Cooperative Exploration

Exploring an unknown environment with a team of robots

Facundo Benavides Olivera
fbenavid@fing.edu.uy
Facundo.BENAVIDES-OLIVERA@isae.fr

PEDECIBA Informática
Universidad de la República
Uruguay

École doctorale EDSYS
ISAE-Supaero
Université de Toulouse
France

Draft
Toulouse, 16/09/2015

Abstract

In this document a description of the cooperative exploration problem is presented and a set of surveyed jobs -since September/14 until April/15- is briefly described.

Additionally, certain future works and research lines are expressed at the end.

Contents

Chapter 1

Introduction

Many robotic missions can be executed more efficiently and robustly if they are conducted by teams of robots[1]. Together, the robots can achieve the goal faster than a single robot. In addition, working in team adds redundancy. By eliminating the single point of failure, it makes the system more tolerant of damage or malfunction of any of its members. Another advantage is that it is possible to exchange superimposed information helping to reduce mapping and localization uncertainty performed on each robot. [BMSS05]

By contrast, in presence of active sensors, the existence of multiple robots can cause -due to interference between readouts- a worsening overall system performance. In addition, the more members the team has, the more complex path planning and the longer detours (in order to avoid collisions) will be necessary. [BMSS05]

In this sense, to be able to build a global map of the environment (mapping) through the exchange of local representations of the environment (local maps) is a fundamental skill that a system must have to enhance the effective implementation of cooperative activities. [VCBL09]

In turn, the construction of maps of unknown environments is only solvable accomplishing the total coverage of those environments. In this case, the exploration problem can be viewed as a subsidiary task to the mapping problem.

However, the problem of exploration -by itself- is considered as one of the fundamental problems in autonomous mobile robotics. There are several scenarios where achieving complete coverage of a zone is one of the main parts of the mission: planetary exploration, reconnaissance, rescue, agriculture, cleaning or the exploration of dangerous places as mined lands, radioactive zones, etc. [Bau13]

In these cases, it is essential that the robots be able to keep a record of the visited places in order to maximize the utility of the system avoiding revisit known places. Moreover, the robots must perform the mapping of the environment since it is necessary to build navigation plans and coordinate their actions. [BMSS05]

The mapping task can be identified as a capability that supports the exploration task.

1.1 Outline

The document is organized as follow. Chapter 2 presents a description of the cooperative exploration problem. There, some definitions are given and the main aspects to be analyzed are expressed. In chapter 3, a set of surveyed jobs -grouping the most important or widely known ones- is briefly described. Finally, the future work (including certain research lines) is presented in chapter 4.

[1]Henceforth, the terms robot, agent or member of a team will be used interchangeably.

Chapter 2

Problem definition

The exploration of a unknown environment can be defined from the interaction between the following three problems: localization, mapping and motion planning. In figure 2.1 those interactions can be seen.

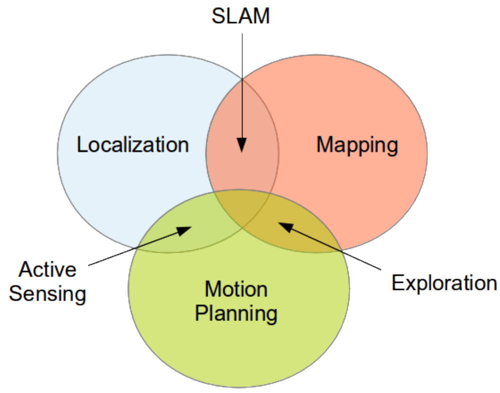

Figure 2.1: Sub problems of exploration.

The Exploration problem is composed of Mapping and Motion Planning tasks. SLAM implies that Localization and Mapping are carried out simultaneously. Active sensing or Active Localization refers to motion planning guided by localization through the locations that increase the accuracy of localization.

Additionally, the cooperative exploration problem is concerned to the full covering of an unknown environment -by the hand of a team of robots- spending the minimum amount of time. [BMSS05]
In fact, the overall exploration time is the most commonly used criteria to measure the quality of proposals.

However, to do that, the robots need to coordinate their actions and build a common map of the whole environment. Thus, the robots must be able to build maps on-line while they are in motion. As a consequence, mapping is constantly interleaved with motion planning, and vice versa. In fact, to take good decisions they need a map but to build a good map they need to take good decisions.

Therefore, and in the context of multi-robot exploration, two key coordination problems must be solved:

- task assignment: how to coordinate the robots in order to efficiently cover the environment.
- collaborative mapping: the robots have to construct a global map in order to plan their paths and coordinate their actions.

Furthermore, other aspect must be considered: *communication*. For cooperating during exploration, the robots must to be communicated. So, communication is necessary to perform both of the coordination tasks and depending on its characteristics will be the proposal solutions that be applicable in each case.

As a conclusion, by the one side, the Mapping and Motion Planning problems become more complex than in the single robot case and, by the other, many issues can arise because of communication constraints.

2.1 Task assignment

Choosing the next place to be visited often takes into account, by the one side, the potential utility that would be obtained by visiting this place and, by the other side, the cost of performing the path. In the case of a multi robot the purpose of coordination system is to determine which task (place to visit) must be performed by each member to maximize the overall utility of the system while minimizing the amount of overlapped information obtained by all of them. [BMSS05]

When multiple robots are involved, to avoid several of them moving to the same place is desirable. To determine appropriate target location for individual robots, usually, the task assignment is addressed in the following order [WSB08]:

1. task identification: is concerned to identify the points of interest.
   It strongly depends on the sensory capacities of the robots and on the environment representation.

2. task allocation: is concerned to task assignment pursuing a distribution of tasks to robots that maximize the overall system utility.

2.2 Cooperative mapping

It concerns to the ability to build a single map from the "local" maps that is built by each of the team members.

Be able to build a global map of the environment (mapping) through the exchange of local representations of the environment (local maps) is a fundamental skill that a system must have to enhance the effective implementation of cooperative activities. Therefore, building and sharing environment representations are key features.

For this purpose, besides having the ability to build a map, each robot must also be able to communicate with the other ones (this possibility could be not always present) and, even being able to communicate, it must be defined which are the appropriate moments to exchange information. [VCBL09]

2.2.1 Centralized

Centralized mapping is performed by the asynchronous integration of sensory readings sent by the agents, building a single map with information from all the explored areas. This approach requires that a initial condition must be respected: all the robots must be together and able to exchange relative position between each others. Furthermore, some of the agents must play the role of cartographer. So, under the computing point of view, this approach is conducted in a partially distributed manner, but, at the end, a unique global map is built and delivered by the cartographer. [SAB+00, BMSS05]

2.2.2 Distributed

Distributed mapping, by contrast, does not requires neither the robots starting in the same place nor one of them playing the role of cartographer. However, the situations (or events) where be appropriate and feasible to make the exchange of information must be identified. [VBSL11, VCBL09]

Chapter 3

State of art approaches

In this chapter a brief abstract of each one of the surveyed approaches is presented. However, to facilitate the reading, in some cases, they are grouped in sections to describe common aspects, advantages or disadvantages.

In that sense, and as is proposed in [Bau13], the analysis will be conducted considering the hypothesis that are often made, objectives pursued and the methods employed. In addition, the approaches can be classified according to: hypothesis, objectives, identification methods, allocation methods and communication methods.

3.1 Hypothesis of work

3.1.1 Environment and initial conditions

At first, it is important to note that the environments heavily impact on the exploration strategies. Generally, the environments defined to develop those proposals and experiments have the following common characteristics: unknown, partially observable, stochastic, sequential, continuous, static and multi-agent[1].

However, probably the most important aspect is related to the distinction between indoor and outdoor environments.

The indoor environments[2] are commonly structured and the risk of collisions between robots is usually bigger. Furthermore, the use of certain kind of sensor could be infeasible or problematic (Global Positioning System - *GPS*, Ultra Sound - *US*).

By contrast, the outdoor environments[3] use to be bigger than the former ones and unstructured. The use of *GPS* is available and recommended but the exchange of mapping information is usually harder because of the absence of many common points of reference.

Another important hypothesis refers to initial conditions. More precisely, the team of robots can begin the exploration task from a common place or not. In the first case, it is normally assumed that they have a common frame of reference and that they are able to exchange positioning information as part of an initialization task of the system. In this scenarios, the exchanging of mapping information during exploration is usually simple.

On the contrary, if the robots start in different locations, a more complex process will be needed to allow an appropriate exchange of mapping information.

3.1.2 Robot capabilities

When the robot capabilities are under analysis, there are many aspects that can affect the choice of the exploration strategy. The computing power of the robots -often smaller compared with other computing stations- can limit the mapping algorithms that can be used, battery capacity and consumption level can be problematic if long distances will be covered, communication aspects (as band width, scope and reliability) can definitively condition the task allocation process and the exchanging of mapping information as well. Finally, the sensory capabilities (sensors nature,

[1]Considering the definition provided in [RN10].
[2]Typically bounded.
[3]Commonly unbounded but could be also bounded.

scope, accuracy, resolution and acquisition frequency) play a very important role in every task which depends on observation (localization, mapping, recognition).

3.2 Exploration Objectives

Depending on the underlying problem that the exploration comes to solve, the objectives can vary pretty much. If the main purpose is related to searching objects, persons or places of interest (survivors or fire sources for example) the priority will be put over coverage to increase -by this way- the probability of find them quickly.

However, if surveillance is the main purpose, the environment must be explored several times. On the other side, when cartography is the main purpose, instead of minimize the total spent time, the accuracy of the resultant map is the main objective.

3.3 Task assignment

Most of the time, the problem of assigning tasks is often dealt mainly considering the resolution of both task identification and task allocation -in that order. [WSB08]

3.3.1 Identification methods

Task identification concerns to the identification of environment points or areas that be potentially interesting for the system. It heavily depends on the sensory capacities of the robots and on the underlying environment representation too. The most widely used representations are occupancy grids and *Voronoi* diagrams. While the first one is characterized by being independent of the shapes of obstacles and only usable in closed environments, the second one can also be used in open environments but is dependent on the forms.

A method for identifying points of potential interest was proposed by *Yamauchi* in 1998 [Yam98] and consists of identifying the environmental points that are in the borderline between the known and unknown areas. Since then, those points are well known as *frontier points*. In addition to this method -which is very popular until our days, currently, there are others which are having a lot of attention and are based on the segmentation of the environment. These approaches divide the known environment in areas or segments for which the robots be allocated for their recognition. [WSB08]

3.3.2 Allocation methods

Task allocation algorithms have to determine the optimal distribution of tasks among team members. Typically, these algorithms seek to maximize an objective function where the utility associated with each task is represented. The utility is usually considered as a trade off between the hope of the information gain at task's completion time and the cost of undertaking it.

3.3.2.1 Auctions

It consists in the centralized selection of tasks considering the information gain in descending order. Each robot computes the information gain and cost per task and sends a bid communicating which is their order of preference to perform the pending tasks. The bids are sent only when robots receive a new global map update. Therefore, the cost of computing the relative utility of each task from each robot is distributed throughout the system. Only the allocation decision is processed centrally by a greedy algorithm. For this reason, at least one agent must assume the role of executive, distributing the tasks centrally. Every time he receives a bid, he waits some time -during which he expects more bids- to make the decision including as many robots as possible. To avoid several of them moving to nearby places, when a point is assigned, the gain of all their neighbors are discounted by an amount proportional to the distance to the point allocated. [SAB+00]

In figure 3.1 an example where three frontier points were identified. The circles indicate sensors range.

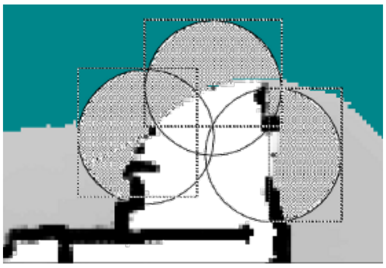

Figure 3.1: Example of modeled environment under Auction approach. [SAB+00]

3.3.2.2 Genetic algorithms

In spite of having some problems, the approaches based on auctions are widely used. Those strategies, in one end (single-item auctions), depending on the existence of synergies between team members and the environmental objects during the exploration may result in a highly sub-optimal distribution of tasks, and in the other (combinatorial or bundle-item auctions), which consider multiple items each time, could present performance issues since the number of combinations to be evaluated grows up exponentially with the number of tasks. Thanks to well known properties, particularly random search, the genetic algorithms approach is presented as an alternative to the ones based on auctions (see 3.3.2.1). However, this is a centralized approach and requires that the number of robots during all the exploration process be invariant. [MZL07]

3.3.2.3 Occupancy grids

The allocation of tasks is performed centrally, directly based on the knowledge of the environment (provided by a global map: the occupancy grid) and the location of robots and frontier points detected so far.

The main difference between this approach (presented in [BMSS05]) and the one presented in [SAB+00] is that all computation is done centrally.

3.3.2.4 Hungarian method

It is a centralized method of task allocation proposed by Kuhn in 1955 [Kuh55]. Originally, it was proposed to distribute jobs in machines but is easily adaptable to the problem of task allocation to robots. Requires that the number of machines (robots) and jobs (tasks) be identical. It is an iterative algorithm that process a matrix whose order of execution time is $O(n^3)$, being n the number of machines. [WSB08]

Additionally, this method is used in combination with a environmental segmentation technique. The underlying idea is that the allocation of tasks is performed considering certain segments of the environment where there are unknown areas. The working hypothesis is that in highly structured environments is more convenient performing the exploration after dividing the environment into disjoint segments. Then, each robot is sent to explore different segments. This way, is expected to achieve full exploration decreasing the overlap between agents as much as possible.

In [WSB08] the environmental segmentation is done by the hand of *Voronoi* diagrams. *Voronoi* diagrams are geometric structures widely used in robotics for building road maps. [Roq96, FGLM01, RD05, WvH07, BG08] Particularly, the fact that -by construction- the structure of the environment is reflected in the diagram itself (depending on the shape of the present objects) is interesting.

The procedure consists of:

- to build an occupancy grid
- to build a *Voronoi* diagram from the occupancy grid
- to process the *Voronoi* diagram eliminating some nodes and retaining only the ones -called critical points [Thr98]- which allow to partitioning the diagram in segments.

In figure 3.2 can be appreciated an segmented environment.

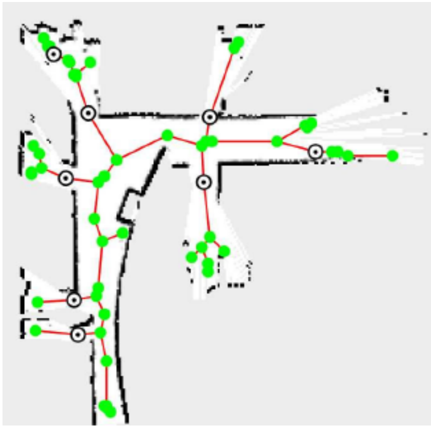

Figure 3.2: Segmentation example. [WSB08]

3.3.2.5 *MinPos* method

MinPos is a novel approach fully presented in [Bau13]. It uses a strategy very similar to the one used in *Occupancy grids* (described in section 3.3.2.3), but introducing some improvements. As all the strategies that do not compute all the robot-task combinations (optimal calculus), the strategy in this approach attempts to distribute the robots over the unexplored locations as much as to be possible. The underlying hypothesis is that if this could be done along the time until the end of exploration, the exploration time would be smaller.

So, the main contribution of this approach consist in providing a better distribution of robots over the terrain minimizing the overall cost of tasks for a bigger set of scenarios. Particularly, instead of allocating the tasks just taking into account the cost, the robots are stacked under every single task -sorted by cost in ascending- and finally assigned to the one in which stack it has resulted best ranked (on top).

In figure 3.3 the task allocation result of applying *MinPos* method can be appreciated. There, exists four tasks $\{F_0, F_1, F_2, F_3\}$ and three robots $\{R_0, R_2, R_5\}$. It can be seen that despite R5 is nearest to the frontier points F_1, F_2 and F_3, it is assigned to F_0 attempting the ranking of R_5 related to each task.

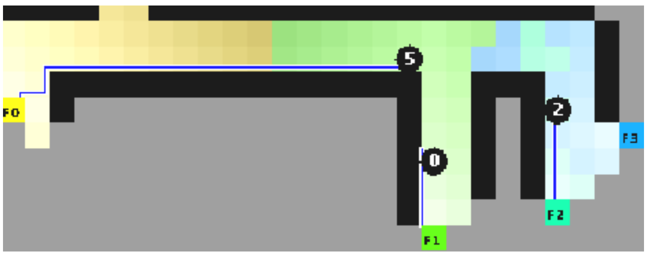

Figure 3.3: *MinPos* allocation method. [Bau13]

3.3.3 Objective function: Utility

In most of the surveyed jobs, the environment's representation used for computing the utilities is based on occupancy grids[4]. Basically, the occupancy grids allow distinguishing discreetly as well the known areas (the ones reached by the robot sensors) as the unknown ones (potential places to visit). Then, a classification can be defined and the cells can be labeled as *{busy, free, unknown}*. As the sensed information has a degree of uncertainty, the cells are often filled -progressively as the robots move in the environment- with probabilities representing the knowledge about the environment in that location. Thus, it is common to associate values less than 0.5 with the busy

[4]Also named as belief grids or belief maps.

class, values greater than 0.5 with the free class and values equal to 0.5 meaning that nothing can be said about those cells (unknown cells). Also, the utility functions often comprise a term where the potential information gain is expressed and another where the associated costs of navigation are computed.

3.3.3.1 Information gain and cost estimation

In this section there are two approaches that attempt to estimate the information gain and cost of the frontier points.

Information gain regions This is a partially distributed approach where the information gain estimation is computed by each robot considering all the frontier points, the observation radius and the number of unknown cells around each frontier point. After sending a bid with their preferences (based on euclidean distance), the robots still waiting until the *Executive Robot* starts to allocate the tasks. Attempting to avoid that many robots are moving to nearby locations, once a frontier point is allocated, the information gain of the neighbor frontier points is decreased proportionally with the euclidean distance between them. The path planning is performed using the iterative value iteration algorithm.[SAB+00]

Entropy Considering the centroids of the frontier regions, the entropy of the path between the robot's current position and the final destination is calculated. The iterative *value iteration* algorithm is used to compute the entropy of every path. Paths with very high entropy are filtered. Finally, the entropy is used to estimate the usefulness of the path while the cost is estimated through metric information. [VVDAC12]

Even when this approach has been applied in single robot systems, it is interesting because it provides a different idea about how to evaluate the paths -towards a set of possible destinations-during exploration.

3.3.3.2 Cost discounting

Other approaches do not attempt to estimate the information gain of frontier points assuming that in fact it is inestimable and, furthermore, the risk of estimation errors not compensates the effort. Instead, it is assumed that all the frontier points start with the same theoretic utility value to which the cost of the path is deducted.

Usually, two main strategies are used to implement this approach. One of them, as in [SAB+00], is based on the competition between the neighbor frontier points (taking into account a notion of visibility or relative euclidean distance). The other one, as in [WSB08], is based on the *Voronoi* diagram representation and the discounting is performed over the critical points of the diagram. This way, while the number of robots is less than the number of segments, there is no need to decrease the utility of the frontier points. Furthermore, this proposal allow to manage an heterogeneous team of robots. It is easy to include the feasibility of a task can be made by a particular robot into his cost function. When a robot can not reach a location because of his size, the cost function returns an infinite value. This way, the system can adapt to a bigger set of scenarios.

3.4 Communication methods

Into the context of cooperative exploration problem, communication is a central aspect in any proposed solution. As task allocation as cooperative mapping depend strongly on it.

3.4.1 Communication purposes

The communication between robots can take several forms. In fact, on the one hand, the robots can exchange mapping information (implicit cooperation) that is useful for have a global map updated but is also important to take decisions in a cooperative way. Additionally, if is necessary, they can also exchange useful information for coordination purposes (explicit coordination).

3.4.1.1 Implicit cooperation

There are several approaches where the cooperation is carried out in an implicit way. [VBSL11, VCBL09, Yam98]
In these cases, the cooperation between robots arises from a shared resource: the global map. Thus, each time a robot updates his map, it is sent to the other team members. Following this protocol, every member can always keep his global map updated -quite similar to the others- with similar landmarks, obstacles and targets. Therefore, if the same task allocation method is running in all members, they can cooperate without the need of explicit coordination.

3.4.1.2 Explicit coordination

Depending on the task allocation methods, sometimes explicit communication is required. This is the case of auctions (see 3.3.2.1) where the robots exchange bids with their preferences or the approach followed in [FFOV09] where the robots need to coordinate when they are close.

3.4.1.3 Information flow

Although this aspect is quite related to computing requirements, it has important implications on the communication ones. Depending on the quantity and the way the information flows among the robots will be the requirements to the network infrastructure. In general, the more centralized is the multi-robot system, the more important are reliability and band width.

Centralized approach Many approaches are based on the existence of an executive or distinguished agent that is responsible for collect all the information and send a final result back (updated global map or tasks assignment) to every team member. However, even when the role assignment can be done dynamically, the existence of a unique failure point is ever a risk. These cases are described in sections 3.3.2.2 and 3.3.2.3.

Partially distributed Other approaches (centralized or not) attempt to take advantage of the computing capacities present in every single robot. Thus, a part of the calculus is distributed over each of the team members. Typically, estimation of path cost and local maps are the tasks chosen to be made by themselves. These cases are described in sections 3.3.2.1 and 3.3.2.4

3.4.2 Communication constraints

In general, most of the jobs assume ideal communication[5]. However, this is usually an unrealistic assumption for several reasons. Some environment characteristics such as size or structure (the number of obstacles in indoor environments: buildings, it is usually large). The possibility of failure of communication systems should also be taken into account.

On the contrary, some proposals attack the problem without requiring the existence of ideal communication (more realistic). In that sense, the proposal presented in [BMSS05] considers both possibilities. One operating mode best suited to the existence of perfect communication and another if that hypothesis is not verified.

If there are restrictions on communication, the system is divided many times as robots subgroups have formed because of the lack of communication. Many experiments was conducted (varying the scope of communication devices) and one of the most interesting results was that the system does not show significant improvements if the scope of devices reaches 30% of the maximum distance measured in the environment.

[5] Without errors nor losses, with unlimited bandwidth and scope.

Chapter 4

Future work

4.1 Research lines

The following list represent (ordered in descending) the main topics over which could be interesting to keep working on. In some cases, they are quite related between each others.

- task assignment method
  - keep working on time optimization
  - bringing support for more realistic scenarios: not assuming ideal communications
  - dealing with group subdivision, open environments?
  - heterogeneous systems: size, capabilities (sensors, actuators)
- multi-objective function: distance, safety, level of battery
- explicit cooperation: messages and protocols to ask for help
- cooperative motion planning, active localization

4.2 Scheduling

- development of the *minpos* approach (until Oct/15)
  - select and deploy a *SLAM* algorithm
  - reuse (with some adaptations) a motion planning algorithm
  - develop an identification method
  - develop *minpos* allocation method
  - build a virtual scenario (environment + purpose)
  - simulation tests will be conducted
- keep working on this document (until Sep'15)
  - at least two new versions will be released
- 3th stay in Toulouse (Sep/Oct '15)

Bibliography

[Bau13] Antoine Bautin, *Stratégie d'exploration multirobot fondée sur le calcul de champs de potentiels*, Ph.D. thesis, Département de formation doctorale en informatique, École doctorale IAEM, 2013.

[BG08] Priyadarshi Bhattacharya and Marina L. Gavrilova, *Roadmap-Based Path Planning: Using the Voronoi Diagram for a Clearance-Based Shortest Path*, IEEE Robotics & Automation Magazine **15** (2008), no. 2, 58–66.

[BMSS05] W. Burgard, M. Moors, C. Stachniss, and F.E. Schneider, *Coordinated multi-robot exploration*, IEEE Transactions on Robotics **21** (2005), no. 3, 376–386.

[FFOV09] Antonio Franchi, Luigi Freda, Giuseppe Oriolo, and Marilena Vendittelli, *The sensor-based random graph method for cooperative robot exploration*, IEEE/ASME Transactions on Mechatronics **14** (2009), no. 2, 163–175.

[FGLM01] M. Foskey, M. Garber, M. C. Lin, and D. Manocha, *A Voronoi-based hybrid motion planner*, Proceedings of the 2001 International Conference on Intelligent Robots and Systems, vol. 1, 2001, pp. 55–60.

[Kuh55] H. W. Kuhn, *The Hungarian method for the assignment problem*, Naval Research Logistics **2** (1955), 83–97.

[MZL07] Xin Ma, Qin Zhang, and Yibin Li, *Genetic algorithm-based multi-robot cooperative exploration*, 2007 IEEE International Conference on Control and Automation, ICCA, vol. 00, 2007, pp. 1018–1023.

[RD05] Waldir L. Roque and Dionsio Doering, *Trajectory planning for lab robots based on global vision and Voronoi roadmaps*, Robotica **23** (2005), no. 4, 467–477.

[RN10] S. Russell and P. Norvig, *Artificial Intelligence: A modern approach*, Pearson, 2010.

[Roq96] Waldir L. Roque, *Introducao a Técnicas de Planejamento de Trajetórias de Robôs Móveis*, vol. 1, pp. 125–150, -, 1996.

[SAB+00] Reid Simmons, David Apfelbaum, Wolfram Burgard, Dieter Fox, Mark Moors, Sebastian Thrun, and Håkan Younes, *Coordination for Multi-Robot Exploration and Mapping*, Aaai, Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, 2000, pp. 852–858.

[Thr98] Sebastian Thrun, *Learning metric-topological maps for indoor mobile robot navigation*, Artificial Intelligence **99** (1998), no. 1, 21–71.

[VBSL11] Teresa Vidal, Cyrille Berger, Joan Sola, and Simon Lacroix, *Environment Modeling for Cooperative Aerial / Ground Robotic Systems*, STAR, 2011, pp. 681–696.

[VCBL09] Teresa a. Vidal-Calleja, Cyrille Berger, and Simon Lacroix, *Event-driven loop closure in multi-robot mapping*, 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems, Ieee, October 2009, pp. 1535–1540.

[VVDAC12] Rafael Valencia, Jaime Valls Miro, Gamini Dissanayake, and Juan Andrade-Cetto, *Active pose SLAM*, IEEE International Conference on Intelligent Robots and Systems, 2012, pp. 1885–1891.

[WSB08] K.M. Wurm, C. Stachniss, and W. Burgard, *Coordinated multi-robot exploration using a segmentation of the environment*, 2008 IEEE/RSJ International Conference on Intelligent Robots and Systems (2008), 1160–1165.

[WvH07] Ron Wein, Jur P. van den Berg, and Dan Halperin, *The visibility–Voronoi complex and its applications*, Computational Geometry: Theory and Applications **36** (2007), no. 1, 66–87.

[Yam98] Brian Yamauchi, *Frontier-based exploration using multiple robots*, Proceedings of the second international conference on Autonomous agents AGENTS 98 (1998), no. May, 47–53.