Cambridge University Press
0521573319 - The Map-Building and Exploration Strategies of a Simple Sonar-Equipped Robot: An Experimental, Quantitative Evaluation
David Lee
Excerpt
More information

Chapter 1

Question, Context and Method

1.1 The Question - What is This Thesis About?

This thesis examines the process by which an autonomous mobile robot constructs a map of its operating environment. This process can be considered as two distinct topics. First, the robot has to interpret the findings of its sensors so as to make accurate deductions about the state of its environment. This is the problem of 'map-building'. Second, it has to select its viewpoints so that the sensory measurements contain new and useful information. This is the problem of 'exploration'. This thesis describes a practical and experimental investigation into both of these issues.

This document is structured as a large number of short chapters. This reflects the wide range of subjects which had to be examined in order to build an effective working robot for map-building and exploration experiments. For ease of reading, the chapters are grouped into three parts; Part I (Chapters 2 to 4) examines the principal areas of previous research upon which this thesis is built; Part II (Chapters 5 to 10) describes the components of the map-building system; and finally Part III (Chapters 11 to 20) reports on experiments to evaluate the effectiveness of a range of exploration strategies. The closing chapters of Part III summarise the results and conclusions and suggest directions for further research.

The remaining sections of this introductory chapter serve as an overview of the thesis and put the later chapters into context.

Section 1.2 begins with a brief review of the history of mobile robots and then describes some of the issues which are currently occupying researchers. Section 1.3 outlines the hardware and software which make possible the experiments described in this thesis and also explains some of the key implementation decisions. Section 1.4 summarises the contributions of this research.

1.2 The Context - Why Make Maps?

It can be difficult to devise watertight definitions of research topics. Take, for example, 'robotics'. Many researchers have suggested definitions, usually agreeing about the core of the subject but disagreeing about the inclusion of topics such as teleoperation and prostheses.

Cambridge University Press
0521573319 - The Map-Building and Exploration Strategies of a Simple Sonar-Equipped Robot: An Experimental, Quantitative Evaluation
David Lee
Excerpt
More information

This thesis makes no attempt to provide a definition, but it may prove worthwhile to examine some views of the topic and to see what they have in common.

The fascination of robotics lies in its attempt to create machines which have something in common with human beings. The layman's idea of a robot is dominated by the fictional examples seen in films and on television, ranging from the first movie robot in 'Metropolis' in 1926 through 'C3PO' in 'Star Wars' to Data in 'Star Trek'. These machines not only have human skills such as language and reasoning but they also look like people. With this expectation it can be disappointing to visit a robot lab and see mobile robots that look more like dustbins on wheels.

If humanoid appearance is not important, then what is? Consider a few attempts at definition:

> A robot is a programmable, multi-function manipulator designed to move material, parts, or specialized devices through variable programmed motions for the performance of a variety of tasks. (Schlussel 1983)

> A robot is a machine which can be programmed to do a variety of tasks, in the same way that a computer is an electronic circuit which can be programmed to do a variety of tasks. (McKerrow 1991, page 8)

> Robotics is the intelligent connection of perception to action. (Brady 1985)

These definitions raise questions of course (Are mobile robots 'manipulators'? What does 'intelligent' mean? ...) but they share the requirement that a robot must be able to perform a variety of tasks. If a machine blindly repeats the same set of actions, with no possibility of variation, it does not qualify as a robot.

The first industrial robot began operation in the early 1960's. Since then robots have gained wide acceptance in the manufacturing industry, specifically in the manufacture of vehicles and electric machine tools. By 1988 the world population of industrial robots had grown to 280,000 (Kennedy 1993, page 88).

Early industrial robots operated in environments which were specifically designed around the robot. Each component was supplied to the robot in a predefined position and orientation so that the robot knew exactly where to find it. The robot could indeed be programmed to perform different tasks, but the changeover could be a costly and time-consuming process. As the range of potential applications has expanded over the last 30 years, there has been increasing interest in robots which are able to identify variations in their environment and to react to them without human intervention. It would, for example, be useful for an assembly robot to be able to pick up components from a conveyor belt, however they may be positioned. This interest in tolerance of variation has also been fuelled by the trend towards shorter production runs. A manufacturer may need to produce several products and may not be willing to incur the costs of frequent reprogramming and recalibration of a robot. For example, it was recently reported (Hallahan 1994) that these pressures led IBM to scrap the robots in an automated factory and to replace them with human workers. The result was an increase in productivity.

Environmental variation increases rapidly when robots become mobile. The appropriate action for the robot depends upon where it finds itself, it may be uncertain about its exact location, and it may have to share its environment with unpredictable human beings. In recent years, a growing amount of research effort has been invested in the problems peculiar to mobile robots. A rough indication of the rate of growth of this research effort can be obtained by counting the number of published papers about[1] 'robots' and 'mobile robots' in the last 10 years, according to the Bath Information and Data Services (BIDS). In 1983 there were 313 publications about robots, of which 5 were about mobile robots. By 1993 the total number of robot publications had grown to 775 (a factor of 2.5) of which 99 were about mobile robots (a growth factor of 19.8). By this measure 12.8% of robot research in 1993 was concerned with mobile robots.

Mobile robot research is generally taken to have started in the 1960's, although there were occasional earlier examples (e.g. Shannon's maze-runner in 1940 and Grey Walter's 'turtle' robot in 1953). The first mobile robot to use vision was Shakey, built in 1969 at the Stanford Research Institute. Its objective was to use its cameras to recognise objects, approach them, and perform an action such as pushing them over (McKerrow 1991, pages 4–6).

A successor to Shakey was the Stanford Cart (Moravec 1983). This again used vision, and built a world model which it used to plan paths whilst avoiding obstacles. Despite some successes, it was found to be unreliable (being confused by changes in the quality of daylight at different times of day) and was extremely slow.

Until the mid 1980's it was taken as axiomatic that a mobile robot should use its sensors to build a world model and then use the world model to plan its actions. Then several researchers, frustrated by the limited achievements of these robots, began to question the need for a world model. What types of behaviour could arise from a robot which simply reacted to its sensory inputs? 'Behaviour-based' robotics was born.

Early experiments with behaviour-based robots had impressive results. It was shown, by Brooks (1986) and others, that behaviour similar to that observed in simple animals, such as insects, could be produced by robots with very little, if any, internal state. Robots could avoid obstacles, approach targets, and follow walls by reacting rapidly to the input from their senses. An architecture, the subsumption architecture, was designed to make it easier to build these robots so that the most appropriate behaviour would be used at each moment.

Advocates of behaviour-based robots invoked evolutionary theory to support their cause. It clearly took much longer for nature to evolve the 'basic' skills, such as walking and avoiding threats, than the 'higher' skills, such as language and reasoning. Therefore, the argument runs, it makes sense to focus first on the acquisition of the simpler skills. Once these are mastered, cognition will be much easier. The use of evolution to support the cause is very much in keeping with a recent 'back to nature' trend. Supporters of artificial neural networks continue to cite neuroscience as their inspiration and justification; genetic algorithms evolve problem solutions in a Darwinian way; the new field of 'Artificial Life' is concerned with the study of artificial systems which exhibit lifelike behaviours.

The debate between the supporters of 'behaviour-based' robotics and the proponents of world models has been heated, giving the impression that the way forward would be *either*

---

[1] The selection was based simply upon the inclusion of the words 'robot' or 'mobile robot' in the paper's title.

behaviour-based *or* model-based. Although these extreme positions have helped to clarify the issues, opinion now appears to be settling in the middle ground. For example, at a recent workshop provocatively entitled 'Models or Behaviours - Which Way Forward For Robotics?', 12 of the 17 speakers favoured a hybrid approach. (AISB 1994)

World models are only useful if they continue to match the true state of the world. A model is then used to *predict* the state of the environment so that effective plans can be made. The value of a world model is therefore directly linked to the degree of predictability of the robot's environment. If the environment is completely under the control of the robot (an automated warehouse, for example), then a world model would be very useful. If, on the other hand, the robot has very little control over its environment (negotiating a busy high street, for example), then a world model would be much less useful than a quick set of reflexes. Most applications lie somewhere between these two extremes, suggesting the wisdom of a hybrid architecture in which the predictable features of the world are incorporated into a world model and the world model is used to *guide* the behaviour-based components.

The debate about the need for a world model has spawned discussions about the type of world model that is appropriate. In particular, a number of behaviour-based projects have decided to reject detailed metric maps in favour of distributed, topological maps. Chapter 2 reviews the different types of map which have been used by mobile robots and argues that the selection of a type of map depends strongly on the intended application of the robot. The different types of map are presented in a hierarchy, ordered by the 'strength' of the map. 'Strength', in this context, refers to the range of geometric properties which can be derived from the map. The categories are:

Recognisable Locations The map consists of a list of locations which can be reliably recognised by the robot. *No geometric relationships can be recovered.*

Topological Map In addition to the recognisable locations, the map records which locations are connected by traversable paths. *Connectivity between visited locations can be recovered.*

Metric Topological Maps This term is used for maps in which distance and angle information is added to the path descriptions. *Metric information can be recovered about paths which have been travelled.*

Full Metric Maps Object locations are specified in a fixed co-ordinate system. *Metric information can be recovered about any objects in the map*

The preceding discussion argues that the decision whether to use a world model and, if so, what type to use depends strongly on the intended application of the robot. One of the first steps in the research was therefore to choose an application. The selected application was *indoor delivery.* Examples of such an application could be office mail delivery, an intelligent wheelchair for disabled people, or even a domestic robot. Chapter 3 describes the properties of such an application in detail and argues that a full metric map would be needed.

The robot could get a metric map in two ways; a human operator could give it the map or the robot could build its own. Advantages of the latter solution include:

Changes in the Robot's Environment Over time the robot's environment will change. It could periodically re-map the world.

Matching the Map to the Sensors It is difficult for a user to predict which features of the world will be easily recognisable by the robot's sensors. A user-supplied map may therefore be of limited value to the robot.

Ease of Use It would be a more attractive commercial proposition for a purchaser of a robot to be able to put it to work without having to measure its new environment or otherwise obtain a map.

Level of Detail For some purposes the level of detail required might be higher than that obtained from a readily-available architectural drawing.

In the light of these advantages, it was decided to pursue the goal of autonomous map construction.

In all but the simplest environments, some objects will be hidden by other, nearer, objects. The robot will then have to move to gather knowledge about its entire work area. One is therefore left with the question of *how* it should move.

Exploration strategies have not been extensively examined in the literature. It is, however, possible to identify a number of categories into which the existing work falls:

Human Control Many researchers report the results of map construction while the robot was under the control of a human operator.

Reactive Control As mentioned earlier, supporters of behaviour-based robotics have, in situations in which some type of world model *is* needed, favoured topological maps. To obtain such a map, the robot will typically navigate under the control of a reactive algorithm, such as wall-following, which is well-suited to the behaviour-based architectures.

Approaching the Unknown A reasonable exploration strategy is for the robot to approach those regions of its environment about which it knows least.

Optimal Search Strategies Under the heading of 'terrain acquisition', researchers have provided mathematical analyses of strategies which are guaranteed to find all objects in the robot's environment. Emphasis is placed on minimising the length of the path travelled by the robot during exploration. This work typically makes simplifying assumptions about ideal sensors.

Chapter 4 describes previous exploration research in more detail and discusses a number of examples.

One of the objectives of this thesis is to provide quantitative comparisons between a representative sample of exploration strategies. In the light of the discussion of behaviour-based robotics, a key question is:

> How much should the robot use its developing map to guide the exploration? When, if at all, does the robot's map contain enough information to justify using it to guide further exploration, instead of using a reactive, representation-free strategy?

1.3 The Method - How Will the Question be Addressed?

Examination of mobile robot research shows two distinct approaches; simulation and implementation. Some researchers build computer models of the performance of a robot and then use the model to test theories and algorithms. Others choose to build a real-world robot. It was necessary to choose between these two approaches. Simulation has advantages. One can see the results of an algorithm much more quickly by applying it to a computer simulation than to the real robot. The researcher is able to test new ideas without the constraints of time and expense associated with using a real robot. In addition, simulations allow the researcher to focus more tightly on the precise aspect of the problem in which he or she is interested. If, for example, the research is centred on path planning there may be little value in worrying about the mechanical engineering problems of building a physical robot. But the advantages are outweighed by disadvantages. To build a computer model, the researcher has to abstract the essential features of the system being modelled. This abstraction necessarily involves some degree of simplification. In mobile robotics this is most often noticeable in the modelling of sensors. For example, analysis of terrain acquisition problems (Lumelsky, Mukhopadhyay, & Sun 1991) assumes a sensor which can reliably detect the boundaries of any object that falls within a given radius of the robot. This is a highly idealised model of a sensor. Research based on such simplifications may well produce useful results, but there is always the danger that, in the simplification process, one has ignored a vital property of the robot so that the results will not be valid when tested on a real robot. Another important disadvantage has been summarised as 'simulations are doomed to succeed' (Miller *et al.* 1989). Since the same person is modelling both the problem and its solution, it is very tempting to include into the model of the problem just those features which can be handled by the solution. This is not to suggest any dishonesty on the part of the researcher. It may simply be that the model and the solution are built upon the same set of assumptions. The work described in this thesis tests exploration strategies on a real robot. However, in order to keep the time savings of a simulation, a Trace/Replay mechanism was implemented. This meant that all the sensor and movement information which was generated during an exploration by the real robot could be stored and subsequently replayed at will. This was found to be extremely useful throughout the research. Whenever a new idea was being implemented, a large amount of authentic information was available for testing.

Recent research, especially that motivated by behaviour-based robotics, has emphasised the creation of completely stand-alone robots, shunning the use of 'umbilical cords' to connect

Cambridge University Press
0521573319 - The Map-Building and Exploration Strategies of a Simple Sonar-Equipped Robot: An Experimental, Quantitative Evaluation
David Lee
Excerpt
More information

the robot to a stationary computer. The approach adopted in this research was, however, just that. A small mobile robot was constructed with a serial cable link to a Sun workstation on which the map construction and exploration control were performed. The following points are given in support of this decision:

- A real-world robot was implemented to tackle two real-world problems which are often oversimplified in simulations: sensing and localisation. Neither of these problems is diminished by the presence of an umbilical cord.
- Given the objective of testing exploration strategies, the smaller the robot the better. If the robot were enlarged to be able to carry the equivalent power of the Sun workstation on-board, it would not only be much more expensive but would also need a larger area of lab space in which to create varied environments for exploration.
- In parallel with the communications cable, it was practical to supply constant power to the robot. This made it possible to have lengthy experimental sessions without having to worry about keeping batteries charged.
- The graphics workstation was ideal for the display of the generated maps and for the examination of the exploration paths selected by the robot. If the robot had been stand-alone, one would either have had to follow it around to examine the map or implement a periodic download to a static display.
- The separation of hardware mirrored a corresponding separation of function. The workstation operated with high-level commands such as 'Move Forward 1000 mm' or 'Turn Left 90 Degrees', leaving the robot to concern itself with the low-level details such as motor control and obstacle avoidance. The robot itself had no world model. With this separation it was straightforward to test the same high-level software with a different robot, and vice versa.

The robot, ARNE[2], is described in detail in Chapter 5.

A simple dialogue was defined for communication with ARNE. This could then be used for direct control, through a terminal for example, or for control by the exploration and map construction software. The dialogue is described in Appendix D.

A key choice was the type of range sensor to be used on ARNE. The most commonly used range sensors in mobile robotics are vision, laser rangefinders, and ultrasonic time-of-flight sensors (sonar). The use of each of these sensors is an active research topic. Ultrasound was chosen for ARNE partly because of its low cost (many delivery applications are likely to be at the cheaper end of the market) and also because of recent work which suggested that ultrasonic sensing had been under-rated. Many researchers, frustrated by problems of wide beam width and unwanted reflections, have decided that ultrasonic sensing is only suitable for short-range obstacle avoidance. However, recent work (Zelinsky 1991b; Leonard & Durrant-Whyte 1992; Curran & Kyriakopoulos 1993) has suggested that sonar's bad reputation may not be justified and that reliable range readings can be obtained from sonar if a realistic

---

[2] Autonomous Robot for Navigation and Exploration. Also named for Arne Saknussemm, Jules Verne's explorer who was first to reach the centre of the earth. (Pronounced 'Arnie').

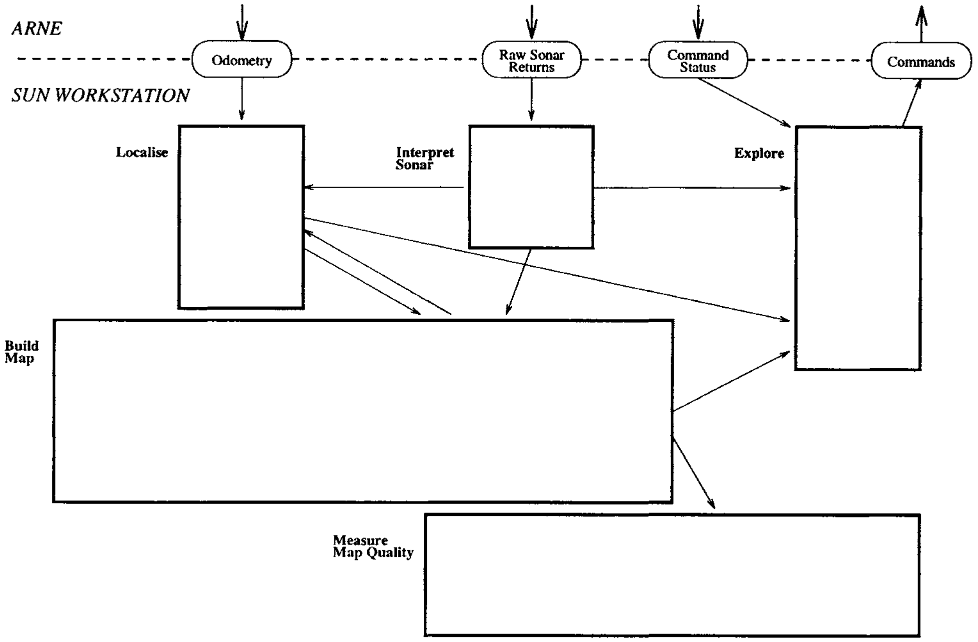

Figure 1.1: The Main Modules of the Workstation Software

The boxes represent the main software modules which were implemented on the workstation. The arrows indicate data flow between the modules and the robot. The uneven size and spacing of the boxes is to allow detail to be added as each module is explained. Path planning does not appear as a separate task on this diagram because, as will be seen later, it plays a part both in exploration and in the map quality calculations.

model of the sensor's behaviour is used. ARNE was therefore equipped with a Polaroid ultrasonic rangefinder.

A suite of software modules was designed and implemented on the Sun workstation to perform the tasks listed below. The block diagram in Figure 1.1 summarises the communication links among these tasks and between them and ARNE.

Interpret Sonar The range values generated by ARNE take the deceptively simple form of an angle (relative to the robot's orientation) and a distance. Unfortunately the sonar beam is wide and there is no guarantee that the object that caused the echo was in the centre of the beam. Preprocessing of the sonar returns can reduce some of the uncertainty. Chapter 6 describes some experiments to test the performance of ARNE's sonar sensor and uses the results of the experiments to model the sensor's behaviour.

Build Map This is the core of the system software. Position and range information are merged to generate a representation of the world which can be used to plan ARNE's actions. The representation is formed in three layers; first the range information is analysed to suggest features which could have caused the given readings; these hypothetical features are then gathered into mutually-supportive clusters to create 'confirmed' features; finally the confirmed features are used to construct a grid-based free-space map. The entire map-building process is described in detail in Chapter 7.

Plan Paths The free-space map can be used for path planning. The problem is to generate a sequence of movement commands which will, according to the map, move ARNE efficiently from a known starting position to a specified target position without colliding with any obstacles. A path planner is implemented using the technique of distance transforms. This is described in Chapter 8. The path planner is used in two places in this thesis. As one would expect, it is used to plan exploratory movements. It is also used when measuring the quality of the maps produced by these movements.

Localise To translate sensor readings into information about the world, it is essential to know where ARNE was when the readings were taken. This information comes from two sources. The first, and simpler, source is the odometry information returned by ARNE which converts measurements of the amount of wheel rotation into an estimate of the distance moved or the angle turned. Odometry is notoriously unreliable because of uneven floors or wheel slippage. It is therefore necessary to augment the odometry by measuring ARNE's position relative to known objects in the environment. Chapter 9 gives details of the chosen localisation method. It uses a Kalman filter to determine the best estimate of ARNE's position, given all the information available from odometry, range sensors, and the latest map.

Measure Map Quality The investigation described in this thesis has placed great emphasis on the need for practical experiments and quantitative, statistical, evaluation of the results. For this to be possible, it was essential to have a clearly-defined measure of map quality. The technique employed here is to define a set of 'benchmark' tasks and predict how successful the robot would be at performing those tasks if it used its latest map. Chapter 10 gives some background to the question of quality measurement and describes the metrics used in this thesis.

Explore How should the robot choose the next position from which to examine its environment? The majority of Part III of the thesis is concerned with this question. A range of exploration strategies are designed, implemented, and submitted to experimental evaluation.

1.4 Contributions

This thesis describes an experimental investigation into the complementary issues of map-building and exploration.

The novel contribution of this research consists of:

- The integration of a physical robot, a sonar model, map construction algorithms, and a localisation algorithm into an effective working system;
- The definition and implementation of a novel quantitative measure of map quality;

Cambridge University Press
0521573319 - The Map-Building and Exploration Strategies of a Simple Sonar-Equipped Robot: An Experimental, Quantitative Evaluation
David Lee
Excerpt
More information

- A thorough quantitative and statistical evaluation of the map-building and exploration capabilities of the system, using the quality metric and a variety of exploration strategies. Each strategy is tested in a range of environments.

The system components and the quality metric have been outlined in the previous section and will be described fully in Part II of this thesis. The exploration strategies and the experimental results are described in Part III.

But first, Part I reviews the previous research upon which this thesis is built.