating the investigation of basic research topics. We presented work from the IUB rescue robots

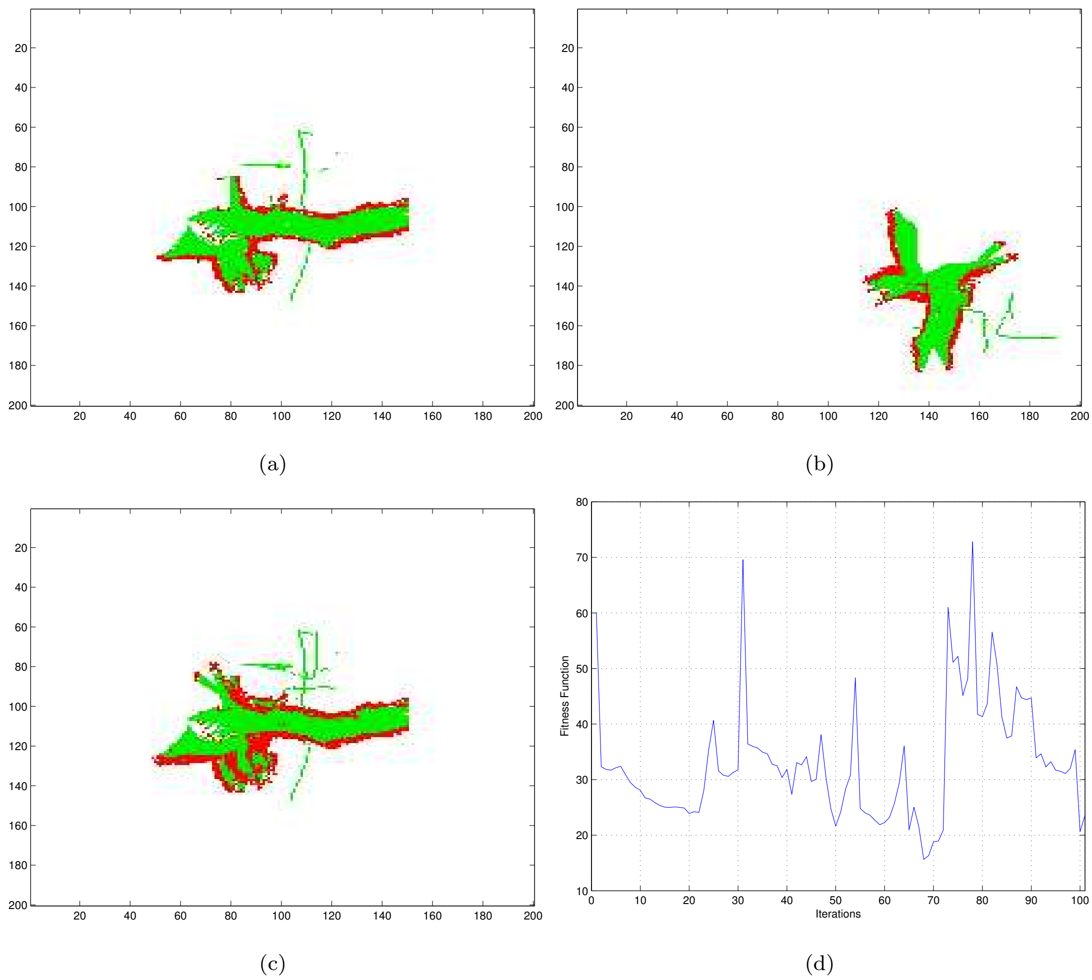

Figure 3: Subfigures a and b illustrate the maps created by two robots while exploring two different parts of the same environment. To make the matching task more challenging the magnetic compass and the odometry system were differently calibrated. Subfigure c shows the best matching found after 200 iterations of the search algorithm and subfigure d plots the trend of the dissimilarity function during the exploration

team which illustrates some approaches in this direction. Concretely, results from mapping as a core form of world modeling in this unstructured and dynamic environment are presented. Furthermore, we addressed the important issue of cooperation by discussing a novel approach to multi-robot mapping.

REFERENCES

[1] A. Birk, S. Carpin, and H. Kenn. The IUB 2003 rescue robot team. In *Robocup 2003*. Springer, 2003.

[2] Andreas Birk and Holger Kenn. A control architecture for a rescue robot ensuring safe semi-autonomous operation. In Gal Kaminka, Pedro U. Lima, and Raul Rojas, editors, *RoboCup-02: Robot Soccer World Cup VI*, LNAI. Springer, 2002.

[3] Andreas Birk, Holger Kenn, and Thomas Walle. On-board control in the robocup small robots league. *Advanced Robotics Journal*, 14(1):27 – 36, 2000.

[4] S. Biswas, B. Limketkai, S. Sanner, and S. Thrun. Towards object mapping in dynamic environments with mobile robots. In *Proceedings of the IEEE/RSJ International Conference on Intelliegent Robots and Systems*, pages 1014–1019, 2002.

[5] S. Carpin and A. Birk. Stochastic map merging in rescue environments. In *Robocup 2004*. Springer, 2004. Accepted for pubblication.

[6] S. Carpin, V. Jucikas, and A. Birk. Multi-robot mapping for rescue robotics. In *Proceedings of the 2004 international workshop on safety, security and rescue robotics.*

[7] S. Carpin, H. Kenn, and A. Birk. Autonomous mapping in the real robot rescue league. In *Robocup 2003*. Springer, 2003.

[8] G. Dissanayake, P. Newman, S. Clark, H.F. Durrant-Whyte, , and M. Csorba. A solution to the simultaneous localisation and map building (slam) problem. *IEEE Transactions of Robotics and Automation*, 17(3):229–241, 2001.

[9] A. Jacoff, B. Weiss, and E. Messina. Evolution of a performance metric for urban search and rescue (2003). In *Performance Metrics for Intelligent Systems*, 2003.

[10] K. Konolige, D. Fox anx B. Limketkai, J. Ko, and B. Steward. Map merging for distributed robot navigation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 212–217, 2003.

[11] H.P. Moravec. Sensor fusion in certainty grids for mobile robots. *AI Magazine*, 9(3):61–74, 1988.

[12] M. Micire R. Murphy, J. Casper and J. Hyams. Potential tasks and research issues for mobile robots in robocup rescue. In Tucker Balch Peter Stone and Gerhard Kraetszchmar, editors, *RoboCup-2000: Robot Soccer World Cup IV*, Lecture Notes in Artificial Intelligence 2019. Springer Verlag, 2001.

[13] L.E. Parker. Current state of the art in distributed autonomous mobile robots. In L.E. Parker, G. Bekey, and J.Barhen, editors, *Distributed Autonomous Robotic Systems 4*, pages 3–12. Springer, 2000.

[14] L.E. Parker, K. Fregene, Y. Guo, and R. Madhavan. Distributed heterogeneous sensing for outdoor multi-robot localization, mapping, and path planning. In A. Schultz, editor, *Multi-Robot Systems: From Swarms to Intelligent Automata*, pages 21–30. Kluwer, 2002.

[15] S.I. Roumeliotis and G. Bekey. Distributed multirobot localization. *IEEE Transactions on Robotics and Automation*, 18(5):781–795, 2002.

[16] Rolf Dieter Schraft and Gernot Schmierer. *Service Robotics*, A K Peters, 2002.

[17] Rosalyn Graham Snyder. Robots assist in search and rescue efforts at WTC. *IEEE Robotics and Automation Magazine*, 8(4):26–28, December 2001.

[18] S. Thrun. Robotic mapping: A survey. In G. Lakemeyer and B. Nebel, editors, *Exploring Artificial Intelligence in the New Millenium*. Morgan Kaufmann, 2002.

[19] S. Thrun, W. Burgard, and D. Fox. A real-time algorithm for mobile robot mapping with applications to multi-robot and 3d mapping. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 321–328, 2000.

[20] S. Thrun, D. Fox, W. Burgard, and F.Dellart. Robust monte carlo localization for mobile robots. *Artificial Intelligence*, 2001.