4th International Conference on Intelligent Robotics and Applications (ICIRA), Aachen, Germany, December 2011

Exploration Strategies for Building Compact Maps in Unbounded Environments

Matthias Nieuwenhuisen[1], Dirk Schulz[2], and Sven Behnke[1]

[1] Autonomous Intelligent Systems Group, University of Bonn, Germany
nieuwenh@ais.uni-bonn.de, behnke@cs.uni-bonn.de
[2] Fraunhofer Institute for Communication, Information Processing and Ergonomics FKIE, Wachtberg, Germany
dirk.schulz@fkie.fraunhofer.de

Abstract. Exploration strategies are an important ingredient for map building with mobile robots. The traditional greedy exploration strategy is not directly applicable in unbounded outdoor environments, because it decides on the robot's actions solely based on the expected information gain and travel cost. As this value can be optimized by driving straight into unexplored space, this behavior often leads to degenerated maps. We propose two different techniques to regularize the value function of the exploration strategy, in order to explore compact areas in outdoor environments. We compare exploration speed and compactness of the maps with and without our extensions.

Keywords: Outdoor Exploration, Compact Map-Building

1 Introduction

The autonomous exploration of unknown indoor and outdoor environments is an active research topic in mobile robotics. In the literature, solutions to special instances of the problem have been proposed over the last years. The overall goal of any exploration approach is to plan a robot's motion in such a way that a map of the robot's environment can be build efficiently, solely based on the robot's sensor data. Most of the approaches developed so far, decide on the robot's motions based on a trade-off between the expected information gain, e. g. the expected increase in newly observed area, and the cost of the exploration. In bounded spaces, like buildings or predetermined areas of interest, greedy motion strategies based on this idea exist that fully explore the environment with low (travel) costs (cf. [8], [6]). If the robot operates in a virtually unbounded outdoor environment, like it is the case e. g. in planetary exploration, it is generally not feasible to explore the whole reachable area. A strategy that greedily maximizes the information gain can lead to degenerated maps in this situation, as it is often possible to obtain a maximum of new information by stubbornly driving on a straight line away from the starting position. The resulting map cannot provide useful information about the local shape or the topology of the environment. For this reason, some additional means are required to constrain the robot's motion, while it explores an open outdoor environment. One simple possibility to achieve this is to artificially bound the area to explore, e. g. by manually defining a bounding polygon. This approach can lead to problems as the topology of the environment is a priori unknown, e. g. with concave inclusions. Although it is possible to reach all free space parts contained in the bounding polygon, it may require extensive exploration to discover a path in more complex scenarios.

In this article, we propose two different regularization techniques that – in combination with a greedy exploration strategy – allow to acquire *compact* maps of the environment. In the scope of this article, *compact* maps are maps with a width-to-height ratio close to one. We achieve this goal by introducing additional terms into the value function of the underlying decision-theoretic motion planner that decides on the robot's actions. The first regularization, called spiral exploration, biases the value function towards favoring motion commands that lead to a spiral exploration path, if possible. The second one, called distance-penalized exploration, directly penalizes motion commands that lead to maps with an unbalanced width-to-height ratio.

The remainder of the article is organized as follows: after discussing related work in the next section, we will briefly introduce the basics of greedy decision-theoretic exploration approaches in Section 3. Section 4.1 then introduces the spiral exploration extension, followed by Section 4.2 that explains the distance-penalized exploration technique. We present simulations in Section 5, where we evaluate the performance of both exploration approaches using experiments in environments of varying structure.

2 Related Work

A major element of exploration systems is to find optimal sensing positions where a robot can gather as much new information about the world as possible, also known as next best views. As the world is unknown a priori, strategies to estimate the next best positions are needed. A widely used approach in 2D exploration is to choose sensing positions at the frontier between known and unknown space and to weight them according to the travel costs to reach these positions [15]. Burgard, Moors, and Schneider described a strategy to coordinate frontier-based exploration in multi-robot systems [1]. González-Baños and Latombe use the maximal visible unexplored space at a scan position as an approximation of the art gallery problem [4]. Surmann, Nüchter, and Hertzberg applied this strategy in a 3D exploration system [12]. Pfaff et al. use virtual laser scans, simulated by 3D ray-casting, to calculate the expected information of scan positions to build a 2.5D world model [11]. Although our approach is similar to these techniques, as it uses a greedy exploration strategy, none of these works addresses the problem of building compact maps in unbounded environments. Several geometrical approaches for following motion patterns to cover unknown areas have been proposed for this purpose, e. g. Cao et al. [2] and Hert et al. [5]. They differ from our approach as their main objective is to cover a given bounding polygon completely. In these approaches, the expected information gain of scan positions is not evaluated. Moorehead [9] proposed a decision-theoretic exploration approach for planetary surfaces that uses additional metrics to guide the robot to areas of scientific interest, like dry waterbeds. In contrast to that, we are interested in maximizing the known area in a way that enables the robot to plan efficient paths afterward. Topological information is used by Morris et al. [10] to explore abandoned mines. We use an implicit selection of the next view positions, as this does not rely on any assumptions on the environment topology.

3 Greedy Exploration

The compact exploration strategies, we propose in this article, are an extension to the greedy exploration approach for autonomous map building. We briefly introduce this underlying approach first, following the notation used in [13].

The goal of greedy exploration is to control the robot's next measurement action in such a way that its uncertainty about the global state of the (static) world is minimized. Formally, this knowledge is represented by a probability distribution over possible world states that is called the robot's belief b. In the context of autonomous map building, the robot acquires information about the shape of the environment using, for example, laser range sensors. Here, the next action to be planned is a motion to a new location u, where a maximum of additional information about the shape of the environment is expected. This *information gain* can be expressed by the expected reduction of the entropy of the robot's belief after integrating a measurement at the new location u, i. e.

$$I_b(u) = H_b(x) - E_z\left[H_b(x'|z,u)\right], \tag{1}$$

where $H_b(x)$ denotes the entropy of the prior belief, and $E_z\left[H_b(x'|z,u)\right]$ is the expected conditional entropy of the belief after carrying out the motion and integrating a new sensor measurement z obtained at the new location. A greedy exploration strategy is now a decision-theoretic approach to choose the action $\pi(b)$ with the best trade-off between information gain and expected travel cost using the value function

$$\pi(b) = \mathrm{argmax}_u\, \alpha I_b(u) + \int r(x,u)b(x)dx. \tag{2}$$

Here, $r(x,u)$ is the travel cost function and the factor α weighs the influence of gain and costs. Note that in the context of mapping, a probabilistic map representation, like occupancy probability grids [3] or multi-level surface maps [14], is generally used to represent the belief.

In order to achieve compact exploration, we add an extra cost term $C(b,u)$ to the value function $\pi(b)$ that penalizes motion commands which are likely to lead to non-compact maps. The resulting value function used is

$$\pi_C(b) = \mathrm{argmax}_u\, \alpha\left(\beta C(b,u) + I_b(u)\right) + \int r(x,u)b(x)dx. \tag{3}$$

The map representation of the belief allows to efficiently approximate the integral for the expected travel costs in Eq. 3 using value iteration (see [13] chapter 17 for details). $I_b(u)$ is generally determined by estimating the expected amount of newly explored area at the location reached after carrying out the motion command u. The new area can be estimated by determining the expected visibility polygon at a candidate view position and subtracting the already explored part of it. The expected visibility polygon can, for example, be computed using a ray sweep algorithm that treats unexplored space as free [4]. If the belief is represented by grid maps, virtual laser scans can alternatively be used. Beams are sent out from a candidate view position and followed through the grid. The estimation of yet unexplored space is reduced to counting the traversed explored and unexplored grid cells. As the simulation can mimic the characteristics of a real laser-range finder, the estimate is generally better than the one based on the visibility polygon.

It remains to describe, how we achieve an exploration behavior that leads to compact maps. For this purpose, we introduce two different utility or cost functions $C(b, u)$ for weighing candidate view positions u, each one leading to a different compact exploration behavior.

4 Compact Exploration Strategies

4.1 Spiral Exploration

In environments without obstacles, a simple exploration technique that leads to a compact map is to drive on a spiral path, with a radial distance between two turns that maximizes the information gain. In outdoor environments with a low to average obstacle density, a spiral trajectory approach can still be reasonable, although obstacles may occasionally block the robot's path. In such situations, the robot obviously needs to deviate from the spiral in a way that still leads to an efficient compact exploration. In order to achieve such a behavior, we implement $C(b, u)$ as a function that rewards motion commands leading to spiral trajectories. For this purpose, we keep track of the center of the map built so far, and we determine the angle θ between the radial vector $\boldsymbol{r}$ that connects the center of the map with the robot's current position and the motion vector $\boldsymbol{v}$ that connects the robot's current position with its intended scan position u. To achieve a spiral trajectory, $\boldsymbol{v}$ should be kept orthogonal to $\boldsymbol{r}$, because the motion direction is tangential to the map's border then. For this reason, we choose to reward considered scan positions with a value proportional to $|\sin(\theta)|$. This procedure is summarized as pseudo code in Algorithm 1. Integrated into the value function $\pi_C(b)$, the robot tends to explore the environment on a spiral path, but the robot starts to deviate from the spiral, if view points exist that, depending on the weighting constants α and β, lead to a sufficiently higher information gain with tolerable motion costs. This will for example happen, if an obstacle blocks the robot's path. To ensure a trajectory circling the convex hull of the map equally close in all directions, a floating center of the map is maintained during exploration by continuously computing its center of mass.

Algorithm 1: Compute_spiral_exploration_cost

```
Input: map, scanPos, robotPos
com ← CenterOfMass (map)
robotToScan ← Normalize (scanPos - robotPos)
robotToCom ← Normalize (com - robotPos)
sinAngle ← sin (arccos (cosAngle))
return Abs (sin (arccos (cosAngle)))
```

4.2 Distance-Penalized Exploration

The spiral exploration strategy effectively leads to compact maps in environments known to have a low obstacle density like, e. g., in planetary exploration. However, with an increasing number of obstacles present in the environment, the spiral strategy tends to leave unexplored holes inside the convex hull of the already explored area. The information gain received by closing such holes is often small, compared to the costs of leaving the tangential exploration direction and the additional travel costs. In this section, we propose an exploration strategy that effectively avoids leaving such holes at the cost of a potentially lower exploration speed. As in the case of the spiral exploration strategy, this is achieved by introducing a cost term into the greedy value function, without explicitly considering to close exploration holes.

To mitigate the discrepancy between the lower information gain received at holes and the high information gain at the frontier to the unexplored space, we penalize view positions depending on their distance from the center of the map. We define the radius of a map as the maximum distance of an explored point from the center of the map. Every view point that potentially expands this radius is penalized. The penalty is proportional to the maximum difference between the current radius of the map and the potential new radius of the map after a range scan taken at this position has been integrated.

In the case of a balanced width-to-height ratio, expansions in every direction are equally penalized and the expected information gain dominates the calculation of the value function. Expansions in directions in which the map has not yet reached its maximum radius are not penalized, even if there are holes left. For this reason, holes inside the explored map have to provide a specific minimum gain to be chosen as an area to explore. Otherwise, the quality of the explored maps can decrease. Hence, the distance penalization criterion resembles our definition of a compact map. The computation of costs for distance-penalized view point selection is described in Algorithm 2.

5 Evaluation

We implemented both strategies for compact exploration in the context of a system for autonomously acquiring three-dimensional environment models using a 3D laser range scanner. All experiments used to get quantitative results are

Algorithm 2: Compute_distance_penalized_exp._cost

```
Input: map, maxRange, scanPos
com ← CenterOfMass (map)
direction ← Normalize (scanPos - com)
maxDistBefore ← 0
foreach exploredCell in map do
    maxDistBefore ← Max (maxDistBefore, Distance (exploredCell,com))
end
maxDistAfter ← Distance (com, scanPos + direction * maxRange)
if maxDistAfter > maxDistBefore then
    return 1 - α∗(maxDistAfter - maxDistBefore)
else
    return 1
end
```

performed in Gazebo [7], the 3D simulation environment from the Player/Stage project. The 3D scanner simulated is a SICK LMS200 scanner, mounted on a turn-table in such a way that it continuously receives vertical 2D scans while it rotates. The horizontal angular resolution of the 3D scan was set to approximately 1 degree by adjusting the turn-table speed.

Following ideas of Triebel et al. [14], our system represents 3D models of the environment by multi-level surface maps, which are grid maps that can store multiple height values, called *surface patches*, representing surfaces at different heights, per cell. Each patch has a vertical extent and can therefore represent man-made structures like walls and floors with a memory complexity comparable to occupancy grid maps. As the traversable plane reachable by a robot in such a map is a 2-manifold, 2D planning algorithms can still be used.

Candidate view positions are determined with inverse ray-casting from up to 10.000 frontier points. To avoid candidate positions caused by very small frontiers, only clusters of five or more frontier points are considered during the ray-casting operation. From every considered frontier point, 100 beams are followed through the environment. As nearby candidate points are considered to gain a similar amount of information, only candidate points exceeding a distance threshold to neighboring points are added. For the information gain calculation, only patches not exceeding a distance threshold to the robot are considered. Due to the fixed angular distance between consecutive laser beams, the Cartesian distance of the endpoints hitting the ground plane exceed the size of the surface patches. Hence, they are not connected.

In the following, we compare the proposed compact exploration strategies with the plain greedy exploration strategy. As 3D range measurements are costly operations, in our setup the exploration speed is measured as the number of surface patches compared to the number of scans taken. The compactness of a map is the number of explored patches in comparison to the total number of grid cells in the map; this includes unexplored cells, because the size of the grid is chosen as the bounding rectangle of all laser scans.

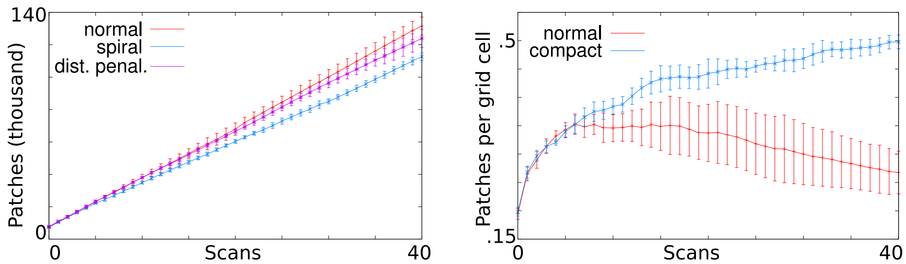

(a) Number of patches in the map for different exploration strategies.

(b) Patches per grid cell for different exploration strategies.

Fig. 1. In Experiment 1 the non-compact and compact strategies have nearly equal exploration speed, but the ratio of patches per grid cell is significantly higher when using the compact strategies.

5.1 Experiment 1: Low Obstacle Density

In the first experiment, we chose an environment with a low obstacle density, in order to compare the maximum exploration speed of the strategies with and without the proposed extensions. The exploration is stopped after 40 3D scans. The trajectory of the non-compact exploration consists of nearly linear paths as these maximize the ratio between newly seen environment and travel costs. As the probability of a change of exploration direction caused by obstacles is low in this experiment, most of these changes are caused by the randomized sampling of candidate scan positions.

The spiral exploration strategy is able to gather compact maps comparable to the maps generated by the distance-penalized strategy. This is explained by the fact, that a spiral exploration extends the map uniformly in every direction without causing high travel costs, as long as the robot is not deflected by an obstacle. Therefore, the maximum distance of a patch to the center of the map is increasing. Figure 1 shows that the non-compact exploration exhibits a slightly higher exploration speed compared to the two compact strategies, which show no significant difference in exploration speed in maps with low obstacle density. During the first scans, there is no significant difference between the strategies, until the robot changes its exploration direction for the first time. The patches per grid cell-ratio decreases for the non-compact strategy then, because the number of occupied cells can only increase linearly with the distance traveled, while the total number of grid cells increases quadratically. This effect is typical for a non-compact strategy. For the compact strategies, the occupancy ratio continues to increase, because the robot's motions between measurements only have a small influence on the size of the bounding rectangle of the map during a compact expansion. This is also visible in a lower variance in the data averaged over several test runs.

5.2 Experiment 2: High Obstacle Density

The purpose of the second and third experiment is to evaluate the exploration behavior of the strategies in an unbounded outdoor environment with a high obstacle density like, e. g., urban environments. We evaluated the strategies in two types of outdoor environments. The first one is an environment containing mainly smaller obstacles, i. e. obstacles that can be passed without long detours, while the second type of environment contains many elongated obstacles, which force the robot to drive longer detours. All the experiments are stopped after integrating 100 scans into the map.

Typical trajectories of the three strategies for Experiment 2 with the convex obstacles are shown in Figure 4a. The non-compact exploration strategy leaves the area where the obstacles are and moves into wide open space, leading to a degenerated map. This effect is typical for unconstrained greedy exploration. In order to maximize information gain, it tends to move away from obstacles, which occlude large portions of the scan. The proposed strategies, instead, direct the robot through passages between obstacles, in order to reach the goal of compact exploration.

In most cases, the distance-penalized strategy reached an exploration performance comparable to the spiral exploration strategy. For this reason, we combined the two compact strategies in Figure 2. Figure 2a shows that the non-compact exploration strategy is able to gather more new information per scan than the compact extensions. The effect is more pronounced than in Experiment 1, where the exploration speeds were nearly equal. This is explained by scan positions close to obstacles, which are avoided by the non-compact exploration strategy. The compact strategies tend to surround obstacles instead.

5.3 Experiment 3: Elongated Obstacles

The second type of outdoor environment, evaluated in Experiment 3, contains many elongated obstacles. These obstacles constrain the possible driving directions to the corridors between obstacles frequently. For the spiral exploration strategy, this can lead to problems, e. g. elongated obstacles orthogonal to the spiral direction force the robot to increase the exploration radius without being able to follow the spiral. After passing an obstacle, the robot follows the spiral again, now with the increased radius. This can lead to significant holes in the convex hull of the explored space and a less compact map.

Example trajectories for the two compact strategies are shown in Figure 4b, maps gathered with these strategies in Figure 5. The spiral exploration proceeds with a too large radius once it leaves the area where many obstacles are. As shown in Figure 3, the distance-penalized strategy achieves more compact maps in this type of environment, when compared to the spiral exploration strategy, at the cost of a slightly lower exploration speed.

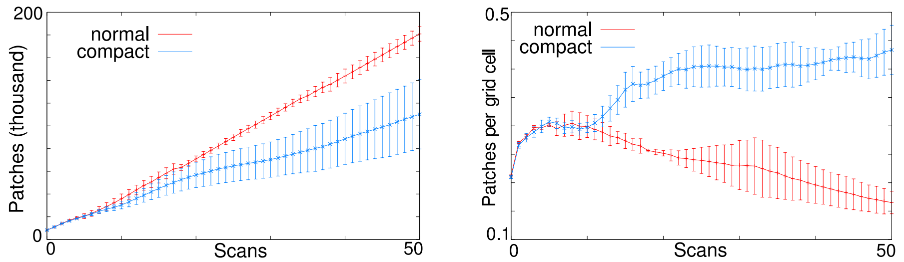

(a) In Experiment 2, the exploration speed of the non-compact exploration is much higher than in the compact case.

(b) The non-compact strategy leads in Experiment 2 to sparse maps, the compact strategies strategies fills the map more homogeneously.

Fig. 2. In an environment with a high density of small objects, the exploration speed of the compact strategies is lower than the speed without compact extensions, but the compactness is significantly higher.

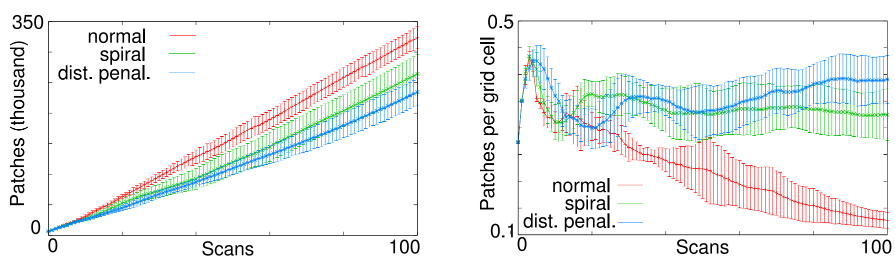

(a) In contrast to the experiments before, the exploration speed of the two compact strategies is not similar. The speed of the distance-penalized exploration is lower than the speed of the spiral exploration.

(b) As the selection of motion commands is constrained by obstacles over long parts of the trajectory, spiral exploration leaves holes resulting in less compact maps.

Fig. 3. In environments with elongated obstacles, our proposed compact strategies perform differently. Compactness has to be traded-off against exploration speed.

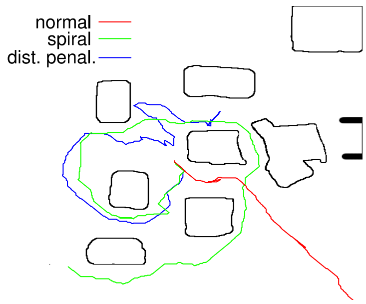

(a) Experiment 2. The edge length of obstacles is chosen to be small compared to the robot's sensing range here. Both compact strategies perform equally well.

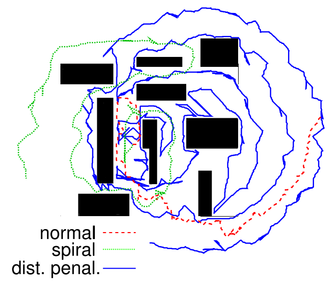

(b) Exploration with elongated obstacles. The long obstacles force the spiral strategy to increase its exploration radius too fast.

Fig. 4. Example trajectories of the exploration strategies.

6 Conclusions

In this article, we have shown that greedy exploration strategies, which only maximize the ratio between expected information gain and travel costs, are not directly applicable in unbounded outdoor environments, because they can lead to degenerated non-compact maps.

To produce compact maps that expand around a point of interest, additional control of the exploration is needed. For this purpose, we proposed two different extensions to greedy exploration strategies: spiral exploration and distance-penalized exploration. These regularizations extend the value function of a greedy strategy with an additional term rewarding sensing positions leading to more compact maps.

In simulation experiments, we compared the exploration speed between a greedy exploration strategy with and without these extensions resulting in the conclusion that the compact strategies are able to reach an exploration speed comparable to the plain greedy strategy in environments with a low to medium obstacle density. In environments with a high obstacle density, compactness has to be traded-off against exploration speed. Spiral exploration is suitable in environments where approximate spiral trajectories are possible. This is the case when the obstacles in the environment are nearly convex and small, compared to the sensing range. The distance-penalized exploration is able to produce maps that are compact, without relying on environmental characteristics, but possibly at a lower exploration speed. Hence, if the environment can be expected to contain obstacles relatively small compared to the measurement range a spiral shaped exploration should be performed. Otherwise, the distance-penalized exploration strategy leads to compact maps without prior knowledge about the area to explore.

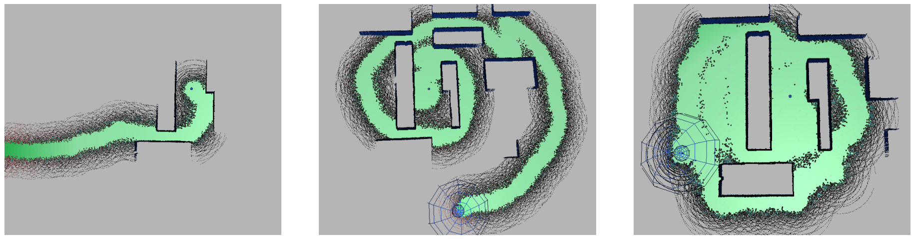

(a) The robot following the frontier-based exploration strategy leaves the area with obstacles quickly, leading to a map not suitable for navigation.

(b) The area with obstacles is surrounded with the spiral exploration strategy, but unexplored enclosures are left in the case of elongated obstacles.

(c) Distance-penalized exploration resembles the compactness criterion, inducing compact maps in arbitrary environments.

Fig. 5. Resulting maps of frontier-based exploration and our proposed strategies after 170 consecutive 3D scans. The green area is traversable, the gray area is unexplored.

References

1. Burgard, W., Moors, M., Schneider, F.: Collaborative Exploration of Unknown Environments with Teams of Mobile Robots. In: Dagstuhl. pp. 52–70 (2002)
2. Cao, Z., Huang, Y., Hall, E.: Region Filling Operations with Random Obstacle Avoidance for Mobile Robots. Journal of Robotic systems 5(2) (1988)
3. Elfes, A.: Using Occupancy Grids for Mobile Robot Perception and Navigation. IEEE Computer 22, 46–57 (1989)
4. González-Baños, H.H., Latombe, J.C.: Navigation Strategies for Exploring Indoor Environments. Intl. J. of Robotics Research 21(10-11), 829–848 (2002)
5. Hert, S., Tiwari, S., Lumelsky, V.: A Terrain-covering Algorithm for an AUV. Autonomous Robots 3(2), 91–119 (1996)
6. Holz, D., Basilico, N., Amigoni, F., Behnke, S.: Evaluating the Efficiency of Frontier-based Exploration Strategies. In: Proceedings of ISR/ROBOTIK (2010)
7. Koenig, N., Howard, A.: Design and Use Paradigms for Gazebo, An Open-Source Multi-Robot Simulator. In: Proc. of IROS (2004)
8. Koenig, S., Tovey, C., Halliburton, W.: Greedy Mapping of Terrain. In: Proc. of ICRA (2001)
9. Moorehead, S.: Autonomous Surface Exploration for Mobile Robots. Ph.D. thesis, Carnegie Mellon University (2001)
10. Morris, A., Silver, D., Ferguson, D., Thayer, S.: Towards Topological Exploration of Abandoned Mines. In: Proceedings of ICRA. pp. 2117–2123 (2005)
11. Pfaff, P., Kümmerle, R., Joho, D., Stachniss, C., Triebel, R., Burgard, W.: Navigation in Combined Outdoor and Indoor Environments using Multi-Level Surface Maps. In: WS on Safe Navigation in Open and Dynamic Environments, IROS '07
12. Surmann, H., Nüchter, A., Hertzberg, J.: An Autonomous Mobile Robot with a 3D Laser Range Finder for 3D Exploration and Digitalization of Indoor Environments. Robotics and Autonomous Systems 45(3-4), 181–198 (2003)
13. Thrun, S., Burgard, W., Fox, D.: Probabilistic Robotics. MIT Press (2005)
14. Triebel, R., Pfaff, P., Burgard, W.: Multi-Level Surface Maps for Outdoor Terrain Mapping and Loop Closing. In: Proc. of IROS (2006)
15. Yamauchi, B.: A Frontier-based Approach for Autonomous Exploration. In: Proceedings of CIRA. pp. 146–151 (1997)