

A comparative evaluation of exploration strategies and heuristics to improve them

Dirk Holz* Nicola Basilico† Francesco Amigoni† Sven Behnke*

*Autonomous Intelligent Systems Group, University of Bonn, Bonn, Germany

†Dipartimento di Elettronica e Informazione, Politecnico di Milano, Milano, Italy

Abstract—Exploration strategies play an important role in influencing the performance of an autonomous mobile robot exploring and mapping an unknown environment. Although several exploration strategies have been proposed in the last years, their experimental evaluation and comparison are still largely unaddressed. In this paper, we quantitatively evaluate exploration strategies by experimentally comparing, in a simulation setting, a representative sample of techniques taken from literature. From a broader perspective, our work also contributes to the development of good experimental methodologies in the field of autonomous mobile robotics by promoting the principles of comparison, reproducibility, and repeatability of experiments.

Index Terms—Robotic exploration. Exploration strategies.

I. INTRODUCTION

Exploration and mapping are fundamental tasks for autonomous mobile robots operating in unknown environments. Recent work [1] showed that *exploration strategies* largely influence the efficiency with which a robot performs exploration. Broadly speaking, an exploration strategy drives a robot within a partially known environment, determining where to acquire new spatial information. In this work, we focus on the mainstream approach of greedy *Next-Best-View* (NBV) exploration strategies. When employing a NBV strategy, exploration comes down to a sequence of steps where, at each step, a number of candidate observation locations are evaluated according to some objective function and the best one is selected for the robot to reach. Several exploration strategies [2], [3], [8], [16], [19] have been proposed, but their experimental evaluation and comparison constitute a topic that is still largely unaddressed, with few exceptions (e.g., [1] and [12]). In our opinion, filling this gap is an important issue in mobile robot exploration.

In this paper, we aim at contributing to the assessment of the experimental comparison between exploration strategies. In particular, we compare three exploration strategies for a single robot [5], [8], [11] that are a representative sample of the current state of the art and we investigate the reasons for their different performance and the ways in which they can be improved. The original contribution of this paper is not in the proposal of new exploration strategies, but in presenting some insights derived from the quantitative experimental evaluation of both some strategies and some general heuristics that can be used to improve them. These insights can be intended as enabling factors for more complex exploration applications and for developing better exploration strategies. Our work extends the results of [1] by comparing a different set of

strategies within a more realistic simulation framework and by presenting new insights. Furthermore, we extend the work in [9] by also evaluating the heuristic improvements when applied to different exploration strategies.

Our work can be also viewed from the general perspective of the definition of good experimental methodologies for autonomous mobile robotics (for instance, see [6] and [14]). Recent efforts have recognized that experimentation in this field has not yet reached a level of maturity comparable with that reached in other engineering and scientific fields [4]. Among the elements that define a good experimental methodology is the *comparison* of experimental results. With this paper, we contribute toward the definition of a framework for evaluating exploration strategies in different setups. We conduct our comparison in simulation, since it enables performing reproducible and repeatable experiments [4]. *Reproducibility* is the possibility to verify, in an independent way, the results of a given experiment. Other experimenters, different from the one claiming for the validity of some results, should be able to achieve the same results, by starting from the same initial conditions, using the same type of instruments, and adopting the same experimental techniques. *Repeatability* concerns the fact that a single result is not sufficient to ensure the success of an experiment. A successful experiment must be the outcome of a number of trials, performed at different times and in different places. These requirements guarantee that the result has not been achieved by chance, but is systematic. Performing experiments using a standard and publicly available simulation platform (like Player/Stage) is a way to promote comparison, reproducibility, and repeatability of experiments.

II. RELATED WORKS

The definition of strategies for autonomous exploration of environments has been addressed by several works in literature. Besides exploration strategies that make the robots move along predefined trajectories [13] and that attempt to close loops for localization purposes [17], the mainstream approach considers exploration as an incremental process in which the next observation location is selected among a set of candidates on the basis of available information. These *Next-Best-View* (NBV) systems evaluate the candidate observation locations according to some criteria. Usually, in NBV systems, candidate locations are on the frontier between the known free space and the unexplored part of the environment and are reachable from the current position of the robot [21]

(an exception is the feature-based approach of [15]). The exploration strategies analyzed in this paper follow the NBV approach. NBV problems have been also studied in Computer Vision and Graphics. However, the proposed techniques do not apply well to mobile robots [8].

In evaluating a candidate location, single or multiple criteria can be used. For example, [21] presents a strategy that uses a single criterion, the *traveling cost*, according to which the best observation location is the nearest one. Other approaches combine traveling cost with different criteria, for example with *expected information gain* [8]. This criterion is related to the expected amount of new information about the environment obtainable from a candidate location. It is estimated by measuring the area of the portion of unknown environment potentially visible from the candidate location, taking into account the so-far built map and the robot's sensing range. Other examples of combining different criteria are [16], in which the traveling cost is linearly combined with the expected reduction of the uncertainty of the map after the observation, and [2], in which a technique based on relative entropy is used. In [19], several criteria are employed to evaluate a candidate location: traveling cost, uncertainty in landmark recognition, number of visible features, length of visible free edges, rotation and number of stops needed to follow the path to the location. They are combined in a multiplicative function to obtain a global utility value. The above strategies are based on *ad hoc* aggregation functions (linear combination, multiplication, ...) that combine criteria. In [3], the authors dealt with this problem and proposed a more theoretically-grounded approach based on multi-objective optimization, in which the best candidate location is selected on the Pareto frontier. In [3], besides traveling cost and expected information gain, also *overlap* is taken into account. This criterion is related to the amount of already known features that are visible from a candidate location. It accounts for the precision of self-localization of the robot: the larger the overlap, the better the localization of the robot.

III. EXPERIMENTAL SETTING

We now introduce our experimental setting in which we compared the three exploration strategies described in the next section. The strategies have been integrated into a robot control architecture [11] and simulated runs have been performed in Player/Stage to assess and compare their performance. The system represents a class of widely used wheeled mobile robots and consists of a differential-drive robot platform equipped with a SICK LMS 200 laser range scanner with 180 degree field of view and 1 degree angular resolution. The goal of the robot is to fully explore an initially unknown environment.

Robot localization and mapping are performed by incrementally registering raw 2D laser range scans as described in [10]. The robot continuously updates the map as it moves. The map is represented as an unordered point cloud where duplicate storage of measurements is avoided by adding to the map only points that provide new information. They are determined according to a minimum distance from the already stored

points. In addition, we update a grid map that represents, for each cell $c^{[xy]}$, its reflection probability

$$p(c^{[xy]}) = \frac{\#hits}{\#hits + \#misses},$$

where $\#hits$ is the number of range beams that have been reflected by an object in the corresponding region and $\#misses$ is the number of range beams that have passed through the cell without being reflected. Initially, a value 0.5 is assigned to each cell, i.e., a cell's reflection is initially unknown. Path planning is accomplished by computing a *reachability map* which stores, for every cell, both the length of the shortest path to reach it from the current location of the robot and the preceding cell along this path. It is built by iteratively applying Dijkstra's algorithm on the grid map without specifying any goal location to fully explore the reachable workspace. Therefore, once a candidate location is selected, the shortest obstacle-free path for navigating to it can be recursively looked up in the reachability map. To guarantee safe navigation, we consider as *traversable* only cells $c^{[xy]}$ such that $p(c^{[xy]}) \leq 0.25$ and whose distance to the closest obstacle is less than 30 cm.

Finally, note that running all the exploration strategies in the same experimental setting provides a fair way to compare them. Furthermore, using the architecture described above allows for directly applying the implemented exploration strategies on real mobile robots.

IV. EXPLORATION STRATEGIES

In this section, we present the three exploration strategies that we compared in our experiments. These methods constitute a representative sample of different classes of NBV exploration strategies proposed in literature. The first one, a closest-frontier strategy [11], is simple, both in its definition and computation, and considers a single criterion for candidate selection – the traveling cost. The second technique [8] combines traveling cost and information gain with an *ad hoc* exponential function. The third technique [5] is based on a more principled way for aggregating multiple criteria in a global utility function.

A. Closest-Frontier Exploration Strategy

The idea of frontier-based exploration strategies is to detect borders between already explored regions of the environment and those regions where the robot has not yet acquired information. Hence, the robot searches for regions that are traversable in the map built so far and that are adjacent to unexplored regions and holes in the map.

A simple frontier-based exploration strategy is the *closest-frontier strategy* (CF). It has been proposed in [21] and can be briefly described according to the following steps:

- 1) determine the set T of traversable cells;
- 2) determine the set R of reachable cells, i.e., compute a reachability map (see Section III);
- 3) determine the set C of cells that are both reachable and traversable: $C = T \cap R$;

- 4) determine the set of frontier cells F by checking for every cell in the set C if it is adjacent to a cell with unknown reflection probability:

$$\begin{aligned} F &= \{c^{[xy]} \mid c^{[xy]} \in C, \\ &\exists c^{[(x+m)(y+n)]} : p(c^{[(x+m)(y+n)]}) = 0.5, \\ &m \in \{-1, 1\}, n \in \{-1, 1\}\}; \end{aligned} \quad (1)$$

- 5) determine $\mathbf{n} = (n^x \ n^y)^T$ as the frontier cell lying closest to the robot's current position $\mathbf{r} = (r^x \ r^y)^T$:

$$\mathbf{n} = \arg \min_{c^{[xy]} \in F} L\left((x \ y)^T, \mathbf{r}\right), \quad (2)$$

where $L(\mathbf{p}, \mathbf{r})$ is the length of the shortest path from \mathbf{p} to \mathbf{r} .

Finally, \mathbf{n} is chosen as the next best observation location and the robot is guided towards it following the minimum path.

B. González-Baños and Latombe's Exploration Strategy

The second exploration strategy we decided to evaluate is the strategy by *González-Baños and Latombe* (GBL) presented in [8]. It selects the next best observation location according to traveling cost and information gain.

Given the current partial map of the environment, this strategy generates a set of candidate locations by randomly sampling cells in the vicinity of frontier cells F . Then, given a candidate location \mathbf{p} , the corresponding utility $u(\mathbf{p})$ is computed according to two criteria: the traveling cost $L(\mathbf{p}, \mathbf{r})$ for reaching \mathbf{p} and the estimated information gain $I(\mathbf{p})$ when performing a sensing action at \mathbf{p} . The global utility is then computed as

$$u(\mathbf{p}) = I(\mathbf{p})e^{-\lambda L(\mathbf{p}, \mathbf{r})}, \quad (3)$$

and the candidate \mathbf{n} that maximizes $u(\cdot)$ is selected as the next observation position (in our experiments we used $\lambda = 0.2$, as suggested by the authors). Whereas the traveling cost is estimated in the same way as above (using the reachability map), the information gain is estimated as the expected relative change in map entropy. That is, we simulate range scans and corresponding map updates at all candidate locations \mathbf{p} . The information gain $I(\mathbf{p})$ is estimated as the difference between the map's entropy before (H) and after (\hat{H}) the simulated update $I(\mathbf{p}) = \hat{H} - H$. Since the probabilistic reflection maps we used represent, in principle, two probabilities for each cell (being occupied and being free), we estimate the map entropy by:

$$\begin{aligned} H &= - \sum_{c^{[xy]}} \left[\underbrace{p(c^{[xy]}) \log p(c^{[xy]})}_{\cong H_{p(\text{occupied})}} \right. \\ &\quad \left. + \underbrace{(1 - p(c^{[xy]})) \log(1 - p(c^{[xy]}))}_{\cong H_{p(\text{free})}} \right]. \end{aligned} \quad (4)$$

C. MCDM-based Exploration Strategy

This exploration strategy has been introduced in [5] for maps of line segments. Here we summarize it and show its extension to grid maps. This exploration strategy exploits a decision theoretic technique called *Multi-Criteria Decision Making* (MCDM), which constitutes a more principled way to combine the criteria that evaluate a candidate location.

Given a candidate location \mathbf{p} , we consider three criteria for its evaluation. The first one is the traveling cost $L(\mathbf{p}, \mathbf{r})$, computed as the length of the path connecting the current position of the robot with \mathbf{p} . Then, we consider the estimated information gain $I(\mathbf{p})$ and the overlap $O(\mathbf{p})$. These two last criteria should be maximized in order to select good observation locations. Both $I(\mathbf{p})$ and $O(\mathbf{p})$ are computed according to a standard entropy measure. Given the set of cells $V_{\mathbf{p}}$ that are visible from the candidate location \mathbf{p} , i.e., cells falling within the sensing range area centered at \mathbf{p} , we distinguish between *old* and *new* cells using a threshold k over the reflection probability. In particular, a cell $c^{[xy]} \in V_{\mathbf{p}}$ is considered as old if $p(c^{[xy]}) \leq k$ or if $p(c^{[xy]}) \geq 1 - k$, otherwise $c^{[xy]}$ is considered as new. In our experiments we set $k = 0.2$. Then, maximizing $I(\mathbf{p})$ corresponds to maximizing the total entropy over new cells of $V_{\mathbf{p}}$ (\mathbf{p} provides a potentially large amount of new information) while maximizing $O(\mathbf{p})$ corresponds to minimizing the total entropy over old cells of $V_{\mathbf{p}}$ (\mathbf{p} provides a good localization).

We call N the set of three criteria that are considered, $N = \{L(\cdot), I(\cdot), O(\cdot)\}$. Given a criterion $i \in N$ and a candidate location \mathbf{p} , an utility value $u_i(\mathbf{p})$ in the $[0, 1]$ interval is computed in order to evaluate on a common scale \mathbf{p} 's goodness according to every criterion. The utility is normalized over all the candidates in the current exploration step. For example, considering the traveling cost $L(\mathbf{p}, \mathbf{r})$ and called C the set of (current) candidate locations, the utility $u_L(\mathbf{p})$ (with $\mathbf{p} \in C$) is computed with the following linear mapping function:

$$u_L(\mathbf{p}) = \frac{1 - (L(\mathbf{p}, \mathbf{r}) - \min_{\mathbf{q} \in C} L(\mathbf{q}, \mathbf{r}))}{(\max_{\mathbf{q} \in C} L(\mathbf{q}, \mathbf{r}) - \min_{\mathbf{q} \in C} L(\mathbf{q}, \mathbf{r}))}. \quad (5)$$

Analogous normalization functions are used for other criteria, preserving the idea that the larger the utility the better the satisfaction of the criterion.

In order to select an observation location, the robot computes a global utility value measuring the overall goodness of each candidate. For every pair $\mathbf{p} \in C$ and $i \in N$ an utility value $u_i(\mathbf{p})$ is computed. MCDM uses an aggregation technique called *Choquet fuzzy integral*. Let us introduce this concept. We call a function¹ $\mu : \mathcal{P}(N) \rightarrow [0, 1]$ a *fuzzy measure* on the set of criteria N when it satisfies the following properties:

- 1) $\mu(\emptyset) = 0$, $\mu(N) = 1$,
- 2) if $A \subset B \subset N$, then $\mu(A) \leq \mu(B)$.

Given $A \in \mathcal{P}(N)$, $\mu(A)$ represents the weight of the set of criteria A . In this way, weights are associated not only to single criteria, but also to their combinations. Global utility $u(\mathbf{p})$ for a candidate location \mathbf{p} is computed by means of the Choquet

¹ $\mathcal{P}(N)$ is the power set of N .

$$\begin{aligned} \mu(L) &= 0.2 & \mu(O) &= 0.4 & \mu(\{L, O\}) &= 0.6 \\ \mu(I) &= 0.4 & \mu(\{L, I\}) &= 0.9 & \mu(\{I, O\}) &= 0.8 \end{aligned}$$

TABLE I
DEFINITION OF $\mu()$ FOR THE MCDM-BASED STRATEGY

integral with respect to the fuzzy measure μ :

$$u(\mathbf{p}) = \sum_{i=1}^{|\mathcal{N}|} (u_{(i)}(\mathbf{p}) - u_{(i-1)}(\mathbf{p})) \mu(A_{(i)}), \quad (6)$$

where (i) indicates the indices after a permutation that changed their order to have, for a given \mathbf{p} , $u_{(1)}(\mathbf{p}) \leq \dots \leq u_{(|\mathcal{N}|)}(\mathbf{p}) \leq 1$ (it is supposed that $u_{(0)}(\mathbf{p}) = 0$) and

$$A_{(i)} = \{j \in \mathcal{N} | u_{(i)}(\mathbf{p}) \leq u_j(\mathbf{p}) \leq u_{(|\mathcal{N}|)}(\mathbf{p})\}.$$

Different aggregation functions can be defined by changing the definition of μ . For example, weighted average is a particular case of the Choquet integral when μ is additive (i.e., $\mu(A \cup B) = \mu(A) + \mu(B)$). Most importantly, μ can model dependence relationships between criteria. Formally, criteria belonging to a group $G \subseteq \mathcal{N}$ are:

- redundant, if $\mu(G) < \sum_{g \in G} \mu(g)$;
- synergic, if $\mu(G) > \sum_{g \in G} \mu(g)$;
- independent, otherwise.

In summary, what MCDM provides is a sort of “distorted” weighted average, which takes into account dependency between criteria. The next observation location is the candidate location that maximizes $u()$ in Eq. (6).

The MCDM-based strategy used in experiments has been defined according to the weights reported in Table I. Such weights have been manually chosen in order to model a synergy relation between the information gain $I()$ and the traveling cost $L()$, thus favoring candidates that satisfy those criteria in a balanced way. Finally, we note that the computational time of employing MCDM, although longer than that of employing CF, has a negligible impact on the time required to map an environment.

D. Heuristics to Improve the Strategies

The three exploration strategies we considered (and most of those presented in literature) have two main limitations:

- 1) the decision of reaching a selected location is not changed until the location is actually reached,
- 2) evaluation of candidate locations is based only on information relative to the single locations, without considering their relation with other locations.

In the next section, we provide an experimental answer to the question of how much these limitations affect the performance of exploration strategies. Here, we describe two simple heuristics that can be applied to exploration strategies in order to cope with these limitations and to obtain a better performance (thereby extending the initial results from [9]).

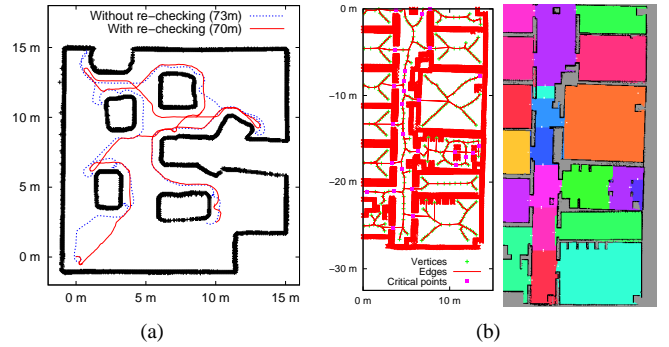


Fig. 1. Example trajectories with and without repetitive re-checking (a) and examples of map segmentation (b).

1) *Repetitive Re-checking*: During navigation to a selected location, the map is continuously updated. As a result, the robot might have fully explored an unknown region before actually reaching the selected frontier location. Hence, continuing to travel to the selected location is unnecessary. We address this problem by using *Repetitive Re-checking* (RR), i.e., the robot checks whether or not the currently approached frontier location \mathbf{n} is still adjacent to at least one cell with unknown reflection probability. As soon as \mathbf{n} is no longer a valid frontier, the robot stops traveling towards it and selects the next best location according to the employed exploration strategy. In Fig. 1(a) we report an example that shows that the robot’s trajectory is shortened by repetitive re-checking, especially when approaching frontiers in the vicinity of corners.

2) *Map Segmentation*: Sometimes it can happen that a single room gets visited multiple times if successively selected locations lie in different rooms. To reduce the number of multiple visits, we applied *Map Segmentation* (SEG), which splits the map built so far into segments representing individual rooms and makes the robot prefer candidates lying in the segment of its current location.

We use an approach based on [18] and [20] that splits map regions at local minima in the Voronoi diagram (critical points) of the map’s free space. We define critical points to be local minima with respect to the distances to the closest Voronoi site, nodes of degree 2, and to be itself adjacent to a junction node or adjacent to another node that is adjacent to a junction node. Using critical points we split previously unassigned map regions into two parts. We assign cells to segments with respect to their distances to critical points. That is, we form clusters of cells being closest to a common split point. This can be performed efficiently by computing an Euclidean distance transform (EDT) for the critical points. For the actual assignment we compute and store both the distance to the closest critical point (as for the EDT) and the closest critical point itself; thus computing a nearest neighbor transform. Then, in an iterative refinement step, we merge segments that are adjacent to each other but not split by the same critical point. An example of the segmentation algorithm is reported in Fig. 1(b).

The map segmentation can be used in the exploration strategies to restrict the set of candidate locations within the scope of the robot’s current segment. If the set of candidates




Fig. 2. The two environments provided by K. Lingemann, A. Nüchter, and J. Hertzberg (a) and by R. Vaughan (b).

belonging to the robot’s current segment is not empty, then the exploration strategy will choose the next best location from that set. In this sense, we will say that, when SEG is used, a “depth-first-like” or *room-by-room* exploration is favored.

V. EXPERIMENTAL RESULTS

We compared the performance of exploration strategies in two office-like indoor environments composed of several rooms and corridors (Fig. 2). Indoor environments present interesting challenges to exploration strategies, mainly due to their intricate structure that makes the selection of the next observation position non-trivial.

We compared the three exploration strategies of Section IV with and without the RR and SEG heuristics. As a baseline for comparison, we report also results obtained with a Random Frontier (RF) exploration strategy. It chooses the next observation location according to a uniform probability distribution over the current candidate locations.

For every configuration in which a particular exploration strategy was tested within an environment, we performed 50 simulation runs with the same initial position for the robot. Each run is considered completed when no more frontiers can be determined in the current map, i.e., when there does not exist any reachable cell adjacent to another cell with unknown reflection probability. To compare the performance obtained in different configurations, we report the mean of the length of trajectories covered by the robot (as in [1]–[3], [5], [16], [19]).

Results obtained in the two environments are reported in Fig. 3. All the strategies perform better than RF, as expected, with more evident differences in the more complex hospital environment. The first interesting comparison that is worth doing is between CF and MCDM strategies. The good performance of CF means that minimizing the traveled distance at every exploration step produces a small global traveled distance in the indoor environments we considered. This fact and recalling that the robot acquires data during its movements explain the good performance of CF. Although CF performs slightly better, MCDM achieves comparable performance with respect to CF. This is not obvious, since in MCDM other criteria ($I()$

and $O()$) are given more importance than traveling cost (see Table I), which is the only criterion adopted by CF. In fact, the MCDM strategy provides, by means of synergy, a good trade-off between $I()$ and $L()$. The close performance of CF and MCDM can be explained also by saying that the latter strategy compensates the potential performance worsening, due to the fact that distance is not minimized, with good information gains. Moreover, we observed that MCDM maps most of the environment following a short path and then travels a relatively long path to complete the map (e.g., filling holes close to corners).

A reduction in the total traveled distance of the three strategies can be observed when enabling Repetitive Re-checking (RR). A strategy with the RR heuristic outperforms the corresponding basic strategy, which needs to reach every selected observation location independently of the sensorial data acquired along the path. Using map segmentation (SEG) reduces the traveled distance especially in the hospital environment, where SEG provides a good quality segmentation. Enabling SEG prevents to leave out corners and occlusions and exploring them in the last steps of the exploration. Without SEG, multiple visits to the same room can be necessary, e.g., when the current robot’s room is not completely explored and the best frontier location happens to be outside that room. Interestingly, the MCDM strategy showed a “depth-first” behavior with respect to unknown regions, even without using SEG. The main reason is the presence of the overlap criterion, which leads to a more conservative exploration by imposing to have a certain amount of old information in each sensorial acquisition. Comparing the two heuristics, SEG appears to reduce the traveled distance slightly more than RR. This result can be explained by considering that RR only stops following an already made decision, while SEG helps in making a better decision.

The GBL strategy (with and without RR and SEG) is outperformed by MCDM and CF in both the environments. This means that using more criteria does not guarantee by itself to obtain a better exploration strategy and suggests that the way in which criteria are combined is fundamental. In this sense, general aggregation techniques such as MCDM appear more suitable to design multi-criteria exploration strategies. This is in accordance with the results of [5], where exploration strategies defined with MCDM and GBL are compared using maps composed of line segments.

Finally, we considered a variant of the GBL strategy, in which the information gain is computed as in MCDM, i.e., by using the entropy only over the new cells visible from a candidate location (data are not shown here). With this different $I()$, GBL shows a slightly better performance, but the above considerations still hold. This suggests that the way in which criteria are combined could be even more important than the methods used to compute the criteria themselves.

An evaluation criterion that has been largely neglected so far, is the *completeness* of the maps after exploration. In our experiments, we compared the map entropies resulting from exploration with those computed on (manually) fully explored maps ($c = (1 - (H_{\text{full}}/H_{\text{expl}})) * 100$ to obtain completeness c in percent). By having the same termination criterion, all basic




Fig. 3. Traveled distances in m (average and standard deviation).

strategies (with and without SEG) achieve a map completeness of roughly 99%. RR lowers the completeness to 98% in the AVZ environment, and 96% in the hospital environment. This is primarily caused by taking less close range measurements in corners, which can be seen in Fig. 1(a). That is, there is a trade-off in RR between shortening the robot's trajectory and increasing its map uncertainty.

VI. CONCLUSIONS

In this paper, we addressed the experimental comparison of frontier-based exploration strategies for an autonomous mobile robot that maps an unknown environment. A representative sample of three strategies proposed in literature has been evaluated in combination with two improvement techniques in a common simulated experimental setting. Some insights obtained from our analysis, like the influence of the function used to combine criteria in evaluating candidate locations, can help in developing better exploration strategies. Our work is intended to constitute another step toward the definition of good experimental methodologies for exploration strategies. In particular, in our experimental framework we used a standard simulation platform in order to support the comparison, reproducibility, and repeatability of experiments.

Several additional issues can be considered to improve the experimental framework of this paper. For example, it would be interesting to compare performance of exploration strategies with that of optimal offline coverage strategies (in which the map is known) [7]. Another issue worth considering is the metric used to measure performance. In this paper, we have considered the traveled distance to account for the energy and time effort, but also the number of map updates and the entropy of the final map can be considered to account for, respectively, the computational effort and for the quality of the produced map (which could also involve loop closures). Moreover, the relationships between the performance of the strategies and the particular setting (robot locomotion, speed, etc.) deserve more attention. Finally, extensions to 3D and flying robots is a matter of future work.

REFERENCES

- [1] F. Amigoni. Experimental evaluation of some exploration strategies for mobile robots. In *Proc. ICRA*, pages 2818–2823, 2008.
- [2] F. Amigoni and V. Caglioti. An information-based exploration strategy for environment mapping with mobile robots. *Robotics and Autonomous Systems*, 58(5):684–699, 2010.
- [3] F. Amigoni and A. Gallo. A multi-objective exploration strategy for mobile robots. In *Proc. ICRA*, pages 3861–3866, 2005.
- [4] F. Amigoni, M. Reggiani, and V. Schiaffonati. An insightful comparison between experiments in mobile robotics and in science. *Autonomous Robots*, 27(4):313–325, 2009.
- [5] N. Basilico and F. Amigoni. Exploration strategies based on multi-criteria decision making for an autonomous mobile robot. In *Proc. ECMR*, pages 259–264, 2009.
- [6] EURON GEM Sig. <http://www.heeronrobots.com/EuronGEMSig/>, 2007.
- [7] Y. Gabrieli and E. Rimon. Competitive complexity of mobile robot online motion planning problems. *International Journal of Computational Geometry and Applications*, 20(3):255–283, 2010.
- [8] H. González-Baños and J.-C. Latombe. Navigation strategies for exploring indoor environments. *International Journal of Robotics Research*, 21(10-11):829–848, 2002.
- [9] D. Holz, N. Basilico, F. Amigoni, and S. Behnke. Evaluating the efficiency of frontier-based exploration strategies. In *Proc. ISR 2010*, pages 36–43, 2010.
- [10] D. Holz and S. Behnke. Sancta simplicitas – on the efficiency and achievable results of SLAM using ICP-Based Incremental Registration. In *Proc. ICRA*, pages 1380–1387, 2010.
- [11] D. Holz, G. K. Kraetzschmar, and E. Rome. Robust and Computationally Efficient Navigation in Domestic Environments. In *RoboCup 2009: Robot Soccer World Cup XIII*, volume 5949/2010 of *Lecture Notes in Computer Science*, pages 104–115. Springer, Germany, 2009.
- [12] D. Lee and M. Recce. Quantitative evaluation of the exploration strategies of a mobile robot. *International Journal of Robotics Research*, 16:413–447, 1997.
- [13] J. Leonard and H. Feder. A computationally efficient method for large-scale concurrent mapping and localization. In *Proc. Int'l Symposium on Robotics Research*, pages 169–176, 1999.
- [14] R. Madhavan, C. Scrapper, and A. Kleiner. Special issue on characterizing mobile robot localization and mapping. *Autonomous Robots*, 27(4):309–481, 2009.
- [15] P. Newman, M. Bosse, and J. Leonard. Autonomous feature-based exploration. In *Proc. ICRA*, pages 1234–1240, 2003.
- [16] C. Stachniss and W. Burgard. Exploring unknown environments with mobile robots using coverage maps. In *Proc. IJCAI*, pages 1127–1134, 2003.
- [17] C. Stachniss, D. Haehnel, and W. Burgard. Exploration with active loop-closing for FastSLAM. In *Proc. IROS*, pages 1505–1510, 2004.
- [18] S. Thrun. Learning Metric-Topological Maps for Indoor Mobile Robot Navigation. *Artificial Intelligence*, 99(1):21–71, 1998.
- [19] B. Tovar, L. Muñoz-Gomez, R. Murrieta-Cid, M. Alencastre-Miranda, R. Monroy, and S. Hutchinson. Planning exploration strategies for simultaneous localization and mapping. *Robotics and Autonomous Systems*, 54(4):314–331, 2006.
- [20] K. Wurm, C. Stachniss, and W. Burgard. Coordinated Multi-Robot Exploration using a Segmentation of the Environment. In *Proc. IROS*, pages 1160–1165, 2008.
- [21] B. Yamauchi. A frontier-based approach for autonomous exploration. In *Proc. CIRA*, pages 146–151, 1997.