On evaluating performance of exploration strategies for an autonomous mobile robot

Nicola Basilico and Francesco Amigoni

***Abstract*— The performance of an autonomous mobile robot in mapping an unknown environment is strongly related to its exploration strategy. Evaluation and comparison of exploration strategies are still waiting for the definition of standard benchmark methodologies. In this paper, we contribute in this direction by discussing some critical issues we faced in experimentally evaluating exploration strategies.**

I. INTRODUCTION

The task of robot mapping [1] concerns the construction of spatial representations of environments by means of autonomous mobile robots. In mapping, autonomous robots move within the environment where they have been deployed for acquiring data. *Exploration strategies* determine the locations reached by the robots, within the partially explored environment. The definition of a good strategy is an important aspect which strongly influences robot performance in mapping tasks. Different exploration strategies have been proposed. However, their evaluation and comparison have not been fully addressed. The definition of standard benchmarks and experimental methodologies is an hot topic in current robotic research. For example, the Rawseeds project [2] aims at building a standard benchmarking toolkit for Simultaneous Localization and Mapping (SLAM) algorithms. The specific problem of evaluating exploration strategies, independently from the whole mapping systems in which they are embedded, is still open. To the best of our knowledge only few works [3], [4] explicitly addressed the evaluation and comparison of exploration strategies.

In this paper we aim at contributing to the definition of an experimental methodology for exploration strategies by enlightening and discussing some critical issues we encountered in experimentally evaluating them. Our goal is not to compare a number of strategies to find the best one, but to report on some relevant points we faced in our work. The points we raise should be considered by an evaluation methodology for exploration strategies. In particular we consider the following aspects:

- selection of appropriate metrics,
- selection of values for parameters of strategies,
- characteristics of the testing environments.

We assume to have an autonomous mobile robot equipped with a range sensor able to acquire $360°$ spatial information within a range r and we view its exploration as a sequence of observations in different locations of the environment. Each sensing action produces a *partial map*, which is integrated in a *global map*, in order to incrementally build a complete spatial representation of the environment. A strategy allows the robot to choose good observation locations. As in most approaches adopted in literature, the robot evaluates a set of candidate locations in the free space of its current global map and chooses the best one. Candidate locations are evaluated by means of measurable features called *criteria* (examples are the travelling cost and the estimated information gain). Given a candidate location p and a criterion i, $u_i(p)$ is the *utility* of p for the criterion i. It quantifies the goodness of p with respect to i. Utilities referring to different criteria are aggregated in a global utility $u(p)$. The robot selects the location that maximizes $u()$. Different criteria and aggregation functions can be employed to define different exploration strategies. We use the Player/Stage [5] simulator to provide concrete examples of issues arising when evaluating exploration strategies.

This paper is structured as follows. Next section overviews the exploration strategies proposed in literature. The experimental framework of Section III is used to point out some critical issues in Section IV. Section V concludes the paper.

II. EXPLORATION STRATEGIES

The definition of strategies for autonomous exploration of environments has been addressed by several works in literature. The proposed techniques can be classified according to two distinct approaches. In the first one, exploration strategies make use of *predefined trajectories* [6], [7]. Trajectories followed by the robot are static and defined off-line, exploiting some *a priori* information about the environment. The main disadvantage of these strategies is their limited adaptability to different environments and the difficulty in building effective trajectories when no information about the environment to explore is available.

In the second approach, depicted in Fig. 1, exploration is viewed as a sequence of steps, each one composed of a movement towards a location and of an observation with which the robot acquires data about the environment. In this context, the exploration strategy does not address, like in the first approach, *how* to move within the environment, but, at each exploration step, it focuses on *where* to move in order to make the next observation. Systems of this class, called *Next-Best-View* (NBV) systems, choose the next observation location among a set of currently reachable candidate locations, evaluating them according to some criteria. Usually, in the NBV approach, candidate locations are generated on the frontier between the free known space and the unexplored

Nicola Basilico and Francesco Amigoni are with the Dipartimento di Elettronica e Informazione, Politecnico di Milano, Piazza Leonardo da Vinci 32, 20133, Italy {basilico,amigoni}@elet.polimi.it

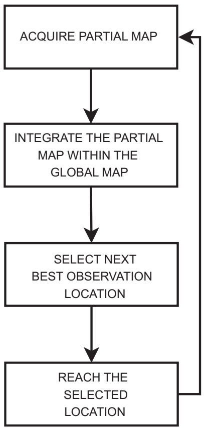

Fig. 1. Next-Best-View exploration approach

part of the environment [8]. By their nature, NBV strategies can easily adapt to different environments. In this paper we consider strategies defined with this approach.

In evaluating a candidate location, different criteria can be used. A simple criterion is the *travelling cost* [9], according to which the best observation location is the nearest one. Other works combine different utilities with *ad hoc* functions. In [10] the travelling cost is combined with *information gain*. This last criterion is related to the amount of new information about the environment, obtainable from a candidate location p. It can be estimated by measuring the area of the portion of unknown environment potentially visible from p, computed according to the global map currently available to the robot. The larger the estimate, the more interesting the candidate location. Given a candidate location p and called $c(p)$ and $A(p)$ the travelling cost and the expected information gain, respectively, these two criteria are combined with an *ad hoc* function in order to compute an overall utility (λ weighs the travelling cost and the information gain):

$$u(p) = A(p)e^{-\lambda C(p)} \quad (1)$$

Other examples are [11], in which the traveling cost is linearly combined with the expected uncertainty of the map after the observation, and [12], in which a technique based on relative entropy is used. Strategies based on *ad hoc* aggregation functions usually are strongly related to the number and meaning of the considered criteria. For this reason, their application in other situations could be difficult. In the work presented in [13], authors dealt with this problem evaluating a theoretically-grounded approach based on multi-objective optimization. The best candidate is selected on the Pareto frontier as the one that better satisfies all the considered criteria. Besides cost and information gain, *overlap* can be taken into account. This criterion is related to the amount of common features between the current global map and the partial map that will be acquired from an observation location. It accounts for self-localization used by the robot to estimate its current position in the environment according to the map: the larger the overlap, the better the localization of the robot.

Exploration strategies have been studied also in multirobot scenarios [14], [15], where the task of exploration and mapping is performed by a team of robots. For example, in [14] each robot evaluates a candidate by means of a weighted aggregation function which combines two criteria: the travelling cost and a general measure of goodness, initially equal for all candidates, that decreases with the proximity of other robots.

More general navigation strategies, of which exploration strategies are a particular case, play an important role also in complex robotic tasks, in which the robot is requested to pursue several, and sometimes conflicting, sub-objectives. An example is the explore-and-search task, involved in search and rescue applications. In this case the robot has not only to explore the environment (a disaster site) but also to find and communicate the position of victims [16], [17]. For example, the navigation strategy proposed in [17] considers the probability of communication success from a candidate location, combining it with cost and expected information gain within an *ad hoc* function. In this context, a priori information related to the task can improve the effectiveness of the strategy. An example can be the possibility to know in advance if victims are uniformly spread or concentrated in clusters. An attempt to integrate this kind of information in an explore-and-search strategy has been presented in [18], where an high-level formalism based on Petri nets is used.

III. The Experimental Framework

In order to illustrate issues involved in evaluating exploration strategies, we defined a navigation system with which we provide some concrete example in the following. A C++ software simulator based on Player/Stage [5] architecture and CGAL graphic libraries [19] has been developed. With this tool different exploration strategies can be tested in different environments, changing the initial location of the robot and the maximum range r of its sensor. In the following we describe how we implemented the NBV process of Fig. 1.

A. Environment representation

The robot maintains a map of the environment, built integrating all the observations. We represent the map with 2D line segments. In particular two lists are used. The *obstacle list* contains line segments representing the obstacles detected in the environment. The *free edge list* stores line segments representing the frontiers between known and unknown space. Each time a new observation is performed, a new partial map m is acquired and the current robot's map M is updated by aligning m to M (according to the position of the robot given by the simulator) and by fusing them. A partial map is composed both of obstacle line segments and of free edges. New obstacle line segments are added to the obstacle list and, similarly, new free edges are inserted in the free edge list. Moreover, old free edges which, after the observation, happen to belong to the explored area, are deleted from the corresponding list.

B. Candidate evaluation

Candidate locations are generated as the middle points of the line segments in the free edge list. Given a candidate location p, we consider four different criteria for its evaluation. The travelling cost $c(p)$ is computed as the total length of the path connecting the current position of the robot with p. In order to facilitate path-planning, a reachability graph for all the available candidate locations is maintained during the exploration. Two different estimates of expected information gain are considered. The first one, $i_{area}(p)$, is computed as the amount of area visible from p, i.e., falling within sensor range and not belonging to the already mapped space. In the second estimate, the amount of new information potentially obtainable from p, $i_{seg}(p)$, is computed as the length of free edge line segments visible from p. The last considered criterion is the overlap $o(p)$ between the current map and the area visible from p. It is calculated as the length of obstacle line segments which are visible from p. In our experiments we used $r = 8m$ in order to force the robot to make a significant number of steps to complete exploration. For each one of these criteria an utility value, in the $[0, 1]$ interval, is computed in order to evaluate on a common scale p's degree of satisfaction according to every criterion. The utility is defined by normalization over all the candidates in the current exploration step. For example, considering the travelling cost $c(p)$ and called $\mathcal{L}$ the set of candidate locations available at the step k, the utility related to this criterion for $p \in \mathcal{L}$ is computed with the following linear mapping function:

$$u_c(p) = 1 - (c(p) - \min_{q\in\mathcal{L}} c(q))/(\max_{q\in\mathcal{L}} c(q) - \min_{q\in\mathcal{L}} c(q)) \quad (2)$$

Similar normalization functions are used for other criteria. The use of relative normalization expressed by (2) is justified by the independence between different robot's choices at different steps. Indeed, due to the greedy nature of the NBV approach, the result of the robot's decision at step k only depends on $\mathcal{L}$ and not on previous decisions and previous sets of candidate locations.

IV. Some issues on the experimental evaluation of exploration strategies

In this section some relevant issues arising in experimentally evaluating exploration strategies are discussed. We remark that our goal is not to provide a comprehensive evaluation to select the "best" strategy. We aim at enlightening some critical issues that must be dealt with when evaluating exploration strategies and that should be considered in developing a methodology.

A. Metrics

A first issue, which plays a central role in performance evaluation, is related to metrics and indicators used to assess the goodness of an exploration strategy and to compare different exploration strategies. In order to discuss how to choose such metrics, we need to clearly define the objectives that a good exploration strategy has to meet.

Efficiency is the first obvious goal and can be expressed through the minimization of two types of costs. The first one is the time needed to map the environment (or a given percentage of it), while the second one is related to the power consumption of the robot. In general, these two objectives are not independent, namely they cannot always be optimized separately. It seems adequate to use two metrics that are related to both of them: the *number of sensing actions* performed by the robot and the total amount of *distance travelled*. An exploration strategy providing good performance according to both these indicators can be considered efficient with respect to time and power consumption. Importantly, these indicators are general, and abstract away from the technical details of the employed robot, making comparison of strategies rather independent from the particular navigation system.

A second objective is represented by effectiveness, and can be assessed by evaluating the map produced by the robot. It has to represent in a consistent way the spatial information collected during the exploration and to reliably describe the environment. Differently from efficiency, the quality of a map does not depend only on the exploration strategy employed. Different components of the robot's navigation system influence it. For example, it is evident that using a poor registration algorithm to align a partial map to the global map can result in an inaccurate final map. Since results can depend on several aspects, it is somehow unfair to use metrics that measure the quality of the map to compare exploration strategies. However, it is necessary to qualitatively assess the consistency of produced maps by means of direct observation.

Another point is that performance metrics are usually calculated on aggregated data, reporting, for example, the total amount of travelled distance or the total number of sensing actions needed to map a given environment. However, although this description shows a good level of synthesis and is widely employed [4], [11], [13], a *cumulative* description allows more qualitative insights on the measured performance.

For instance, it is interesting to relate the amount of the explored area to the distance travelled or the number of sensing actions. For example, Fig. 2 reports the covered area versus the travelled distance for a strategy based on the minimization of the travelling cost. The three curves represent the behavior of the robot when starting from 3 different locations in the same environment (locations 1, 3, and 4 of Fig. 5(a)), while symbols on the curves are sensing actions. The mapping ends when the 90% of the overall area has been covered. As it can be seen, the best performance is obtained from starting location 1, despite in the first steps the robot performed better when starting from location 3. Around step 500 the amount of covered area starting from location 1 increases while travelling short distances. From the graph it is also possible to observe that the robot, when starting from location 4, made a significant number of perceptions without obtaining any new spatial information and, in order to find better observation locations, it had to travel long distances.

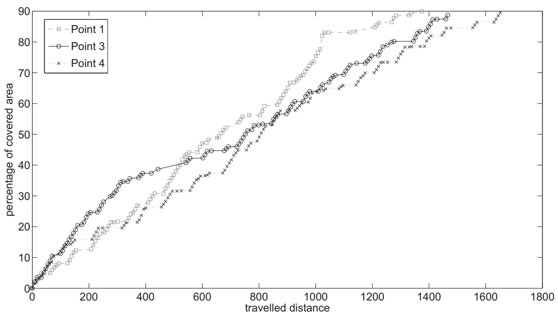

Fig. 2. Cost-minimization strategy from starting locations 1, 3, 4 of Fig. 5(a)

This can be observed looking at the horizontal portions of its associated curve. All of these features are evident from cumulative data but are not visible from metrics calculated on aggregated data.

Performance metrics should depend only on exploration strategies. A cumulative description can evidence interesting qualitative insights on performance.

B. Parameters

Another critical issue is related to the values for parameters. It is not surprising that parameters are involved in the definition of every exploration strategy. For example, in our sample context of Section III-B, a strategy is defined by a set of criteria and by the aggregation function used to compute the overall utility for candidate locations. Let us assume that we want to compare strategies with different criteria sets, using a weighted average as aggregation function. Therefore, called $\mathcal{C}$ the set of considered criteria, and $u_i(p)$ the utility of the candidate location p with respect to the criterion $i \in \mathcal{C}$, the overall utility of p is obtained as:

$$u(p) = \sum_{i \in \mathcal{C}} \omega_i u_i(p) \tag{3}$$

Parameters are represented, in this situation, by coefficients ω_i. Each ω_i represents the weight associated to a criterion i, namely its relative importance. Another example is represented by the strategy proposed in [10] which employs the *ad hoc* aggregation function (1). Parameter λ is still related to the relative importance of the criteria, since it weighs the travelling cost against the expected information gain.

Selection of values for these parameters could depend on the particular task the mobile robot is facing and on the type of the criteria considered. Intuitive selection of values may be misleading and produce unexpected results. For example, let us consider as criteria the travelling cost and the two information gain estimates, i.e. $\mathcal{C} = \{c, i_{area}, i_{seg}\}$. We adopt the aggregation function (3) and we compare the performance obtained with two different settings of values for parameters, reported in Table I. The first set of values give more importance to the travelling cost, defining a strategy called *CostFirst*, while the other set of values give more importance to information gain, defining a different strategy called *InfoGainFirst*. Given an environment, which strategy will map it faster? *CostFirst*, which evaluates better close candidate locations, or *InfoGainFirst*, which evaluates better candidate locations that offer good views on unexplored areas? In Fig. 3 we report an example of comparison of the two strategies. Robot started from the same location within the same environment (location 5 of Fig. 5(a)). Due to the characteristics of the exploration task we would have expected better performances to be achieved by *InfoGain-Fisrt* strategy. Indeed, it seems reasonable to consider the acquisition of new spatial information as the primary way to conduct a fast exploration. However, as it can be seen from the graph, better performaces are achieved with the *CostFirst* strategy. When giving more importance to the information gain estimates, the robot tends to travel longer in order to reach the most promising locations in terms of new obtainable information. In Fig. 3 this is represented by the horizontal portions of the associated curve.

	CostFirst	*InfoGainFirst*
c	0.4	0.2
i_{area}	0.3	0.4
i_{seg}	0.3	0.4

TABLE I

WEIGHTS FOR *CostFirst* AND *InfoGainFirst* STRATEGIES

Another consideration can be called for to explain this result: the difference between the criteria. The travelling cost is an exact criteria, while expected information gains are estimates. For instance, a location with a good expected information gain can become, after the sensing action, a bad observation location. Therefore, it is important when choosing weights for criteria to consider their nature and to distinguish between exact criteria, which do not introduce any error, and criteria that are more or less precise estimates. In this sense, giving more importance to former ones makes exploration strategy more “robust”. A final aspect to consider is related to the meaning of criteria. The criteria represented by i_{area} and i_{seg} are based on different techniques to estimate the same feature i.e., information gain, therefore, giving a large weight to both of them can lead to overevaluating information gain. For this reason, weight values should also consider that some criteria can be redundant.

Parameters usually set the relative importance of criteria that are used to evaluate candidate locations. Giving high importance to criteria which are based on estimates can worsen the performance.

C. Overlap

As already discussed, the evaluation of exploration strategies can be influenced by particular aspects of the navigation system. Let us show this aspect by considering the overlap criterion that is related to the ability of the robot to localize itself. We define the strategy *CostFirst with Overlap* with the weights reported in Table II. The highest importance is still given to the travelling cost, but also overlap is taken into

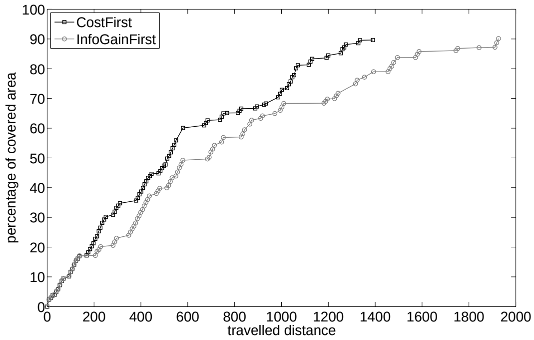

Fig. 3. Performance of *CostFirst* and *InfoGainFirst* strategies from starting location 5 of Fig. 5(a)

	CostFirst with Overlap
c	0.4
i_{area}	0.2
i_{seg}	0.2
o	0.2

TABLE II

WEIGHTS FOR *CostFirst with Overlap* STRATEGY

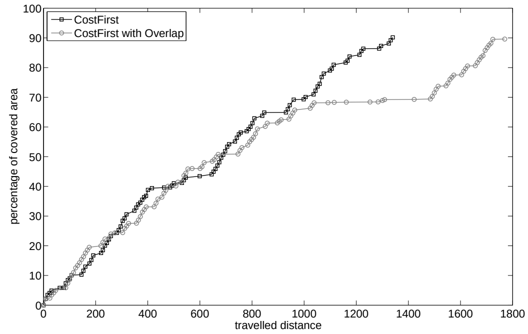

(a) Starting location 3

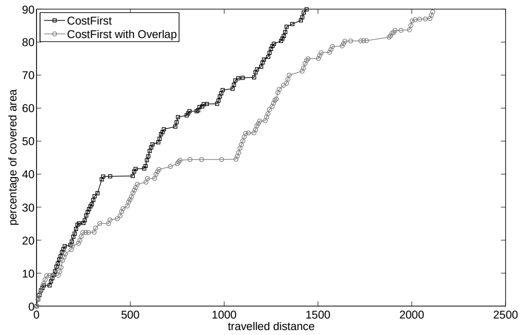

(b) Starting location 6

Fig. 4. Performance of *CostFirst* and *CostFirst with Overlap* strategies from starting locations 3, 6 of Fig. 5(a)

account. In Fig. 4 two significant examples are reported, where *CostFirst with Overlap* strategy is compared with *CostFirst* when the robot starts from locations 3 and 6 of Fig. 5(a). In Fig. 4(a), referred to starting location 3, we can observe that the two strategies performed similarly until a certain point, where the growth of the mapped area started to differ significantly. What happened is that after travelling 1000 units, the robot, when executing the *CostFirst with Overlap* strategy, decided to cover large distances in order to reach locations which well satisfy the overlap criterion. Despite this kind of behavior can introduce improvements for localization, it introduces a worsening in the adopted performance metrics. This is even more evident when considering the results obtained from starting location 6, reported in Fig. 4(b). Several good observation locations, from which a large amount of new area is expected to be visible, are not visited with the consequence of a lower growth rate of the mapped space. Exploration towards unknown areas is limited by the fact that good observation locations are constrained to have a significant amount of already mapped space that is visible. Since the advantages introduced by overlap are not visible when using our metrics (recall Section IV-A), comparison of strategies with and without overlap is not sound. Therefore, we suggest that overlap should be addressed directly as a part of the navigation system, considering it as a threshold under which a candidate location is not taken into account for evaluation. This amounts to say that candidate locations from which localization is difficult are not considered.

Overlap is an example of a criterion that can introduce advantages for the navigation system e.g., better localization. These advantages are not reflected by the chosen performance metrics. This criterion should be embedded as a feature of the navigation system.

D. Environment

The last issue we consider is the selection of the environments. Since the adaptability of a strategy is a more-than-desirable characteristic, performance evaluation has to be conducted in different kinds of environments. For example, consider the two indoor environments in Fig. 5. The first environment (Fig. 5(a)) is characterized by the presence of rooms and corridors, while the second one (Fig. 5(b)) is an open space. The main issue in dealing with different environments is to understand where differences between strategies can be seen more clearly.

In many experiments we found that differences between strategies are more likely to emerge in open spaces. Generally in these environments the number of candidate locations at each step is larger than in more cluttered environments. Therefore, there is a significant number of steps in which the decision is not trivial and where different strategies can lead to very different performances. This can be seen in Fig. 6, which compares the performance of *CostFirst* and *Latombe* strategies (defined as (1) with $\lambda = 0.2\mathrm{m}^{-1}$) in the environments of Fig. 5.

Open spaces evidence the different performance of exploration strategies.

(a) Cluttered environment (approx. $7200\mathrm{m}^2$)

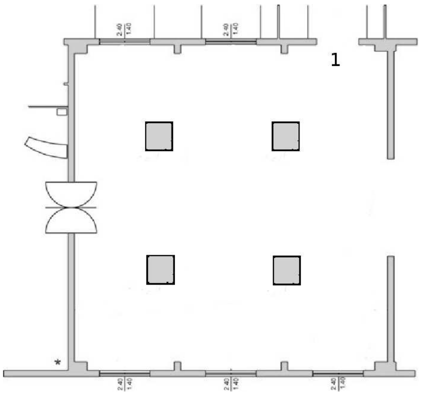

(b) Open space (approx. $7000\mathrm{m}^2$)

Fig. 5. Testing environments (numbers indicate starting locations)

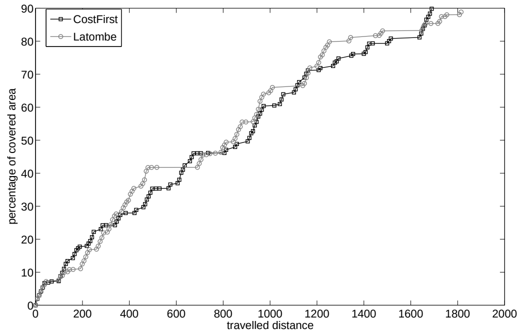

(a) Cluttered environment (starting location 2)

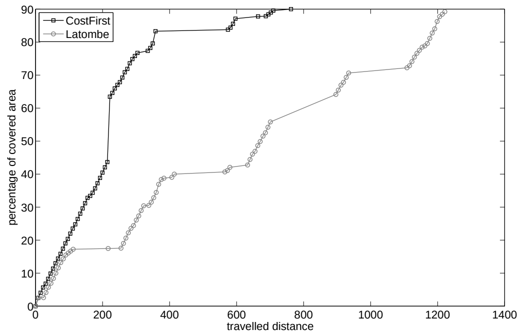

(b) Open space (starting location 1)

Fig. 6. Performances in different environments

V. Conclusions

In this paper we have put forward some critical issues related to the experimental evaluation of NBV exploration strategies. We discussed selection of appropriate performance metrics, selection of values for parameters, and some aspects regarding the characteristics of the environments where evaluation is performed. As stated at the beginning, the purpose of this paper is not to give final solutions for these critical issues, but to contribute to the definition of a methodology for evaluating exploration strategies. In this sense, we believe that the issues we discussed should be accounted for in such a methodology.

References

[1] S. Thrun, "Robotic mapping: A survey," in *Exploring Artificial Intelligence in the New Millenium*. Morgan Kaufmann, 2002.

[2] "The rawseeds project." [Online]. Available: http://rawseeds.elet.polimi.it/

[3] D. Lee and M. Recce, "Quantitative evaluation of the exploration strategies of a mobile robot," *International journal of robotic research*, vol. 16, pp. 413–447, 1997.

[4] F. Amigoni, "Experimental evaluation of some exploration strategies for mobile robots," in *IEEE International Conference on Robotics and Automation, 2008. ICRA 2008*, 2008, pp. 2818–2823.

[5] "The player project - free software tools for robot and sensor applications." [Online]. Available: http://playerstage.sourceforge.net/

[6] E. Bourque and G. Dudek, "Viewpoint selection-an autonomous robotic system for virtual environment creation," in *IROS*, vol. 1, 1998, pp. 526–532.

[7] J. Leonard and H. Feder, "A computationally efficient method for large-scale concurrent mapping and localization," in *International Symposium on Robotics Research*, 1999, pp. 169–176.

[8] B. Yamauchi, "A frontier-based approach for autonomous exploration," in *Proceedings of the IEEE International Symposium on Computational Intelligence in Robotics and Automation*, July 1997, pp. 146–151.

[9] B. Yamauchi, A. Schultz, W. Adams, and K. Graves, "Integrating map learning, localization and planning in a mobile robot," in *Intelligent Control (ISIC)*, 1998, pp. 331–336.

[10] H. H. Gonzalez-Banos and J. C. Latombe, "Navigation strategies for exploring indoor environments." *International Journal of Robotics Research*, vol. 21, pp. 829–848, 2002.

[11] C. Stachniss and W. Burgard, "Exploring unknown environments with mobile robots using coverage maps," in *Proceedings of the International Conference on Artificial Intelligence (IJCAI)*, 2003, pp. 1127–1134.

[12] F. Amigoni, V. Caglioti, and U. Galtarossa, "A mobile robot mapping system with an information-based exploration strategy," in *First International Conference on Informatics in Control, Automation and Robotics (ICINCO2004)*, 2004, pp. 71–78.

[13] F. Amigoni and A. Gallo, "A multi-objective exploration strategy for mobile robots," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA2005)*. IEEE Press, 2005, pp. 3861–3866.

[14] B. Yamauchi, "Frontier-based exploration using multiple robots," in *Proceedings of the Second International Conference on Autonomous Agents*, 1998.

[15] W. Burgard, M. Moors, and F. E. Schneider, "Coordinated multi-robot exploration," *IEEE Transactions on Robotics*, vol. 21, no. 3, pp. 376–378, 2005.

[16] Y. Jin, Y. Liao, M. Polycarpou, and A. Minai, "Balancing search and target response in cooperative uav teams," in *Proceedings of the 43rd IEEE Conference on Decision and Control*, 2004.

[17] A. Visser and B. A. Slame, "Including communication success in the estimation of information gain for multi-robot exploration," in *International Workshop on Wireless Multihop Communications in Networked Robotics*, 2008.

[18] D. Calisi, A. Farinelli, L. Iocchi, and D. Nardi, "Multi-objective exploration and search for autonomous rescue robots: Research articles," *Journal of Field Robotics*, vol. 24, no. 8-9, pp. 763–777, 2007.

[19] "Cgal - computational geometry algorithms library." [Online]. Available: http://www.cgal.org/