

Moving from ‘how to go there?’ to ‘where to go?’: Towards increased autonomy of mobile robots

Francesco Amigoni, Nicola Basilico, and Alberto Quattrini Li

Abstract Autonomous mobile robots have seen a wide spread development in recent years, due to their possible applications (e.g., surveillance and search and rescue). Several techniques have been proposed for solving the path planning problem, in which a user specifies spatial targets and the robots autonomously decide how to go there. In contrast, the problem of where to go next, in which the targets themselves are autonomously decided by the robots, is largely unexplored and lacking an assessed theoretical basis. In this work, we make a step towards a framework for casting and addressing this problem. The framework includes the following dimensions: the amount of knowledge about the environment the robots have, the kind of that knowledge, the criteria used to evaluate the success of the decisions, the number of decision makers, and the possible adversarial nature of the settings. We focus on applications relative to exploration and patrolling.

Key words: Navigation strategies, autonomous mobile robots, exploration, patrolling

1 Introduction

It is widely recognized that autonomous mobile robots will be increasingly employed in tasks that would be difficult, dangerous, or simply boring for humans. Several steps have been made towards this goal. For example, most mobile robots

Francesco Amigoni
Politecnico di Milano, Italy, e-mail: francesco.amigoni@polimi.it

Nicola Basilico
Dalla Molle Institute for Artificial Intelligence (IDSIA), Switzerland, e-mail: nicola@idsia.ch

Alberto Quattrini Li
Politecnico di Milano, Italy, e-mail: alberto.quattrini@polimi.it

deployed in service applications are currently able to plan a path between their current pose and a target pose, exploiting a plethora of algorithms that have been developed over the years. Indeed, *path planning* has been one of the long-standing topics of mobile robotics and, although some extensions and variants are still actively investigated (e.g., multirobot path planning), has a solid and relatively stable set of pivotal techniques [20, 21]. A naïve path planning problem formulation involves: a starting pose s for the robot, a target pose t for the robot, some knowledge about the environment, and a criterion O to measure path quality. The outcome is a feasible and safe path that connects s to t satisfying O (at least to some degree) and that, basically, answers the question: ‘how to go there’?

Comparatively much less effort has been devoted to investigate the techniques that allow a mobile robot to autonomously determine the target poses t . In some applications, the target pose is provided by the user or imposed by the task (e.g., reach a screw to turn it). However, in many applications of service robots, the target pose is not known in advance but largely depends on the current history of interactions between robots and environments and on the goal of the robots. In such situations, the robots themselves have to autonomously decide where to go, and usually they need to do so at runtime while performing a task.

In this paper, we address the problem of making such decisions, which we call the problem of developing *navigation strategies* for autonomous mobile robots. Navigation strategies answer the question: ‘where to go?’. The main contribution of this work is the proposal of a framework for casting and addressing the issues related to the development of effective navigation strategies and, as such, it is more a conceptual, rather than a technical, contribution. With this work we provide some insights in addressing the problems about navigation strategies, without claiming to entirely solve the problem.

This paper is structured as follows. The next section analyses some examples of navigation strategies with the aim of both surveying some significant areas in which they are employed and showing that each area is currently using its own techniques and methods for defining them. Section 3 introduces our framework for navigation strategies, which attempts at representing a common ground independent of any specific application. The framework is then applied in practice to suggest new research directions for developing better navigation strategies in Section 4. Finally, Section 5 concludes the paper.

2 Navigation Strategies

As motivated in the previous section, the development of fully autonomous mobile robots requires the development of adequate navigation strategies that allow the robots themselves to select the next locations to visit, according to the specific application. In this section, we discuss some examples of navigation strategies, with the aim of introducing the general idea and not of providing an exhaustive survey of the works presented in the literature.

To better focus the situations we address, we provide a very general abstract model of the behavior of a fully autonomous mobile robot while executing a generic task:

- (a) perform some action in the current location,
- (b) decide a location of the environment where to move,
- (c) reach the selected location,
- (d) return to Step (a).

Although over-simplified, the above model evidences some interesting issues. Step (c) involves low-level planning (e.g., path planning and localization), while Step (a) relates to actions specific to the task that the robot executes at the reached location (e.g., acquire sensorial data or move some object). The navigation strategy is involved in Step (b) and, as already said, we will refer to it as the set of techniques that allow an autonomous mobile robot to answer the question ‘where to go next?’, given the knowledge it possesses so far.

In general, a navigation strategy significantly impacts on the task execution’s performance. Therefore, the problem is to define *good* navigation strategies, i.e., strategies that allow the robot to perform its task maximizing some performance metric or criterion.

The major challenges related to this problem mainly derive from two issues. The first one is that the definition of a navigation strategy strongly depends on the robot’s particular task. Compare, for example, a robot employed for exploration of unknown environments with another robot that has to patrol an environment to prevent intrusions (these two cases are discussed in detail below). In the first case, the robot should select the locations to reach such that it can obtain good views of the surroundings to be integrated in its current knowledge of the environment. In the second case, the robot has to account also for tactical issues such as preventing an intruder to predict its movements and elude it. As this example suggests, completely different navigation strategies could be developed for different tasks.

The second issue is related to the goodness of a navigation strategy. Sometimes this concept is intuitively easy to define. For example, in the case of the patrolling robot, one could define the optimal strategy as the one that minimizes the probability for an intruder to break in without being detected. However, in other situations the goodness of a navigation strategy is harder to capture. In the exploration example, different criteria contribute to the goodness of a strategy, e.g., the amount of mapped area in a given time, the total traveled distance, the quality of the obtained map, the time spent to map a given area, and many others. In general, searching for the best strategy according to some metric increases the problem’s difficulty with respect to the case in which a sub-optimal strategy is employed.

Although these application-dependent specific issues, in the next section we will attempt to define a general framework for developing navigation strategies. In the remaining of this section, we discuss more deeply two specific examples of navigation strategies. They have been selected only based on the expertise of the authors and are not intended to endorse any claim of generality or importance. Other examples include continuous monitoring of temporal-spatial phenomena in partially known

environments (e.g., ocean monitoring [28]), olfactory exploration [22], tactile exploration of object properties [25], information gathering [27], and pursuit/evasion [13].

2.1 Exploration of Unknown Environments

Consider the following problem: an autonomous mobile robot is deployed in an initially unknown (static) environment with the aim of discovering its physical structure, namely of building its map [30]. A map of an environment represents the location of the obstacles and of the free space. The robot is equipped with sensors (e.g., laser range scanners) that allow it to acquire spatial data in its surroundings. It moves, repeatedly sensing different portions of the environment, and builds a corresponding map. This exploration problem (and its immediate variants, also including multiple robots) is fundamental for several applications, ranging from classical map building, to search and rescue [29], and to coverage [11].

One of the main challenges of exploration of unknown environments is related to the navigation strategy, namely where to move to perform the next sensing action in partially-known environment? Consider, as an example, the situation represented in Fig. 1.

The mainstream approach to define navigation strategies for exploration is to pick, from a set of candidate locations (usually obtained by sampling the frontier between known and unknown space, as shown in Fig. 1), the best one that maximizes an utility function that combines different evaluation criteria. This selection




Fig. 1 A partially explored environment. Black line segments are the known part of the environment, while the grey ones are still unknown to the robot (red dot). Dotted red lines are the frontiers between known and unknown portions of the environment. The robot has to decide which location it should reach for performing the next perception among the three available possibilities (green, blue, and orange dots).

is performed at each step of the exploration process, within the so-called *next-best-view* approach.

Several proposals have been made for evaluation criteria and utility functions. For example, the work in [16] proposes the following utility function to evaluate a candidate location p :

$$A(p) \exp(-\lambda d(p)) \quad (1)$$

where $A(p)$ is the expected information gain that the robot can get in p (measured as the maximum amount of unknown area visible from p), $d(p)$ is the distance between the current position of the robot and p , and λ is a parameter that weights the two criteria.

The system proposed in [32] uses the following utility function:

$$\frac{A(p)P(p)}{d(p)} \quad (2)$$

that considers also the estimated probability $P(p)$ of a successful communication between the robot (once at p) and a fixed base station.

In [9], the authors use the following linear utility function:

$$U_{t'} - \beta \cdot V_{t'}^{i'} \quad (3)$$

where $U_{t'}$ and $V_{t'}^{i'}$ are similar to the expected information gain and travelled distance, respectively, used by the works presented above, and β is the relative weight of benefit versus cost of candidate location t' for robot i' .

Authors in [31] use an utility function that evaluates a sequence of m candidate locations (poses) with the following multiplicative function:

$$\sum_{i=1}^m (\exp(lv_i - sv_i) \prod_{j=1}^{q_i} (\frac{\exp(-|\theta_j|)}{\sqrt{s_j} + 1}) \times (\frac{1}{n_i} \sum_{k=1}^{n_i} p_k + Ne_i) fmin_i(d_l)) \quad (4)$$

where q_i is the total number of robot stops to reach location i , lv_i is the length of the closest frontier edge at i , s_j is the distance from the robot to the next possible location j , sv_i is the distance from the next possible location i to the closest frontier edge, θ_j is the orientation change to reach the next configuration j , p_k identifies the probability of viewing landmark k at i , n_i and Ne_i are the number of landmarks and of corners inside a visibility region at i , and $fmin_i$ is a function that penalizes i when it is too close to an obstacle, according to distance d_l .

While all the above navigation strategies (and most of others proposed in the literature) are rather *ad hoc*, we recently developed a more general and flexible approach based on Multi-Criteria Decision Making (MCDM) that provides a solid approach, grounded in decision-theoretical field, for selecting the next candidate location in exploration [7].

2.2 *Patrolling Environments to Prevent Intrusions*

A patrolling problem involves a robot that is deployed in a (usually known) environment with the objective of protecting it from intrusions. Intruders can come from outside the environment (e.g., a building). The robot is equipped with sensors (e.g., cameras) that allow it to detect the presence of a possible intruder. It moves, repeatedly monitoring different portions of the environment. The problem of patrolling environments with autonomous mobile robots has attracted the attention of many researchers both within the Autonomous Robotics and the Multiagent Systems communities [2, 14]. It can be considered as an instance of a general class of problems related to surveillance and security [19] and can be generalized from detecting intruders to detecting any kind of threats.

The challenge of finding the navigation strategy that maximizes the protection level of a given environment has been mainly tackled adopting an explicit model of the intruder (its preferences over the possible targets, its knowledge on the patroller's strategy). Note that, the stance that is often taken is relative to a worst-case scenario, in which the intruder knows the strategy of the patroller (but not its actual actions, since the strategy is usually randomized).

Given this model, a game theoretical perspective is adopted [19]: the patroller (i.e., the autonomous mobile robot) and the intruder are considered to play a strategic game that develops along a discrete timespan. At each time step, one of the two players performs an action (for the patroller, moving from one location of the environment to another location; for the intruder, entering or not the environment) and the combination of their actions determines the final outcome of the game (e.g., the intruder has been detected or the intruder successfully entered the environment). It can be shown that the optimal navigation strategy for the patroller (namely, the navigation strategy that minimizes the probability of an intruder to break in) is equivalent to a particular equilibrium, called *leader-follower*, of such game [24].

In this context, we recently provided an approach that calculates deterministic and randomized strategies in environments with any topology, allowing the adversary to always know the current position of the patroller. A set of heuristic techniques were also defined with the objective of reducing the computational effort in settings of realistic size [8].

3 Towards a Framework for Navigation Strategies

In this section, we introduce a framework in the attempt to capture the main issues of navigation strategies in a way that is independent of the specific applications. In general, a robot needs a navigation strategy to select its target locations using its current knowledge about its environment and task (Fig. 2). Almost all navigation strategies take as input a *state* enclosing task-related information about the environment (e.g., in exploration usually it is a map of the currently explored space and the current position of the robot) and provide the locations to reach as output.

As the examples of Section 2 partially show, several techniques have been proposed to tackle the problem of the definition of good navigation strategies, including decision-theoretical techniques, (PO)MDPs, game-theoretical techniques, and many others, but no general framework for the navigation strategy problem has been proposed so far.

Navigation strategies can be characterized along many dimensions, here we discuss some of those, being aware that the list is far from being definitive and complete.

- Offline vs. online strategies.
- Global vs. partial knowledge of the environment.
- Kind of knowledge used to make decisions.
- Optimality criterion.
- Number of decision makers.
- Adversarial vs. non-adversarial scenario.

A first distinction can be done between *offline* and *online* strategies. In the first case, the locations to reach are computed for every possible input state before the robot actually executes the task, i.e., before the robot actually employs the strategy. With an online strategy, instead, the decision is computed during the task execution for the different situations the robot encounters.

This distinction can also be described with respect to another dimension, namely the amount of available information about the environment. If the environment (or, more precisely, all the information needed for making decisions) is fully known in advance, the robot has a *global knowledge*. Conversely, if only partial or no environment’s information is initially available, the robot has a *partial knowledge* and should increase its knowledge to make more informed decisions. The availability of global knowledge results in the possibility of computing the strategy offline and, possibly, searching for an optimal solution. On the other hand, a partial knowledge is typically associated with the use of an online navigation strategy where sub-optimal algorithms are employed. As an example, consider the two common tasks of coverage and exploration. In coverage [11], the environment is known in advance and the robot should cover (possibly, under some constraints) all the free area with the footprint of its sensors. In exploration, the environment is unknown at the beginning and the robot has to “discover” it. The first case is characterized by a global knowledge and the optimal strategy, e.g., the shortest route, can be computed offline. The




Fig. 2 The intuitive idea of navigation strategy.

second case is an example of partial knowledge situation for which decisions have to be made online, due to the impossibility to predict the states that the robot will face. The optimal strategy cannot be found in general and sub-optimal greedy algorithms (e.g., next-best-view approaches) must be employed. Note that in both cases the optimal strategy can be the one that minimizes the travelled distance.

The *kind of knowledge* used by the robot to make decisions can be of different natures. A specific kind of knowledge induces how decisions are made. For example, in exploration, criteria used in the utility function usually derive from the robot's knowledge about the spatial features that has been collected in the map built so far. However, in principle, there are other types of knowledge that can be used, like high-level information coming from human users.

The *optimality criterion* is a crucial dimension for the design of a navigation strategy, which is evaluated according to the selected optimality criterion. It is strongly dependent on the task the robot should perform, but also on the robot capabilities. Sometimes, this concept is easy to define, like in the patrolling case, in which the optimal solution is the one that minimizes the probability for an intruder to break in. However, in some applications, like search and rescue, it is hard to define which optimality criterion should be selected to measure the performance of the system, because there are more criteria that should be optimized. They are sometimes conflicting, as some measure the benefits obtained, while others measure costs to perform the assigned task (e.g., explored area vs. robot's battery and travelled distance). In such a case, a trade-off between benefits and costs should be addressed. Furthermore, in other cases the optimality criterion is not naturally evident. For example, in a pursuit/evasion scenario, there is not any main optimality criterion, but only feasibility (i.e., either the evader is caught or not caught).

The *number of decision makers* (robots) is another dimension. The presence of multiple agents can pose significant difficulties. Determining a navigation strategy for a team of robots has to deal with the exponential growth (in the number of robots) of the number of possible actions, and usually involves a task-assignment problem [15], where the task is a location to be reached by a robot. Multiple robots can cooperate to make a globally optimal decision or can compete to make individually optimal decisions.

The presence of multiple agents introduces another dimension represented by the *adversarial nature* of the setting, which is related to the possible presence of adversaries. An adversary can be defined as a rational agent acting against the robot's objectives, and whose interaction has to be considered when computing the navigation strategy. Intuitive examples of these last two dimensions can be found in patrolling. An adversary, i.e., a possible intruder, can be considered by the patrolling robot in deciding where to move for protecting the environment. In this case, a competitive interacting scenario emerges and game theoretical techniques are employed to find navigation strategies (see Section 2.2).

Let us apply the general framework just defined to the two examples of Sections 2.1 and 2.2. In the case of exploration, the navigation strategy is online and based on a partial knowledge of the environment. Decisions are made from knowledge embedded in a metric map of the environment (consisting of obstacles, free

space, distances, ...) and the optimality criteria considered include travelled distance or number of perceptions to fully explore the environment, area discovered in a given time interval, In most of the examples presented in Section 2.1, there is a single decision-maker, but there can be several ones, and the scenario is clearly non-adversarial.

In the case of patrolling, the navigation strategy is calculated offline based on the global knowledge of the environment. Decisions are made from knowledge of the environment and of the adversary and the optimality criteria adopted are related to the probability of an intruder to break in in the locations of interest of the protected environment. There is a single decision-maker that operates in an adversarial scenario.

4 Generating Novel Research Directions

The framework proposed in the previous section can stimulate novel ideas and research directions for enhancing the development of navigation strategies by facilitating the identification of significant open issues.

For example, according to what we discussed above, in developing navigation strategies for exploring initially unknown environments, the decision about the next destination location is usually based only on the information provided by a *metric map* of the environment. A metric map represents the geometrical and spatial features of an environment, like the position of obstacles. All the works presented in Section 2.1 use a metric map to extract the values for criteria they use in the utility function, like the expected information gain or the distance between a robot and a candidate location. One interesting, and potentially effective, extension can be to consider also information contained in *semantic maps* of environment. A semantic map associates spatial concepts to spatial entities. For example, a room identified in the metric map can be labeled as ‘kitchen’ or as ‘bathroom’ using one of the several methods that have been proposed (e.g., [23, 33]). Moreover, *a priori* information about the task can be provided by the user and usually biases the behavior of the robot. For example, in a search and rescue operation, such information can be used to prefer the selection of candidate locations in offices if a disaster happened during working hours, because the user thinks that it is more likely to find victims there than in other places (e.g., in a canteen). Combining high-level knowledge of places in buildings (provided by semantic maps) and information coming from users for developing better exploring robots proves to be an interesting and largely uninvestigated research line, which we have preliminarily tackled in [12].

Another research direction stimulated by the framework is about experimental methodologies. A lively debate on good experimental methodologies is currently ongoing in the autonomous robotics community, which has recognized that the lack of sound procedures, globally accepted benchmarks, and well-established metrics for autonomous robotic systems slow down research and make the transfer of existing research results to market products rather difficult [3, 5]. For example, ex-

perimental evaluation of navigation strategies for exploration is problematic [6]. One relevant question is about which data should be used for evaluation. In real environments, evaluation can be clearly performed only *ex post*, once the environment has been completely mapped. Data sets acquired by real robots (like those on Rawseeds [10] and Radish [18]) are usually not so dense to cover every perception in every pose of the environment and cannot be used to evaluate navigation strategies for exploration. Simulation is sometimes the only viable solution [4]. Another question is about which metric should be used for evaluation. This is related to the difficulty of establishing an optimality criterion that we evidenced for the corresponding dimension of our framework. The number of sensing actions performed by the robot and the total amount of distance travelled to completely explore an environment are two common metrics, but there currently is no agreement on this [6]. Another important issue, which has been largely overlooked so far, is the availability of a reference (optimal) performance for conducting the experimental evaluation. Without such a reference performance, navigation strategies for exploration can be compared only in a relative way, against each other, and not in an absolute way, according to their distance from optimum. An initial attempt to overcome this problem has been presented in [26].

Other possible research directions that are expected to help developing better navigation strategies include the integration of human decisions within the framework, leading to mixed-initiative systems [17]. For example, in a search and rescue scenario, a robotic system could exchange knowledge and cooperate with human rescuers to more effectively find victims in a collapsed building; while, in exploration, an human operator can remotely bias the behavior of a team of robots by dynamically assigning priorities to portions of the environment (see above). Some of the challenges holding in these scenarios involve the design of human-robot interaction frameworks as well as the definition of mixed decision-theoretic layers, where human decisions could be combined with those made by robots in an effective way [1].

5 Conclusions

In this work, we proposed a general abstract framework to coherently address the analysis and the development of navigation strategies, which allow mobile robots to autonomously decide target poses. The framework has been defined and shown useful to stimulate new research directions, towards better navigation strategies.

The proposed framework will be further enhanced, especially providing a more formal decision-theoretical definition of its basic elements in order to pave the way towards the easy development and the theoretical and experimental assessment of good navigation strategies that make mobile robots more autonomous.

References

1. Adams, J., Rani, P., Sarkar, N.: Mixed initiative interaction and robotic systems. In: AAAI Workshop on Supervisory Control of Learning and Adaptive Systems (2004)
2. Agmon, N., Kaminka, G., Kraus, S.: Multi-robot adversarial patrolling: facing a full-knowledge opponent. *J Artif Intell Res* **42**(1), 887–916 (2011)
3. Amigoni, F., Reggiani, M., Schiaffonati, V.: An insightful comparison between experiments in mobile robotics and in science. *Auton Robot* **27**(4), 313–325 (2009)
4. Amigoni, F., Schiaffonati, V.: Good experimental methodologies and simulation in autonomous mobile robotics. In: L. Magnani, W. Carnielli, C. Pizzi (eds.) *Model-Based Reasoning in Science and Technology, Studies in Computational Intelligence*, vol. 314, pp. 315–332. Springer (2010)
5. Amigoni, F., Schiaffonati, V.: Autonomous mobile robots as technical artifacts: A discussion of experimental issues. In: L. Magnani (ed.) *Model-Based Reasoning in Science and Technology, Studies in Applied Philosophy, Epistemology and Rational Ethics*, vol. 8, pp. 527–542. Springer (2014)
6. Basilico, N., Amigoni, F.: On evaluating performance of exploration strategies for an autonomous mobile robot. In: IROS Workshop on Performance Evaluation and Benchmarking for Intelligent Robots and Systems (2008)
7. Basilico, N., Amigoni, F.: Exploration strategies based on multi-criteria decision making for searching environments in rescue operations. *Auton Robot* **31**(4), 401–417 (2011)
8. Basilico, N., Gatti, N., Amigoni, F.: Patrolling security games: Definition and algorithms for solving large instances with single patroller and single intruder. *Artif Intell* **184–185**, 78–123 (2012)
9. Burgard, W., Moors, M., Stachniss, C., Schneider, F.: Coordinated multi-robot exploration. *IEEE T Robot* **21**(3), 376–386 (2005)
10. Ceriani, S., Fontana, G., Giusti, A., Marzorati, D., Matteucci, M., Migliore, D., Rizzi, D., Sorrenti, D., Taddei, P.: Rawseeds ground truth collection systems for indoor self-localization and mapping. *Auton Robot* **27**(4), 353–371 (2009)
11. Choset, H.: Coverage for robotics: A survey of recent results. *Ann Math Artif Intel* **31**(1–4), 113–126 (2001)
12. Cipolleschi, R., Giusto, M., Quattrini Li, A., Amigoni, F.: Semantically-informed coordinated multirobot exploration of relevant areas in search and rescue settings. In: *Proc. ECMR* (2013)
13. Durham, J., Franchi, A., Bullo, F.: Distributed pursuit-evasion without mapping or global localization via local frontiers. *Auton Robot* **32**(1), 81–95 (2012)
14. Fang, F., Jiang, A., Tambe, M.: Optimal patrol strategy for protecting moving targets with multiple mobile resources. In: *Proc. AAMAS*, pp. 957–964 (2013)
15. Gerkey, B., Mataric, M.: A formal analysis and taxonomy of task allocation in multi-robot systems. *Int J Robot Res* **23**, 939–954 (2004)
16. González-Baños, H., Latombe, J.C.: Navigation strategies for exploring indoor environments. *Int J Robot Res* **21**(10–11), 829–848 (2002)
17. Horvitz, E.: Reflections on challenges and promises of mixed-initiative interaction. *AI Mag* **28**(2), 3 (2007)
18. Howard, A., Roy, N.: The robotics data set repository (Radish). <http://radish.sourceforge.net/> (2003)
19. Jain, M., An, B., Tambe, M.: An overview of recent application trends at the AAMAS conference: Security, sustainability and safety. *AI Mag* **33**(3), 14–28 (2012)
20. Latombe, J.C.: *Robot Motion Planning*. Kluwer Academic Publishers (1991)
21. LaValle, S.: *Planning Algorithms*. Cambridge University Press (2006)
22. Marjovi, A., Marques, L.: Multi-robot topological exploration using olfactory cues. In: *Proc. DARS*, pp. 47–60 (2013)
23. Mozos, O., Stachniss, C., Burgard, W.: Supervised learning of places from range data using AdaBoost. In: *Proc. ICRA*, pp. 1742–1747 (2005)

24. Paruchuri, P., Pearce, J., Tambe, M., Ordonez, F., Kraus, S.: An efficient heuristic approach for security against multiple adversaries. In: Proc. AAMAS, pp. 311–318 (2007)
25. Pezzementi, Z., Plaku, E., Reyda, C., Hager, G.: Tactile-object recognition from appearance information. *IEEE T Robot* **27**(3), 473–487 (2011)
26. Quattrini Li, A., Amigoni, F., Basilico, N.: Searching for optimal off-line exploration paths in grid environments for a robot with limited visibility. In: Proc. AAAI, pp. 2060–2066 (2012)
27. Singh, A., Krause, A., Guestrin, C., Kaiser, W.: Efficient informative sensing using multiple robots. *J Artif Intell Res* **34**(1), 707–755 (2009)
28. Smith, R., Schwager, M., Smith, S., Rus, D., Sukhatme, G.: Persistent ocean monitoring with underwater gliders: Towards accurate reconstruction of dynamic ocean processes. In: Proc. ICRA, pp. 1517–1524 (2011)
29. Tadokoro, S.: *Rescue Robotics*. Springer-Verlag (2010)
30. Thrun, S.: Robotic mapping: A survey. In: *Exploring Artificial Intelligence in the New Millennium*, pp. 1–35. Morgan Kaufmann (2002)
31. Tovar, B., Munoz, L., Murrieta-Cid, R., Alencastre, M., Monroy, R., Hutchinson, S.: Planning exploration strategies for simultaneous localization and mapping. *Robot Auton Syst* **54**(4), 314–331 (2006)
32. Visser, A., Slamet, B.: Including communication success in the estimation of information gain for multi-robot exploration. In: Proc. WiOPT, pp. 680–687 (2008)
33. Wolf, D., Sukhatme, G.: Semantic mapping using mobile robots. *IEEE T Robot* **24**(2), 245–258 (2008)