

Chapter 1

Introduction

robot

From translation of 1920 play “R.U.R.” (“Rossum’s Universal Robots”), by Karel Čapek (1890-1938), from Czech *robot* “forced labor, drudgery”, from *robotiti* “to work, drudge”, from Old Church Slavonic *rabota* “servitude”, from *rabu* “slave”, from a Slavic stem related to German *arbeit* “work”.

1.1 Overview and Motivation

Since the late 1960s, with the development of SRI’s Shakey robot [54], artificial intelligence (AI) and mobile robotics have been closely related. A mobile robot must be autonomous, deal with uncertainty, plan and decide what to do, react to unpredicted situations, that is, it has to overcome really hard problems if we want it to act in an intelligent and autonomous way. Thus, mobile robots pose one of the biggest challenges for AI.

Although impressive successes have been obtained since Shakey, it cannot still be said that the objective of having truly autonomous robots has been achieved. One of the fields in which there is still much to do is on mobile robotics for unknown environments.

Robotic systems for navigating through unknown environments are a focus of research in many application areas including, among others, spacecraft (rovers for Mars and the Moon) and search and rescue robotics. These systems have to perform very different tasks, from looking for rocks, picking them up and analyzing them, to assessing damages or looking for survivors after a natural disaster or accident has happened. However, they all share two key characteristics: first, they have to achieve their goals *autonomously*, and second, they have to *navigate* in unknown environments.

The first key point in these applications, *autonomy*, arises from the impossibility of always having a human operator controlling the robotic system. Although the ideal situation would be to have an expert operator controlling the robot, the necessary conditions cannot always be met. These conditions are usually related to the communication between the operator and the robot. In many situations it is very difficult to guarantee that the communication link will be robust, in terms of speed and availability. A clear

example is found on planetary exploration missions. A major problem in such missions is the distance between the robot and the control station (usually located on the Earth); the time of sending an order to the robot and having it executed can be in the order of minutes. In the case a fast reaction were needed (changing the trajectory of the robot, selecting a new scientific target that might be more relevant to the mission, etc.), this time would not be acceptable at all. Another disadvantage of depending on external agents (be it a human or any other device — e.g. a GPS device for localisation) is that the robot can get blocked if any of these agents providing basic information for accomplishing the task crashes. This would leave the robot with no means to continue with its mission. Therefore, all the decisions should be taken on-board, without having to exchange commands or information with any external agent, or at least, make this exchange minimal, such as sending only information about task initialisation (e.g. target selection, task description, etc.).

The other important point for such applications is *navigation*. The robot must be able to start in an unknown location and navigate to a desired target. Navigation in unknown unstructured environments is still a difficult open problem in the field of robotics. The first difficulty of such an environment is that there is no a priori knowledge about it, and therefore a map can only be built while exploring. Secondly, unstructured environments are characterized, precisely, by the lack of structure among the different objects of the world. This is usually the case for outdoor environments. On the other hand, in structured environments, such as offices, buildings, etc. many suppositions about their structure can be done. For instance, walls and corridors are straight, they are usually orthogonal, most of the doors have the same size, etc. These characteristics are very helpful when building a map of the environment, as its structure can be inferred without the need of sensing the whole environment. Contrarily, in unstructured environments such suppositions do not hold, so the robot can only rely on the information it is able to gather from its sensors. This makes the task of map building and navigating even more difficult.

This research work is part of a larger robotics project. Another partner (IRI¹) in the project is building a six legged robot with on board cameras for outdoor landmark recognition. The goal of the project is to have a completely autonomous robot capable of navigating in outdoor unknown environments. A human operator will select a target using the visual information received from the robot's camera, and the robot will have to reach it without any intervention of the operator. The robot could also have an image or description of the target, so the human intervention would not even be needed for selecting the target. A first milestone for achieving the final goal of the project is to develop a navigation system for indoor unknown unstructured environments for a wheeled robot. Moreover, the environments of this first stage are composed of easily recognizable landmarks, since the vision system for outdoors is not yet available.

We propose a robot architecture to accomplish this first milestone. The approach used and the results obtained are the subject of this thesis. The robot architecture is composed of three systems: the *Pilot* system, the *Vision* system and the *Navigation* system. Each system competes for the two available resources: motion control (direction of movement) and camera control (direction of gaze). The three systems have the

¹Institut de Robòtica i Informàtica Industrial, <http://www.iri.csic.es>

following responsibilities. The Pilot is responsible for all motions of the robot. It selects these motions in order to carry out commands from the Navigation system and, independently, to avoid obstacles. The Vision system is responsible for identifying and tracking landmarks (including the target landmark). Finally, the Navigation system is responsible for choosing higher-level decisions in order to move the robot to a specified target. This requires requesting the Vision system to identify and track landmarks (in order to build a map of the environment) and requesting the Pilot to move the robot in various directions in order to reach the goal position or some intermediate target.

From the brief description of the robot architecture given above, it can be observed that the three systems must *cooperate* and *compete*. They must cooperate because they need one another in order to achieve the overall task of reaching the target position. But at the same time they are competing for motion and camera control.

The Navigation system is implemented as a multiagent system, where each agent is competent in a specific task. Depending on its responsibilities and the information received from other agents, each agent proposes which action the Navigation system should take. Again, we find that the agents must cooperate, since an isolated agent is not capable of moving the robot to the target, but they also compete, because different agents want to perform different actions.

The problem of cooperation and competition between different agents is very common in robotics, and *Behavior-based Robotics* [3] addresses exactly this issue. This approach to robotic systems deals with coordinating, or arbitrating, different behaviors sending requests for actions, usually incompatible with each other, to a robot. The role of the coordination is to select a single action command that will be sent to the robot. When this action is a selection of one of the behaviors' requests, we talk about *competitive* coordination, whereas if the action is a mix of several behaviors' requests, we talk about *cooperative* coordination. In our architecture, each agent plays the role of a behavior, and there is an additional agent playing the role of coordinator.

For both the overall robot system and the Navigation system, we propose the use of a new competitive coordination system based on a *bidding mechanism*. In the overall robot system, the Navigation and the Pilot systems generate bids for the services offered by the Pilot and Vision systems. These services are to move the robot toward a given direction, and to move the camera and identify the landmarks found on its view-field, respectively. The service actually executed by each system depends on the winning bid at each point in time. Similarly, in the Navigation system, each agent bids for the action it wants the robot to perform. These bids are sent to a special agent that gathers all bids and determines the winning action. The selected action is then sent as the Navigation system's bid for the services of the Vision and Pilot systems.

The architecture just described above is actually an instantiation of a general coordination architecture we have developed. In this architecture there are two types of systems: *executive systems* and *deliberative systems*. Executive systems have access to the sensors and actuators of the robot. These systems offer services for using the actuators to the rest of the systems (either executive or deliberative) and also provide information gathered from the sensors. On the other hand, deliberative systems take higher-level decisions and require the services offered by the executive systems in order to carry out the task assigned to the robot. Although we differentiate between these

two types of systems, the architecture is not hierarchical, and coordination is made at a single level involving all the systems. This coordination is based on a simple mechanism: *bidding*. Deliberative systems always bid for the services offered by executive systems, since this is the only way to have their decisions executed. Executive systems that only offer services do not bid. However, those executive systems that require services from any executive system (including themselves) must also bid for them. The systems bid according to the internal expected utility associated to the provisioning of a service. A coordinator receives these biddings and decides which service each of the executive systems has to engage on. In the instantiation for our navigation problem, there is a deliberative system, the Navigation system, one executive system that bids, the Pilot system, and one executive system that only offers services, the Vision system.

To navigate in an unknown environment, the robot must build a map. Existing approaches assume that an appropriately detailed and accurate metric map can be obtained through sensing the environment. However, most of these approaches rely on odometry sensors, which can be very imprecise, due to the wheels or legs slipping, and lead to many errors that grow as the robot moves.

Our approach considers using only visual information. The advantage of using visual information is that it is independent of any past action the robot may have performed, which is not the case for odometry. The robot must be equipped with a robust vision system capable of recognising landmarks, and use them for mapping and navigation tasks. The specific scenario that we are studying assumes that there is a target landmark that the robot is able to recognize visually. The target is visible from the robot's initial location (so that the human operator can select it), but it may subsequently be occluded by intervening objects. The challenge for the robot is to acquire enough information about the environment (locations of other landmarks and obstacles) so that it can move along a path from the starting location to the target.

But even vision-based navigation approaches assume unrealistically accurate distance and direction information between the robot and the landmarks. We propose a fuzzy set based approach that assumes only very rough vision estimation of the distances and, therefore, does not rely on any localisation device.

The main goal of our research is to design a robust vision-based navigation system for unknown unstructured environments. In particular, we want to provide the robot with orientation sense, similar to that found in humans or animals. The rationale of the orientation sense is that the robot does not need to know the exact route from its starting point to the target's location, but it uses landmarks as references to find its way to the target. To make a parallel with humans, when giving directions for going somewhere in our city, no one gives exact distances, turning angles, etc., but gives approximate distances, and more important, reference points (distinctive places such as buildings, squares, etc.) that help us getting to the destination. In our approach, this orientation sense is realized by the use of *landmark-based navigation*, *topological mapping* and *qualitative computation* of landmark locations. We give a brief definition of each of these three concepts:

- *Landmark-based navigation*: A landmark is a visually salient object of the environment the robot is able to identify. Other navigation approaches that do not use vision systems define a landmark as a set of features the robot can detect with its

sensors (usually sonar or laser readings). As the robot explores the environment, it stores the detected landmarks on a map. When the robot needs to locate itself on the map, it can do it by matching the detected landmarks with the landmarks already stored on the map. This approach avoids requiring odometry as the main information source for navigating and building maps.

- *Topological mapping*: this approach to map building has a close relationship with landmark-based navigation. A topological map is usually a graph, where nodes represent landmarks and arcs represent paths or motion instructions for going from one landmark to another. The advantage of this approach is that there is no need for building accurate geometric maps. Storing the topological relationships among the landmarks in the environment is enough.
- *Qualitative computation*: we use the term “qualitative” in the sense that we do not need to work with exact distance or angle information; we can deal with some imprecision about it. More specifically, we deal with it by means of fuzzy numbers and fuzzy arithmetic. Thus, when we talk about qualitatively computing the location of a landmark, it means that we are taking into account the possible imprecision about its location.

Our map representation, however, is slightly different from the one given above. The map is a labeled graph whose nodes, instead of representing isolated landmarks, represent triangular shaped regions delimited by groups of three non-collinear landmarks, and whose arcs represent the adjacency between regions, that is, if two regions share two landmarks, the corresponding nodes are connected by an arc. The arcs of the graph are labeled with costs that reflect the easiness of the path between the two corresponding regions. A blocked path would have an infinite cost, whereas a flat, hard paved path would have a cost close to zero. Since the map is not given, but built while the robot moves, these costs can only be assigned after the robot has moved (or tried to move) along the path connecting the two regions. Although the adjacency of nodes in our graph also represents adjacency of topological places, the arcs contain only cost information, not instructions on how to get from one place to another. But, actually, this information is not missing, it is implicit in the adjacency of regions. Given that two nodes are adjacent only if their regions share two landmarks, it is clear that to go from one region to another the robot has to cross the edge formed by the two common landmarks, unless there is a long obstacle blocking this path.

Although this topological map would be sufficient for carrying out navigational tasks, we also provide the robot with the capability of storing the spatial relationships among landmarks. To realize this capability, we have extended Prescott’s beta-coefficients system [55]. Prescott’s model stores the relationships among the landmarks in the environment. The location of a landmark is encoded based on the relative locations (headings and distances) of three other landmarks. This relationship is unique and invariant to viewpoint. Once this relationship has been stored, the location of each landmark can be computed from the locations of the three landmarks encoding it, no matter where the robot is located, as long as the robot can compute the heading and distance to each of the three landmarks. As the robot explores the environment, it stores the relationships among the landmarks it sees. This creates a network of relationships among

the landmarks in the environment. If this network is sufficiently-richly connected, then it provides a computational map of the environment. Given the headings and distances to a subset of currently-visible landmarks, the network allows us to compute the locations of all of the remaining landmarks, even if they are currently not visible from the robot. Prescott assumed that the robot could have the exact distances and headings to the landmarks, but as we mentioned previously, we need to deal with the imprecision of the real world. To deal with it, we have extended the model using fuzzy numbers and fuzzy arithmetic. With this extension, if the target is ever lost during the navigation, the robot will compute its location with respect to a set of previously seen landmarks whose spatial relation with the target is qualitatively computed, both in terms of fuzzy distance and direction.

We have implemented the overall architecture and the Navigation system and first tested it over a simulator. After obtaining promising results on simulation, we have implemented the algorithm on a wheeled robot and tested it on real environments.

Although the tuning of our system was carried out through the experimentation with the real robot, we also employed simulation to apply machine learning techniques in order to improve the performance of the system. In particular, we have applied Reinforcement Learning and Genetic Algorithms. We have used Reinforcement Learning to have the system learn to use the camera only when appropriate. The camera is a very expensive resource, and it has also a very high demand (the Pilot and several agents compete for its control). Since manual tuning of the parameters controlling the agents' behaviors is very difficult, and the problem we are trying to solve is a quantitative trade-off, Reinforcement Learning is found to be the most appropriate technique to use, as its main goal is to maximize expected reward. We have obtained good results that show the feasibility of applying Reinforcement Learning to improve our system. Nonetheless, we still need further experimentation and tuning of the learning algorithm, in order to integrate the learned policy into the multiagent system. In parallel, we have used a Genetic Algorithm to tune the different parameters of the agents. The tuning of these parameters cannot be done manually, neither can it be done using Reinforcement Learning. Therefore, we have chosen to use a genetic algorithm approach.

1.2 Contributions

The objective of the research carried out during the completion of this PhD thesis has been to accomplish the first milestone of the above mentioned project, that is, developing a navigation system for indoor unknown unstructured environments for a wheeled robot. More precisely, we have focused on the Navigation system and on the overall robot architecture. However, we have also had to implement simple versions of the Pilot and Vision systems in order to realize and test the Navigation system.

As it may have already been noticed, this work has two main research threads: the *control architecture* and the *mapping and navigation method*.

Regarding the *control architecture*, we have proposed a general coordination architecture that uses a bidding mechanism for coordinating a group of systems (and agents) that control a robot. This mechanism can be used at different levels of the control architecture. In our case, we have used it to coordinate two of the systems of the robot

(Navigation and Pilot systems) and also to coordinate the agents that compose the Navigation system itself. Moreover, the multiagent view of the Navigation system could also be applied to other systems, having a multiagent Pilot or a multiagent Vision system. Using this bidding mechanism, the action actually being executed by the robot is the most urgent one at each point in time, and thus, if the agents bid rationally, the dynamics of the biddings would lead the robot to execute the necessary actions in order to reach a given target. An advantage of using such mechanism is that there is no need to create a hierarchy, such as in the subsumption architecture, but it is dynamically changing depending on the specific situation of the robot and the characteristics of the environment. A second advantage is that its modular view conforms an extensible architecture. To extend this architecture with a new capability we would just have to plug in a new system (or agent).

Regarding the *mapping and navigation method*, we have extended the work presented by Prescott [55], so that it can be used with fuzzy information about the locations of landmarks in the environment. This is of great importance when working with real robots, as it is impossible to avoid dealing with the imprecision of real world environments. Together with this extension, we have also developed methods that permit computing diverting targets, needed by the robot when there is no clear path to the goal.

1.3 Publications

The publications related with this research that have been published as journal articles or in conference proceedings are the following:

- C. Sierra, R. López de Mántaras and D. Busquets. Multiagent bidding mechanisms for robot qualitative navigation. *Lecture Notes in Computer Science (Proceedings ATAL'00)*, vol. 1986, pages 198–212, Springer, Verlag, 2001.
- D. Busquets, R. López de Mántaras and C. Sierra. A robust MAS coordination mechanism for action selection. *Proceedings of 2001 AAAI Spring Symposium, Stanford, CA. Robust Autonomy Serie*, pages 38–40, 2001.
- D. Busquets, R. López de Mántaras, C. Sierra and T.G. Dietterich. Reinforcement Learning for Landmark-based Robot Navigation. *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2002)*, pages 841–842. ACM press, 2002.
- T.G. Dietterich, D. Busquets, R. López de Mántaras and C. Sierra. Action Refinement in Reinforcement Learning by Probability Smoothing. *Proceedings of the 19th International Conference on Machine Learning (ICML'02)*, pages 107–114, 2002.
- D. Busquets, T.G. Dietterich, R. López de Mántaras and C. Sierra. A multi-agent architecture integrating learning and fuzzy techniques for landmark-based robot navigation. *Topics in Artificial Intelligence. Lecture Notes in Artificial Intelligence (Proceedings of CCIA'02)*, vol. 2504, pages 269–281, Springer, Verlag, 2002.

- D. Busquets, C. Sierra and R. López de Mántaras. A multi-agent approach to fuzzy landmark-based navigation. *Journal of Multiple-Valued Logic and Soft Computing*, Old City Publishing (In press).
- D. Busquets, C. Sierra and R. López de Mántaras. A Multi-agent approach to qualitative landmark-based navigation. *Autonomous Robots*, Kluwer Academic Publishers (In press).

1.4 Structure of the Thesis

This PhD thesis report is organized as follows:

Chapter 1. Introduction

This chapter gives an overview of this PhD thesis, its motivations, objectives and its main contributions. It also gives a list of the publications originated from the research carried out during the completion of the thesis.

Chapter 2. Related work and state-of-the-art

This chapter is devoted to relevant related work and state-of-the-art on the field of autonomous robots for unknown unstructured environments. The relevant work has been divided in two parts, one for each main thread of research of the thesis: control architectures, and mapping and navigation methods. The relevant work concerning control architectures gives an overview of the different approaches on autonomous robots control, focusing on the behavior-based approach. Regarding the mapping methods, we review and compare the two main approaches for map building, the metric one and the topological one. A comparison between two different localisation approaches (landmark-based localisation and model matching) is also given.

Chapter 3. Mapping and Navigation

In this chapter we firstly describe Prescott's model for storing spatial relationships among the landmarks of the environment. After that, we describe how we have extended this model for dealing with imprecise information about the location of the landmarks. We also present the algorithm for building a topological map of the environment and how it is used to compute diverting targets, needed by the robot when it is blocked.

Chapter 4. The Robot Architecture

In this chapter a general coordination architecture based on a bidding mechanism is presented. We also present the particular instantiation of the general architecture that we have used to solve the navigation problem. A detailed description of the multiagent Navigation system is also given in this chapter.

Chapter 5. Simulation Results

In this chapter the results of the simulated experiments are presented. These experiments include the testing of our architecture and the application of Machine Learning techniques in order to improve the performance of the system. In particular, we present the application of Reinforcement Learning, which we have used to make the system learn how to appropriately use the camera, and an application of Genetic Algorithms, used to tune some of the parameters of the agents of the Navigation system.

Chapter 6. Real Experiments

This chapter is devoted to present the results of the experiments on real environments with a real robot. Firstly, the wheeled robot platform and a simple vision system used for the real environments experiments are described. Then, we describe the different scenarios in which the experiments have been carried out. Finally, the results of the experimentation in such scenarios are given and discussed.

Chapter 7. Conclusions and Future Work

In this chapter, we summarize the main contributions of the thesis, and point out some open problems and future research perspectives that we plan to tackle in the near future.