

A Multi-Objective Exploration Strategy for Mobile Robots

Francesco Amigoni and Alessandro Gallo

Dipartimento di Elettronica e Informazione

Politecnico di Milano

Piazza Leonardo da Vinci 32, 20133 Milano, Italy

amigoni@elet.polimi.it, gallo@airlab.elet.polimi.it

Abstract—Exploration strategies are used to guide mobile robots in building maps of environments. Usually, exploration strategies work greedily by evaluating a number of candidate observation positions on the basis of a utility function and selecting the best one. The utility functions are defined in an *ad hoc* manner as the composition of values measuring different features of a candidate observation position, such as the travelling cost and the estimated information gain. In this paper, we propose a more general way to define an exploration strategy through multi-objective optimization. In our approach, the values of the features are kept separated without combining them in a particular utility function. Experimental results demonstrate the effectiveness of our method.

Index Terms—Exploration strategies. Robot mapping. Multi-objective optimization.

I. INTRODUCTION

Robot mapping deals with building spatial representations of the environments in which the robots operate [1]. A mapping process is often implemented in an incremental way: the map grows as data about the environment are collected at different positions. One of the most important operations that are performed during mapping is the greedy selection of the next best observation position in which the robot will move to acquire new data. An *exploration strategy* determines the observation positions the robot reaches during the mapping process. Usually, exploration strategies evaluate a number of candidate observation positions on the basis of different *features* (including, for example, the travelling cost and the estimated information gain) and select the best one.

In this paper, we propose a multi-objective exploration strategy for mobile robots that determines the next best observation position. Our strategy considers three features to characterize an observation position: the travelling cost, the information gain, and the precision of the localization of the robots. A multi-objective exploration strategy represents a more rational approach to evaluate candidate observation positions than those proposed in literature. Often, these approaches use value functions to measure the features that determine the “interest” of an observation position and then combine the returned values in an *ad hoc* global utility function that is maximized in order to find the next best observation position [2], [3]. Our multi-objective approach provides a more general way to evaluate a position because it does not combine the values of the features in a particular

utility function but keeps them separated. The proposed approach works both for single robots and for multirobot systems. We experimentally validated our exploration strategy by comparing it with other exploration strategies proposed in literature in different kinds of environments.

We represent the map of the environment as a collection of 2D segments. We suppose the robots equipped with range sensors able to acquire spatial information around them. More precisely, the sensors are supposed to be placed at the same height, to be omni-directional with a given maximum range, and to return data represented as collection of segments. For example, the robots can exploit laser range sensors or cameras (with adequate post-processing of the images). The data acquired by a robot in an observation position constitute what we call a *partial map*.

This paper is structured as follows. The next section reviews some exploration strategies to evidence the novelty of our approach. The proposed multi-objective exploration strategy is presented in Section III and is experimentally validated in Section IV. Section V concludes the paper.

II. EXPLORATION STRATEGIES

A “good” exploration strategy for guiding the process of map building should be: *efficient*, to cover the environment as fast as possible, *accurate*, to build a precise and reliable map, and *adaptable*, to work in different kinds of environments (e.g., in large open spaces as well as in cluttered offices). Exploration strategies are usually defined to trade-off the above requirements in choosing the view points during map building. The maps produced by a map building process can have different representations: geometrical (grids [3], points [4], segments [2]) or topological [5]. In this paper, we concentrate on the exploration strategies rather independently from the map representation (we return on this topic in Section V).

Several systems, based on different exploration strategies, have been developed for autonomous robotic exploration; we classify them in two main approaches. Systems of the first class use *predefined trajectories* along which the robot moves and collects observations [6], [7]. The definition of these trajectories needs *a priori* information about the environment to be explored. In general, such information is not available and effective predefined trajectories cannot be constructed.

Moreover, this approach does not easily adapt to different kinds of environments.




Fig. 1. A next-best-view exploration process

Systems of the second class, called *next-best-view* systems [2], [4], [8]–[11], are more flexible in adapting to different environments. In next-best-view exploration, the observation positions, and therefore the trajectories along which the robot moves, are computed during the exploration. Usually, the robot finds, at each step, the next best observation position from which gathering information about the environment. The nature of the process is thus inherently greedy (although some attempts to use a predictive model of the environment have been done [12]). The general structure of a next-best-view exploration process is shown in Fig. 1. The first operation of the process is the selection of the next best observation position. This choice is usually performed by evaluating a set of candidate observation positions through a *utility function*: the candidate with the highest utility value is chosen as the new observation position. The utility function captures in a value the “interest” of an observation position. The next-best-view systems mainly differ in the form of their utility functions.

The *measurable features* that are considered to evaluate a candidate observation position in the utility functions are: the *travelling cost* to reach the position, the expected *information gain* from the position, and, in the multirobot case, the *dispersion* of the robots. Sometimes, since the precise localization of the robot is an important issue in mapping [1], also the expected amount of the *overlap* between the partial map that will be acquired from the observation position and the current global map is considered. For example, this feature is explicitly considered in the utility function in [4] and is used as a threshold to discard candidates in [2] (in [10] exact localization of the robots is implicitly assumed). All these features are measurable in ways that heavily depend on the map representation adopted. For example, information gain is measured as the expected uncertainty reduction on sensed points in [4] or as the expected reduction of entropy of the probabilistic beliefs of the cells involved in a measurement of the sensor [10].

The above features are usually combined in a single utility

function $f(p)$ that globally evaluates a candidate observation position p . The form of $f(\cdot)$ is determined in an *ad hoc* manner or, in few cases, is more mathematically grounded. For example, the simplest utility function is the cost of reaching an observation position [9], [11]. Other systems use utility functions that compose two different features, for example the travelling cost and the expected uncertainty of the map after the observation are composed linearly in [10] or according to a relative entropy approach in [4]. In [2], the cost and the expected new information gained with the observation are considered. More precisely, the utility function is a combination of the cost $c(p)$ of reaching an observation position p from the current position of the robot and of the new obtainable information from p , estimated by the area $A(p)$ of unknown environment potentially visible from p (λ weights travelling cost and information gain):

$$f(p) = A(p) \cdot e^{-\lambda \cdot c(p)} \quad (1)$$

In [8], the travelling cost is split in two distinct features: the Euclidean distance and the angle of the candidate position p to the current position of the robot. These two features are then combined with the estimated information gain from p to obtain the utility function.

The optimization of $f(\cdot)$ gives the next best observation position. The particular structure of $f(\cdot)$ tries to trade-off among the features, because often they cannot be all optimized independently; for instance, given an environment, the information gain is not independent from the travelled distance since “interesting” positions can be far away from the current position of the robot.

Although most of the work in robotic exploration concerns single robot systems, in the last years some multirobot systems for exploration have been proposed, usually employing the next-best-view approach [3], [13]. Multirobot next-best-view systems roughly perform the same operations of Fig. 1 with the difference that pairs robot - position are determined. At each iteration, a robot in the system is chosen and is associated to its next best observation position. As in the single robot case, the new position is the argument that maximizes a utility function. For example, in [3] the utility function is the composition of the cost of reaching an observation position and the benefit of visiting that position, that, in turn, depends on the proximity of other robots. More precisely, the utility function is:

$$f(p, r_i) = u(p) - \beta c(p, r_i) \quad (2)$$

where $c(p, r_i)$ is the cost of reaching the observation position p for the robot r_i , while $u(p)$ is the benefit of visiting p . The benefit decreases with the presence of other robots in the vicinity of p , in such a way that the dispersion of robots in different areas of the environment is stimulated. β weights the relative importance of benefit versus cost.




Fig. 2. An example of map with explored areas (white), unknown zones (grey), obstacle edges (black), and free edges (boundary between white and grey areas)

There are also hybrid approaches that combine predefined trajectories and next-best-view systems. For example, in [14], [15], an exploration strategy that select an exploration path from a parameterized class of paths is proposed.

The multi-objective exploration strategy proposed in this paper aims at setting up a more systematic framework in which the different features can be combined together in a way that is more general than the *ad hoc* utility functions used in literature.

III. A MULTI-OBJECTIVE EXPLORATION STRATEGY

We adopt an approach to exploration that finds the next best observation position with a multi-objective optimization technique. Our technique elegantly keeps separated the features that define the “interest” of candidate observation positions. This separation allows to flexibly change the set of features used to evaluate a position. Our approach fits in the general schema of next-best-view systems of Fig. 1.

A. Map Representation

We store the map in two lists of segments, that represent the obstacles and the boundary of the unexplored area. The first list, called *obstacle list*, stores the segments representing the edges of the observed obstacles; while the second list, called *free edge list*, stores the segments representing the free edges of the map, namely the boundaries between explored and unexplored areas. (Similar data structures are used, for example, in [2].) The map is updated once new sensorial data, namely a partial map (a collection of segments), are acquired by a robot. The lists of segments are updated by inserting the segments of the partial map. During this operation, segments that are collinear and consecutive are fused together to reduce the dimensions of the lists. Additionally, for efficiency purposes, we keep a list of polygons to represent the free space in which the robots can safely move. An example of map is shown in Fig. 2.

B. Candidate Evaluation

The features we use to evaluate a candidate observation position p follow. The *travelling cost* $c(p)$ is simply computed as the total length of the segments forming the path that connects the current position of the robot to p (Fig. 4(a)). A




Fig. 3. Two paths connecting the current position of the robot and a candidate position, in this case the lower path is chosen

number of paths are planned that connect the two positions and that avoid the obstacles by following their edges; among these paths the shortest one is selected, as shown in Fig. 3. The expected *information gain* $i(p)$ is evaluated as the total length of the segments in the free edge list that are visible from the candidate position; considering visible the segments (or their portions) that fall inside the circle modelling the sensorial view of the robot (Fig. 4(b)). Finally, the expected *overlap* $o(p)$ between the new sensorial data and the information already stored in the map is calculated as the total length of segments in the obstacle list that are visible from the candidate position (Fig. 4(c)). Note that all the features can be calculated given the current map.

To adapt our approach to multirobot exploration, we introduce the concept of penalization: in calculating the information gain the length of a segment s belonging to the free edge list is penalized if it is in view of more robots, according to the function:

$$P(r_i) = S(r_i) \cdot \frac{L}{d(s, r_i) + S(r_i)}$$

where $P(r_i)$ represents the penalization due to the robot r_i , $S(r_i)$ is the sensorial range of robot r_i , L is the length of the segment s , $d(s, r_i)$ is the distance between the segment s and the current position of robot r_i . The quantities $P(r_i)$ (for all r_i) are subtracted to the length L of s to obtain the contribution to information gain given by s in presence of other robots. In such a way the dispersion of the robots in different regions of the environment is encouraged. Note that the penalized length of a segment can be negative; although a lower threshold can be imposed, we do not use it in order to maximize the spreading of the robots.

C. The Proposed Exploration Strategy

The algorithm that implements the exploration strategy we propose is described in the following.

(1) The first step is the generation of a set of possible candidate observation positions. We generate these candidates randomly along the free edges of the map. In the implemented system, we generate a candidate observation position for each




Fig. 4. The features used for evaluating a candidate observation position (the circle shows the field of view of the robot sensor)

segment belonging to the free edge list that is longer than a given threshold (1.5 units in our experiments).

(2) Among the candidates, those *Pareto-optimal* [16] are selected, by directly comparing the values of their features. The Pareto-optimal candidates are those that have feature values such that no other candidate has better values for *all* the features. More precisely, we say that a candidate p' is (strictly) *dominated* by a candidate p when $c(p) < c(p')$, $i(p) > i(p')$, and $o(p) > o(p')$. A candidate p is Pareto-optimal if it is not dominated by any other candidate.

(3) Usually several Pareto-optimal candidates exist, but only one new position is needed for an exploring robot, so a choice between the Pareto-optimal candidates must be carried out. We implemented this choice according to the concept of distance from the *ideal solution* [16]. In our case, the ideal solution is an ideal observation position that has the best values for all the features (i.e., best travelling cost, best information gain, and best overlap). Usually, such an ideal position does not exist. We choose the best Pareto-optimal candidate as the one that is nearest to the ideal position, according to the following distance function:

$$D(p) = \sqrt{(c(p) - c_m)^2 + (i(p) - i_M)^2 + (o(p) - o_M)^2}$$

where p is a Pareto-optimal candidate; c_m , i_M , and o_M are the minimum travelling cost, the maximum information gain, and the maximum overlap obtained from the individual optimization (over the Pareto-optimal candidates) of $c(\cdot)$, $i(\cdot)$, and $o(\cdot)$, respectively; and $c(p)$, $i(p)$, and $o(p)$ are the cost, the new information, and the overlap for p , respectively. Note that the selection of the best Pareto-optimal candidate could be done in other ways (see Section V). Moreover, in our approach, this third step is the only one in which the values of the features are fused together to obtain an evaluation of a candidate observation position condensed in a single quantity.

IV. EXPERIMENTAL ACTIVITY

A. Experimental Setting

To test our multi-objective exploration strategy, we included it in a multirobot exploration system. We suppose that a central agent, the strategist, coordinates the actions of several exploring robots, choosing the observation positions that the robots have to visit (this approach resembles that




Fig. 5. Operations of the strategist

of [3]). The operations performed by the strategist during the exploration process are shown in Fig. 5. At each iteration, a robot is selected among those available in the system. Then, the strategist chooses the next best observation position for the robot. The position is assigned to the robot that autonomously navigates to reach it and acquires data from there. Finally, the partial maps collected by the robots in the observation positions are integrated within the global map by the strategist that broadcasts the updated global map to all robots. This multirobot exploration system has been kept as simple as possible because we are interested in comparatively studying the properties of different exploration strategies. Note that, although only a robot is selected at each iteration, this is not a limiting factor. After the initial step and given that the assignment of a position to a robot is faster than robot navigation and data acquisition, it is unlikely that more than one robot will be available at the beginning of an iteration.

We developed a software simulator, in which it is possible to define the environment to be explored, the number of exploring robots, their maximum sensor ranges, and their initial positions. The simulator was implemented in C++, using the LEDA libraries [17] for data structures and graphics.

With this simulator we tested the performances of our multi-objective strategy by comparing it with other next-best-view strategies described in literature. We implemented the three following alternative strategies. The first one, called 'Greedy', uses the simplest utility function: the travelling cost for reaching a new observation position [9], [11]. The second strategy, called 'Latombe', uses the utility function (1). The third alternative strategy we considered, called 'Burgard', is a multirobot strategy: at each step, the strategist chooses the pair observation position p - robot r_i that optimizes the utility function (2). In the original formulation reported in [3], the benefit is reduced according to a probabilistic model, in order to penalize the observation positions which are likely to be within the range of the sensors of more robots. In our implementation of this strategy, after each assignment of an observation position p to a robot r , we penalize the benefit




Fig. 6. The environments used in the experimental activity: *yhdeksan* (a), *kymmenen* (b), and *yksitoista* (c)

of p by subtracting the quantity $\frac{1}{d(p,r)}$, where $d(p,r)$ is the distance between the current position of r and p . We set $\lambda = 0.2$ (in (1)) and $\beta = 0.001$ (in (2)). We experimentally determined these values to obtain reasonable values of the corresponding utility functions (the original papers do not give much details on how to set these values).

The tests were carried out in three different environments. The first one (called *yhdeksan*) is composed of a large room with several small obstacles (Fig. 6(a)), the second environment (called *kymmenen*) is composed of a large room with three large obstacles (Fig. 6(b)), while the last environment (called *yksitoista*) is composed of five rooms connected by narrow passages (Fig. 6(c)). The diagonals of the environments are about 150 units. An example of exploration performed according to our multi-objective strategy with three robots is reported in Fig. 7. Note that the robots spread uniformly in the environment and quickly map it.

B. Experimental Results

We carried out several tests, varying the environment to be explored, the number of robots (1 or 3), their maximum sensors ranges (7 or 15 units), and their initial positions (grouped or sparse in the environment). The list of tests is reported in Table I. For each test, we have performed 10 runs choosing at random the initial positions of the robots and the candidate observation positions (see previous section). The average results over the 10 runs for each test are shown in Table II (best performances are in bold). Two parameters have been considered to compare the performances of the exploration strategies: the number of simulation steps (that is equal to the number of sensing operations) needed to complete the exploration and the total distance travelled by the robots during the exploration (according to [10], [15]). Note that the performances of the exploration strategies drastically improve when the sensor range of the robots is larger. The effects of augmenting the number of exploring robots are not always evident, but this point requires further investigations. In general, the performances are slightly differentiated in the three environments. Our multi-objective strategy is effective when evaluated in terms of number of steps. According to this parameter, the multi-objective strategy is almost always the best one, showing improvements over the other strategies in all the considered exploration situations. This demonstrates the adaptability of our system

test #	environment	# of robots	sensor range	initial positions
1	yhdeksan	1	7	
2	yhdeksan	1	15	
3	yhdeksan	3	7	grouped
4	yhdeksan	3	7	not grouped
5	yhdeksan	3	15	grouped
6	yhdeksan	3	15	not grouped
7	kymmenen	1	7	
8	kymmenen	1	15	
9	kymmenen	3	7	grouped
10	kymmenen	3	7	not grouped
11	kymmenen	3	15	grouped
12	kymmenen	3	15	not grouped
13	yksitoista	1	15	
14	yksitoista	3	7	grouped
15	yksitoista	3	7	not grouped
16	yksitoista	3	15	grouped
17	yksitoista	3	15	not grouped

TABLE I
EXPERIMENTAL TESTS

to different kinds of environments. Considering the distance, our multi-objective strategy does not produce any significant improvement with respect to the other strategies, probably because the multi-objective strategy does not consider the travelling cost as a dominant factor in choosing observation positions. Moreover, only the multi-objective strategy explicitly considers the overlap between the new sensorial data and the old information already in the map. For this reason its performances in terms of travelled distance are penalized. However, the approach can offer great advantages in the localization, which is a fundamental issue in exploration with real robots and which is usually performed by geometrically matching the newly acquired partial map with the current global map. According to a one-way analysis of variance (ANOVA) test [18], the difference between the four methods is not statistically significant ($F = 0.890$ and $p = 0.446$ for the number of steps and $F = 0.931$ and $p = 0.425$ for the distance).

In our experiments, the Pareto-optimal candidates, among which the next best observation position is selected, are about one third of the number of the candidate observation positions that are reachable from the current robot position. The percentage of Pareto-optimal candidates tends to decrease as the exploration goes on. However, these issues require further investigations. The introduction of penalization in our multi-objective strategy generally reduces the total distance travelled since the robots are better dispersed along the free edges of the environment. This is the most evident difference between our exploration strategy and the others we considered.

V. CONCLUSIONS

In this paper we have presented an exploration strategy that determines the next best observation position by combining different features in a multi-objective optimization. The experimental results demonstrated the effectiveness of our approach when compared with other exploration strategies presented in literature.




Fig. 7. An example of exploration (after 0 (a), 3 (b), 9 (c), 14 (d), 15 (e), and 17 (f) steps)

test #	# of steps				distance			
	G	L	B	M	G	L	B	M
1	93.1	92.6	93.4	86.7	709.68	684.31	719.05	681.31
2	22.4	24.2	23.0	20.8	339.91	336.05	340.84	324.86
3	97.4	94.1	104.1	90.9	758.90	707.25	755.64	734.45
4	93.5	91.6	99.0	88.3	740.69	675.65	756.44	721.21
5	24.9	27.3	30.6	25.6	339.84	360.25	351.37	415.84
6	23.0	24.2	24.3	22.8	335.86	338.95	307.94	352.96
7	78.4	77.7	79.2	71.0	589.37	570.11	616.52	549.39
8	22.4	23.3	22.1	22.6	325.18	322.30	325.43	361.20
9	82.0	82.0	87.6	79.2	643.08	646.84	654.21	648.14
10	80.7	80.5	84.0	74.7	680.01	643.55	647.72	629.88
11	24.9	25.5	29.7	25.1	367.18	355.77	358.52	445.82
12	22.1	22.9	24.0	22.1	333.37	318.13	305.82	365.84
13	23.1	25.1	23.1	22.9	438.53	478.27	450.91	559.10
14	95.1	96.5	102.3	94.5	911.00	774.17	900.45	906.97
15	92.0	95.5	98.4	88.6	894.89	823.64	918.11	849.86
16	25.5	26.4	30.6	24.7	502.00	478.29	468.19	513.43
17	22.9	22.9	24.9	21.2	431.28	392.71	402.29	433.89

TABLE II

EXPERIMENTAL RESULTS (G: GREEDY, L: LATOMBE, B: BURGARD, M: MULTI-OBJECTIVE)

Although the experimental activity presented here has been based on simulation (used often for evaluating exploration strategies [3], [10], [15]) we performed also some experiments with a real robot (a Robuter mobile platform equipped with a SICK LMS200 laser range scanner). The real robot substituted one of the simulated robots of our system. To simulate a omni-directional range sensor, the robot rotates on the spot, takes four consecutive scans (with headings separated by 90 degrees), and aligns the scans to obtain a partial map. The main issues that have to be dealt with before using real robots in our current system are: the refinement of the path planning algorithm, with particular attention to collision avoidance, the translation of the simulated navigation to the primitive movements of the robot, and the management of the uncertainty that affects the data acquired by the laser range scanner (this issue impacts, for example, on the fusion of the segments when inserted in the lists that are used to store the map). The preliminary results obtained with real robots confirm the feasibility of our multi-objective approach to autonomous exploration; however, further work is needed. Moreover, future work will address the study of alternative techniques (instead of choosing the candidate nearest to the ideal position) to select the best candidate among the Pareto-optimal ones. An interesting option could be to change the way the best candidate is selected as the map grows; for example, a selection technique could favor the initial dispersion of the robots and another selection technique could increase the precision of the map when a significant part of

the environment has already been mapped. Finally, we aim to investigate the relations between map representation (e.g., segments or grids) and exploration strategies.

REFERENCES

- [1] S. Thrun, "Robotic mapping: A survey," in *Exploring Artificial Intelligence in the New Millennium*, G. Lakemeyer and B. Nebel, Eds. Morgan Kaufmann, 2003.
- [2] H. H. González-Baños and J. C. Latombe, "Navigation strategies for exploring indoor environments," *International Journal of Robotics Research*, vol. 21, no. 10-11, pp. 829-848, October-November 2002.
- [3] W. Burgard, M. Moors, and F. Schneider, "Collaborative exploration of unknown environments with teams of mobile robots," in *Plan-Based Control of Robotic Agents*, M. Beetz, J. Hertzberg, M. Ghallab, and M. Pollack, Eds. Springer-Verlag, 2002, pp. 52-70.
- [4] F. Amigoni, V. Caglioti, and U. Galtarossa, "A mobile robot mapping system with an information-based exploration strategy," in *Proc. ICINCO*, vol. 2, Setubal, Portugal, 2004, pp. 71-78.
- [5] K. Kakusho, T. Kitahashi, K. Kondo, and J.-C. Latombe, "Continuous purposive sensing and motion for 2d map building," in *Proc. SMC*, vol. 2, Vancouver, Canada, 1995, pp. 1472 - 1477.
- [6] E. Bourque and G. Dudek, "Viewpoint selection - an autonomous robotic system for virtual environment creation," in *Proc. IROS*, vol. 1, Victoria, Canada, 1998, pp. 526-531.
- [7] J. Leonard and H. Feder, "A computationally efficient method for large-scale concurrent mapping and localization," in *Proc. ISRR*, Snowbird, USA, 1999, pp. 169-176.
- [8] A. Nüchter, H. Surmann, and J. Hertzberg, "Planning robot motion for 3d digitalization of indoor environments," in *Proc. ICAR*, Coimbra, Portugal, 2003, pp. 222 - 227.
- [9] B. Yamauchi, A. Schultz, W. Adams, and K. Graves, "Integrating map learning, localization and planning in a mobile robot," in *Proc. ISIC*, Gaithersburg, USA, 1998, pp. 331-336.
- [10] C. Stachniss and W. Burgard, "Exploring unknown environments with mobile robots using coverage maps," in *Proc. IJCAI*, Acapulco, Mexico, 2003, pp. 1127-1134.
- [11] C. Tovey and S. Koenig, "Improved analysis of greedy mapping," in *Proc. IROS*, vol. 3, Las Vegas, USA, 2003, pp. 3251-3257.
- [12] B. Stewart, J. Ko, D. Fox, and K. Konolige, "The revisiting problem in mobile robot map building: A hierarchical bayesian approach," in *Proc. UAI*, 2003, pp. 551-558.
- [13] R. G. Simmons, D. Apfelbaum, W. Burgard, D. Fox, M. Moors, S. Thrun, and H. L. S. Younes, "Coordination for multi-robot exploration and mapping," in *Proc. AAAI*, Austin, USA, 2000, pp. 852-858.
- [14] R. Sim and G. Dudek, "Effective exploration strategies for the construction of visual maps," in *Proc. IROS*, vol. 3, Las Vegas, USA, 2003, pp. 3224-3231.
- [15] R. Sim, G. Dudek, and N. Roy, "Online control policy optimization for minimizing map uncertainty during exploration," in *Proc. ICRA*, New Orleans, USA, 2004, pp. 1758-1763.
- [16] C. A. Coello, "An updated survey of GA-based multiobjective optimization techniques," *ACM Computing Surveys*, vol. 32, no. 2, pp. 109-143, 2000.
- [17] LEDA Library, <http://www.algorithmic-solutions.com/enleda.htm>.
- [18] S. A. Glantz and B. K. Slinker, *Primer of Applied Regression and Analysis of Variance*. McGraw-Hill/Appleton & Lange, 2000.