Spatial Function in Animals and Robots

Kenneth D. Harris

Doctor of Philosophy
University College London
1999

Abstract

This thesis describes work aimed at discovering computational processes that could underlie the performance of spatial tasks, both in animals and artificial systems.

In the mammalian brain, the hippocampus is believed to play an important role in spatial function but is also often said to be a storage place for general memories. The first section of the thesis describes a model of mammalian spatial function in which the neocortex processes sensory information to build an egocentric map of the environment, and the hippocampus acts as an autoassociative memory which stores and recalls snapshots of neocortical activity, providing the animal with the ability to recall a complete egocentric map from a few observable landmarks. The model was implemented on a mobile robot equipped with a sonar sensor. The addition of a modelled hippocampus to previously developed map-making software allowed the robot to perform a task similar to the Morris water maze.

The remainder of the thesis describes two novel methods for map construction from sonar sensory data. The first is a neural-network mapping system inspired by models of visual cortex. In this system, the activity of a cell represents the occupancy of a region of space, and lateral connections between cells enforce prior knowledge about probable map configurations. This method produces good results in comparison to other grid-based mapping systems, but is based on heuristics, rather than mathematical principles, and contains many free parameters.

The second is a feature-based mapping system which is derived from Bayes' theorem. It requires an accurate probabilistic model of the robot's sensor, which is derived empirically. The sensor model is used to construct a probability function on the space of all possible maps, and the problem of map construction becomes a search for the most probable map. Global optimisation problems of this size are very difficult, but a "mean-field" approximation allows a tractable solution. Tested with identical sonar data, the new method requires much less sensory data to produce a high-quality map than previous heuristic methods. Although derived without reference to biology, the method has similarities with certain previously proposed models of cortical function, which are discussed.

I have composed this thesis myself and it reports original research that has been conducted by myself unless otherwise indicated.

Acknowledgements

There many people I need to thank for making this thesis possible.

First and foremost is my supervisor, Michael Recce. Michael gave me not only good scientific advice, but also enthusiasm, patience, personal kindness, and generosity, making him not just a good supervisor, but also a good friend.

I have benefited greatly from discussions with the other members of the lab: Giorgio Grasso, Hajime Hirase, David Jack, Adrienne James, Charles King, Asaur Rahman, and John Taylor. Also a big thank-you is due to David Lee, whose work serves as foundation for a large part of the research in this thesis. David Edwards has been very helpful with mechanical and practical issues. David Lee and Neil Burgess very kindly agreed to read my thesis.

On a personal note, I would like to thank my parents, my sister Becky, my girlfriend Helen, and Man Barefoot.

The work presented in this thesis was partially supported by a National Science Foundation Graduate Fellowship. Any opinions, findings, conclusions or recommendations expressed in this publications are those of the author and do not necessarily reflect the views of the National Science Foundation.

Contents

Part I

Chapter 1

Introduction

1.1 Spatial Function

The aim of this thesis is to contribute to the understanding of computational processes that could underlie spatial function, in both animals and artificial systems.

Spatial function will be taken to mean the ability of an agent to control its own movement through space in order to perform useful tasks. This thesis will not be concerned with control of locomotion (the mechanism by which the agent physically moves) although this can be a challenging research problem, particularly in the case of walking. Instead we will concentrate on computational functions necessary for locomotion to be put to a useful purpose.

It is a working hypothesis of this thesis that discovery of computational mechanisms used by the brain to perform spatial functions will facilitate the construction of more useful robots, and conversely, that discovery of techniques that allow better spatial function in robots may shed light on mechanisms of the brain.

1.2 Types of Spatial Ability

This section will present a list of different types of spatial ability. One must be cautious when trying to divide a cognitive ability into separate sub-abilities. It is possible to break down spatial abilities in different ways, and indeed the literature contains many different ways of doing so. However, providing a di-

the thesis, and furthermore, the abilities listed below are sufficiently different that one might expect them to be carried out by different computational mechanisms.

The following sections will define a list of spatial abilities in an approximately hierarchical order, with those listed first being in some sense easier, or necessary for performance of those listed later.

1.2.1 Chemotaxis

Chemotaxis, the motion of a single cell along a chemical gradient, is the simplest form of spatial behaviour, and has been extensively studied. There are quite a variety of mechanisms of chemotaxis, and in particular bacterial and eukaryotic cells seem to employ very different mechanisms.

The mechanisms of bacterial chemotaxis are very well understood, both at a molecular level (see e.g. Alberts et al. 1994), and at a computational level (see e.g. Schnitzer 1993, Strong et al. 1998), and are fascinating because they are so simple. The mechanisms of chemotaxis have been best studied in the bacterium *E. Coli*. In *E. Coli*, locomotion is produced by *flagella*, extracellular spiral filaments that can rotate either clockwise or anti-clockwise. The structure of the flagella is such that, when rotating anti-clockwise, the different filaments line up, and produce a coherent thrust, propelling the bacterium in a forward direction. But when rotating clockwise, the filaments do not align, and the bacterium rotates chaotically. The bacterium's motion alternates between 2 phases: the "run" phase in which the flagella rotate anti-clockwise, and the bacterium moves forwards, and the "tumble" phase, during which the flagella rotate clockwise and the cell spins at random.

The bacterium's cell membrane contains chemoreceptors which detect the rate of change in time of the concentration of beneficial or harmful substances. When the concentration of a beneficial substance is increasing in time, or the concentration of a harmful substance is decreasing, the probability of a tumble phase is low, but when the concentration of a useful substance is decreasing,

is high. Therefore, if the bacterium is heading towards a beneficial substance it will continue to do so, but if it is heading away from one, it change direction to a new, random direction. Notice that the bacterium makes no attempt to measure concentration gradients. In mathematical terms, the bacterium is performing a Monte Carlo search.

Chemotaxis in eukaryotic cells happens by very different mechanisms. For example, amoeba of the slime mould Dictyostelium discoideum have chemoreceptors on their surface that detect the concentration of the chemotactic attractant, cyclic AMP. Where the attractant concentration is high, it causes intracellular actin to polymerise, making the cell extend in that direction. If there is a high concentration on one side of the cell, but a low concentration on the other, the cell will move towards the higher concentration. One might therefore say that the amoeba are performing a gradient ascent search, rather than Monte Carlo.

A study of robots which behave by mechanisms akin to chemotaxis has been given by Braitenberg (1987). Rather than following the traditional engineering methodology of identifying a problem and then devising a solution, Braitenberg follows an approach he calls *synthetic psychology.* He constructs various simple machines, with a small number of sensors and actuators, and examines their behaviour for different patterns of connectivity between the sensors and actuators. The behaviour of these simple machines can be described by words used to describe human behaviour, such as "restless," "cowardly," and "aggressive." This research provides insight into the way that very simple machines may produce complex behaviours; however it also seems clear that machines that are controlled only by simple chemotactic-style mechanisms are unlikely to perform a useful function in the real world.

1.2.2 Stimulus-Response Behaviour

A slightly more advanced form of spatial ability than chemotaxis would involve an agent following a route by making a series of spatial responses to a series

For example a mole navigating through a network of tunnels might turn left when it comes to a junction with a particular smell, but turn right at the next junction, which has a different smell. Alternatively, an agent may perform a given action in response to the observation of certain visual cues. For example, a PhD student navigating through the tunnels of the anatomy department at University College London might know to turn right when he sees the grey lockers, turn left at the top of the stone stairs, and then exit through the glass door.

Following O'Keefe and Nadel (1978), we will refer to the system of navigating along fixed routes by a chain of stimulus-response associations as the *taxon system.* The taxon system is *egocentric*, meaning that for each stimulus-response association the agent performs a motion relative to its current position and orientation, in response to local cues such as a smell, or the presence of a landmark at a certain egocentric position. O'Keefe and Nadel describe the kind of trajectories obtained by this kind of navigation as *routes.*

One might distinguish between two types of stimulus-response association. An association like "When moving fast towards a large rock, stop" might be innate, or acquired very early in life and fixed from then on. However the associations required to follow a route must be learned by experience of traversing that route. The type of learning required is *associative learning*, and has been studied extensively since the early part of this century (Pavlov 1927). There is now considerable evidence for molecular mechanisms that could underlie such learning in both vertebrates and invertebrates (see e.g. Hawkins et al. 1993).

A problem with simple stimulus-response behaviour as a navigational strategy is its inflexibility. If just one of the cues along the route changes, then the entire route is no longer useful. Once off a route, an agent can only find its way back on by random wandering. Furthermore, if a lost agent finds a familiar landmark it will not necessarily be heading in the same direction as it would have been on the route, and so would not know which way to turn.

Stimulus-response association is a general mechanism thought to underlie many different types of behaviour. However, for a task as fundamental as spatial navigation, one might expect techniques to be used that take advantage of special features of the task at hand, such as the Euclidean geometry of space. One such ability is *path integration* or *dead reckoning.* Path integration is the ability to keep track of the total distance and direction moved, even if the animal has followed a complicated or curved path. By requiring animals to leave and return to a nest in the absence of visual or olfactory cues, path integration ability has been demonstrated in insects (Wehner et al. 1996) and mammals (Mittelstaedt and Mittelstaedt 1980).

Path integration is easy to implement in robots – one simply has to keep track of all the movements and turns made by the robot, and perform a vector summation to find the total distance travelled and angle turned. However, for both animals and robots, inaccuracies will build up because the movements and turns cannot be specified with arbitrary precision. Furthermore, because the movement vectors for subsequent movements must be added, errors will increase in time without bound. For robots, the rate of accumulation of path-integration errors can be reduced by using mechanically precise parts, and indeed many robots are designed to be very precise for this reason. But however precise a robot is, path-integration errors will eventually accumulate unless there is some mechanism to reset them.

1.2.4 Localisation

To overcome errors in path integration, an animal must be able to produce or correct an estimate of its position based on the observed location of environmental features. We will call this ability *localisation*, and it has been demonstrated in insects (Cartwright and Collett 1982), and mammals (Morris 1981).

One can distinguish two different types of localisation: *incremental localisation*, the correction of an approximate estimate of position derived from

path integration, and *absolute localisation* which no approximate estimate is available. One might expect incremental localisation to be the easier of the two abilities. Whether this is the case in animals has not yet been addressed experimentally, but it does seem to be the case in robotics, as we shall see later on.

In robotics, localisation is a difficult problem which does not have a general solution in all circumstances. In industrial robot applications, the difficulties of localisation are usually circumvented by artificially enhancing the environment that the robot operates in. For example, routes taken by the robot could be marked by lines painted on the floor that the robot must follow, or the environment could be enriched with artificial beacons or easily identifiable landmarks such as bar-codes.

Unfortunately, this approach leads to a loss of flexibility. A line-following robot can only follow a small number of pre-determined paths, and if new paths are required, new lines must be painted. Artificial beacons do not necessarily limit the paths that the robot can take; but the process of setting up the beacons and programming the robot with their location is time-consuming, and must be done every time the robot is deployed in a new environment.

1.2.5 Mapping

In this thesis, a *map* will mean a description of the nature and location of a set of environmental features, or *landmarks*. We will use the term *mapping* or *landmark identification* to mean the determination of the nature and locations of environmental features from sensory information. Mapping might be better classified as a problem of perception and sensory processing than of spatial function, but it is so clearly relevant for performance of spatial tasks that it is included in this section.

One can make a distinction between two classes of landmark. The first class consists of environmental features which are behaviourally relevant, such as goals or obstacles. The identification of these landmarks is directly necessary for the performance of a task.

vance, such as a distant mountain. While the identification of this second type of landmark is not directly necessary for performance of the task, it could still be useful, for example in localisation.

It is clear that higher vertebrates can perform landmark identification, but not so clear for insects. One argument that insects can perform landmark identification might be that they are capable of localisation, and localisation requires landmark identification. However, it has also been argued that insects localise by simply by memorising entire visual scenes in retinotopic coordinates, rather than memorising the position of identified landmarks (Judd and Collett 1998).

In the case of artificial systems, landmark identification by vision has proved very difficult. Determining the type of an environmental object requires image segmentation and object recognition, and determining its position requires other techniques such as stereo matching. These are all difficult open problems, and most successful robot systems avoid visual landmark identification altogether. With simpler sensors such as sonar, the problem is easier (Leonard and Durrant-Whyte 1992; Lee and Recce 1997). However, with these sensors, the number of identifiable landmark types is small.

1.2.6 Path planning

Once an animal knows its own position, the positions of obstacles in the environment, and the location of a goal, it must plan a route to the goal. In an open field, the animal can simply head along a straight line towards the goal, but in an environment cluttered with obstacles, path planning is a non-trivial problem. There has been a lot of work done on path planning in mobile robotics (see McKerrow 1991, section 8.3.3), most of which makes use of classical techniques such as recursive search.

In animals, there has been much experimental and theoretical research on path planning of limb movements (see e.g. Flash and Hogan 1985) but to my knowledge no studies have addressed the problem of spatial path planning in

animals. Most models of animal navigation are really concerned with some other aspect (such as the role of the hippocampus), and make do with an *ad-hoc* model of path planning.

1.2.7 Behavioural selection

The abilities listed above give an agent the ability to move to a particular location in space. However, before the agent goes anywhere it must decide where, if anywhere, to go.

For an animal, the choice of destination is going to be influenced by many factors such as the internal state of the animal (e.g. whether it is hungry or thirsty), the presence of stimuli (such as a predator or a visible source of food), or even the time of day. Because so many factors are involved, is very hard to characterise neurobiological processes underlying behavioural selection, but for some recent reviews see Kristan and Shaw (1997) and Owen (1997).

For robots, two approaches to behavioural selection have been in common use. The first approach could be called *top-down* control. With this approach, an executive control module combines the task the robot must perform with a world model to derive an appropriate action for the robot to perform. This approach is natural to computer programmers used to working with reliable digital machines, but when used to control robots in the real world, it can lead to a lack of robustness. In order to overcome this, a second approach, called *behaviour-based robotics*, was developed. In this approach, there is no executive control module, and several different modules all contribute to the robot's behaviour in parallel. Behaviour-based robotics will be discussed further in chapter 4.

1.3 Aim of the thesis

The aim of this thesis is to clarify some of the computational processes that could underly spatial function. These will be studied by computational modelling of neuronal structures thought to be necessary for spatial function in the mammalian nervous system, and by theoretical and empirical investigation

spatial tasks by mobile robots. This thesis will not have anything to say about the question, interesting though it is, of spatial function in lower animals.

The previous section listed a set of spatial abilities, in increasing order of difficulty. For the early ones, the computational processes needed to perform the tasks are clear: no one questions the computational mechanisms of chemotaxis, the molecular mechanisms underlying chemotaxis are well understood, and robots have been constructed that perform similar tasks reliably. Stimulus-response behaviour and path integration are also usually regarded as areas in which not much more research is required. But for the more difficult spatial functions, little is known about how animals perform these functions, or how they might be successfully implemented on mobile robots. One might argue that is is necessary to solve the simpler spatial problems before attempting the more difficult ones (Brooks 1986). In this thesis, we concentrate on the intermediate problems of localisation and mapping.

The research described in this thesis falls into two parts. In part II, the focus is on biology. We present a theory for the performance of absolute localisation in mammals with the aim of unifying two apparently conflicting previous ideas about hippocampal function. The theory is qualitatively discussed with respect to previous biological experiments and models, and predictions are made for new experiments. The theory is then tested by implementation on a mobile robot, and implications for the problem of absolute localisation in mobile robots is discussed.

In part III, the focus is on engineering. We describe two new techniques for map building by sonar-equipped mobile robots. The aim here is not to model any particular biological systems or brain structures, but to produce a mapping algorithm that makes maximum use of the available sensory data. The first technique is a neural network based system, inspired by models of visual cortex. The second is derived purely from probability theory, without any reference to biology. However, the resulting algorithm turns out to have some distinct computational similarities to certain previously proposed models

of cortical function.

1.4 Structure of the thesis

Chapters 2 to 4 are literature review. In chapter 2, we discuss experiments aimed at revealing what spatial functions animals are capable of, and how these functions may be performed by the brain. In chapter 3, we discuss models of the hippocampus, a brain structure thought to be centrally involved in spatial function. In chapter 4, we describe approaches which have been taken to the problem of giving spatial abilities to mobile robots.

In chapter 5, we describe a new theory for how spatial functions may be executed by the brain, concentrating particularly on the role of the hippocampus. In chapter 6, we describe how this theory was tested by implementation on a mobile robot.

Chapters 7 to 9 describe new methods for mapping in the case of sonar-equipped mobile robots. Chapter 7 describes a neural network based system inspired by models of visual cortex. Chapters 8 and 9 describe a second system derived from probability theory. Chapter 8 describes the production of a probabilistic model of the sonar sensor necessary for the mapping algorithm. Chapter 9 describes the algorithm itself, and gives the results of experiments testing the algorithm with real data.

Finally, in chapter 10, we summarise the conclusions of the investigations of the thesis, and discuss some possibilities for future work.

Chapter 2

Spatial function and the brain in mammals

This chapter will describe previous experiments which have tried to find out precisely what spatial functions animals are capable of performing, and how they are carried out in the brain. The literature on this subject is vast, and we can only touch on some of it. Due to constraints of space, we must leave invertebrates behind and concentrate on mammals. In particular most of the experiments described here involve rodents.

When describing biological research, it is necessary to use biological terms. Appendix A contains a glossary of some of the more obscure biological and psychological terms used in this chapter, for readers who may not be familiar with them.

2.1 The spatial abilities of mammals

Mammals appear to be capable of all the spatial abilities listed in chapter 1. They are clearly capable of finding the source of an odour, which may be interpreted as an analog of chemotaxis. Early experiments have also clearly shown that they are capable of producing spatial responses to stimuli, and using stimulus-response chains to navigate using routes (e.g. Watson 1907). However, early experiments also showed that rodents can perform tasks not explainable by stimulus-response behaviour alone. For example, Tolman (1948) showed that rats are capable of taking short-cuts across areas of space which

they have not previously explored.

The ability of rodents to perform path integration was demonstrated by Mittelstaedt and Mittelstaedt (1980), in an experiment where a mother gerbil was required to leave her nest to retrieve a pup in darkness, and would return to the nest along a straight line path even if her trajectory in finding the pup was highly non-linear.

Localisation ability in rats was demonstrated by Morris (1981) using a task known commonly, if bizarrely, as the *water maze.* In this task, a rat is placed in a pool of opaque milky water in which there is a submerged platform, which is invisible to the rat. After repeated trials the rat learns to head directly to the platform from any of several starting points on the edge of the pool. In order to solve this task, the rat must use environmental features to decide its location, remember the location of the goal, and move along a vector from this starting point to the goal. Morris initially claimed that localisation was possible from parts of the environment that the animal had not previously visited (Morris 1981; see also Sutherland and Linggard 1982), and called the ability to transfer spatial knowledge gained in one place to another *instantaneous transfer.* Recently, however, the suggestion that animals can perform instantaneous transfer has become more controversial (Sutherland et al. 1987; Keith and McVety 1988; Chew et al. 1989; Keith 1989; Whishaw 1991; Alyan 1994).

2.2 The effect of brain lesions on spatial ability.

Lesions to the neocortex in rats have been shown to have a serious effect on spatial ability, disrupting the ability to travel to both visible and memorised goals at fixed egocentric and allocentric locations (e.g. Kolb et al. 1994; Save and Moghaddam 1996). Data suggest that the cortical areas involved in spatial function include the posterior parietal cortex and medial prefrontal cortex (but see also de Bruin et al. 1994). In humans, lesions to the right

a review), a condition in which the patient has an impaired ability to attend to stimuli in the contralateral hemisphere. The deficit seems to be one of egocentric, rather than allocentric, spatial representation, and effects mental imagery as well as the perception of physical stimuli (Bisiach and Luzzatti 1978). A general consensus is that the posterior parietal cortex is necessary for explicitly spatial functions, such as the determination and representation of object locations relative to the body (in the language of chapter 1, we would say it is necessary for mapping), and that the prefrontal cortex is necessary for more general high-level planning and behavioural functions not only in a spatial context (see e.g. Kolb and Whishaw 1995).

Hippocampal lesions in rats have been shown to impair performance in the water maze task described above (Morris et al. 1982). Recent PET studies have confirmed that the hippocampus is also involved in spatial function in humans (Maguire et al. 1998). However the most obvious effect of human hippocampal lesions is strikingly different – the syndrome of "global anterograde amnesia" (Scoville and Milner 1957). In this syndrome, the patient is able to retain information for a small length of time, for a few minutes or until distracted, but not for longer. Only *declarative* or *episodic* memories, which involve the recall of specific events or facts seem to be affected, whereas *procedural* memories, such as motor skills acquired through practice are not. Also, episodic memories formed a long enough time before the hippocampal damage are not affected.

2.3 The hippocampal theta rhythm

In rats, the EEG activity of the hippocampus falls into two distinct states, depending on the behaviour of the animal.

During locomotion, hippocampal EEG shows an approximately sinusoidal 7-12 Hz rhythm, called *theta* (Vanderwolf 1969). While the animal is awake, the correlation between theta and motion is very strong, and there is also evidence that the frequency of the theta rhythm is correlated with the speed

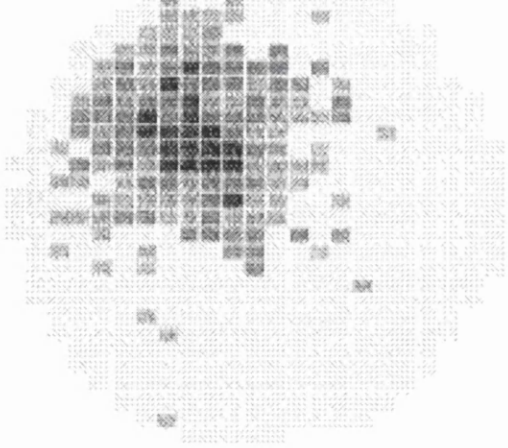

Figure 2.1: The firing rate map of a typical hippocampal place cell, reprinted from Muller et al. (1991). The shading represents the average firing rate of a single CA1 cell at each position in a cylindrical arena.

of the animal (Recce 1994). The theta rhythm is also observed during REM sleep.

During slow-wave sleep and waking behaviours that do not involve movement (such as resting, feeding, or grooming), the hippocampal EEG shows a pattern called *large irregular activity*, which consists of a broader ranger of frequencies, and is punctuated by short-duration *sharp waves*.

2.4 Properties of hippocampal place cells.

Recordings in rat hippocampus have revealed the existence of *place cells* (see figure 2.1), which fire when the rat is in a specific part of an environment (O'Keefe and Dostrovsky 1971; O'Keefe 1976; Muller et al. 1987). Normally place cells have a stable activity pattern, corresponding to a single region, or *place field*, in an environment. In a single environment, place cells have been recorded over many days and appear to have stable place fields (Thompson and Best 1990). When rats randomly search for food in an open field environment, the firing of place cells is largely independent of head direction. However, when a rat is constrained to move back and forth along a fixed one-dimensional trajectory, place cell firing is also modulated by the direction of motion. (McNaughton et al. 1983; O'Keefe and Recce 1993; Muller et al.

Several studies have demonstrated that the firing of place cells is driven by visual sensory input. For example, if all cues in an environment are rotated, the firing areas of place cells rotate with it (O'Keefe and Conway 1978; Muller and Kubie 1987). Place cells can also be controlled by senses other than vision, and indeed normal place cells are found even in blind and deaf rats, provided that the environment is sufficiently rich in tactile cues to allow for localisation (Hill and Best 1981; Save et al. 1996).

If a rat is placed in a well-explored environment, and the lights are switched off, the place cells continue to fire in their usual place fields, although the place fields become slightly broader (Quirk et al. 1990). This broadening may be attributable to a build up of path-integration errors that, in the absence of visual input, cannot be corrected by localisation. If the lights are switched off *before* the rat is placed in the familiar room, a new set of place cells begins to fire, and the new pattern of activity persists after the lights have been switched back on.

When a rat is allowed access to a previously unreachable part of a familiar environment, previously silent place cells begin to fire with place fields in the previously unreachable area. However, place fields in the familiar part of the environment are unchanged (Wilson and McNaughton 1993).

Finally, neurons in other areas near the hippocampus, such as the dentate gyrus (Jung and McNaughton 1993), the subiculum (Sharp and Green 1994), and the entorhinal cortex (Quirk et al. 1992) also show place-dependent firing. The size of spatial firing field of the cells varies among these regions (Barnes et al. 1990). Neurons of the entorhinal cortex show broad place fields, but in contrast to hippocampal place cells, the firing patterns medial entorhinal cells are topologically transformed when the shape of the environment changes (Quirk et al. 1992).

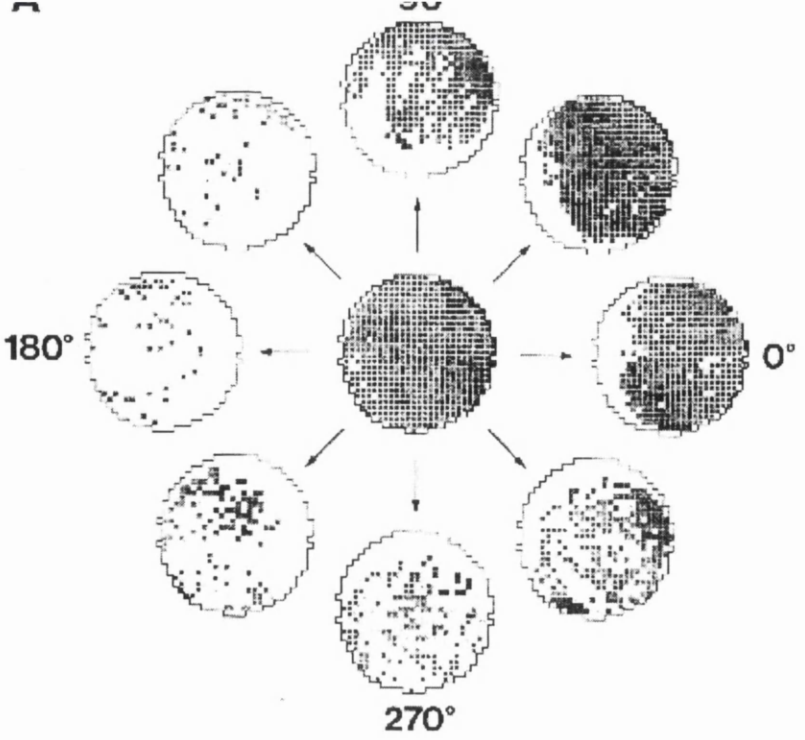

Figure 2.2: A typical head-direction cell, reprinted from Taube et al. (1990a). The central circle shows firing rate as a function of position, averaged over all head directions. It can be seen that the cell has little positional preference. The peripheral circles show firing rate as a function of position when the rat's head is pointing in the appropriate direction. The firing rate is strongly modulated by head direction, but not by position.

2.5 Head direction cells

Recordings in the postsubiculum of rats have identified *head direction cells* (Ranck 1985; Taube et al. 1990a) whose firing rates are tuned to the animals allocentric head direction, but largely independent of the animals position in space (see figure 2.2). Since the initial discovery of head direction cells in the postsubiculum, they have also been found in the anterodorsal nucleus of the anterior thalamus (Taube 1995), the laterodorsal thalamic nucleus (Mizumori and Williams 1993), the retrosplenial and parietal cortical areas (Chen et al. 1990; McNaughton et al. 1991) and the striatum (Wiener 1993).

Like place cells, the firing of head direction cells is determined by visual input (Taube et al. 1990b). Also like place cells, if the lights are switched out when an animal is in a familiar environment, head direction cells initially fire

tial firing directions begin to drift (Mizumori and Williams 1993). This may again be ascribed to the build-up of path-integration errors in the absence of localisation.

If the positions of significant landmarks are changed during an exploration session, the directional preference of head direction cells may follow either the landmarks positions, or path-integrated ideothetic information. In general, it appears that rats use landmarks preferentially over ideothetic input, provided experience has proved those landmarks to be stable (Taube and Burton 1995; Knierim et al. 1995). Multiple cell recordings have shown that the head direction and place cell systems are coupled, in that an environmental manipulation causes a consistent rotation in the preferred directions of all recorded head direction cells (Taube and Burton 1995), as well as causing a rotation in the position of hippocampal place fields (Knierim et al. 1995).

Chapter 3

Theories of hippocampus

In this chapter we will describe some previous theories of hippocampal function in mammals. Again, the literature is vast, and we can only mention some of the many theories.

As described in the last chapter, the evidence from rats points to an explicitly spatial role for the hippocampus, but the evidence from humans points to a more general memory function. This chapter will be in two parts, the first describing explicitly spatial theories and the second more general theories.

3.1 Explicitly spatial theories

3.1.1 The cognitive map theory

The first and most influential theory of hippocampal involvement in spatial function is the *cognitive map* theory of O'Keefe and Nadel (1978). O'Keefe and Nadel distinguished two separate systems involved in navigation: the *taxon* system, previously described in chapter 1 of this thesis, which allows for route following by stimulus-response association; and the *locale* system which involves the use of a *cognitive map*, for which the proposed substrate is the hippocampus. The cognitive map is like a geographical map, in that the location of the animal and the external cues are represented in an fixed, world-centred frame. An important feature of the locale system is that goals are treated in exactly the same way as other environmental features. This is in contrast to taxon navigation, where goals act as reinforcers of learned behaviour.

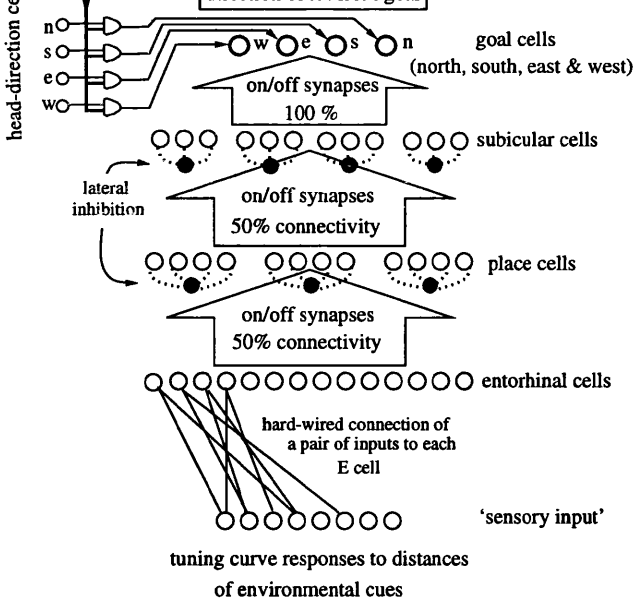

Figure 3.1: The model of Burgess *et. al.*, reprinted from Burgess *et al.* (1996), figure 6. Cells in the input layer code for the distance, but not bearing, from the rat to identified landmarks, and cells of the output layer determine which direction the rat should move in to reach the goal.

During the twenty years since its publication, the cognitive map theory has received wide-ranging experimental support, including data from lesion experiments in rats (see Jarrard 1993 for a review), and the properties of place cells.

The cognitive map theory is a qualitative theory. It ascribes, in words, a function to the hippocampus, but says little about how this function might be computed by the cells of the hippocampus. O'Keefe and Nadel (1978, section 4.8) did also tentatively sketch a neural model which could act as a spatial map, but it their own words, it was "far off the mark" and would be "hopefully modified or rejected in the near future." Nevertheless, the cognitive map theory has been vastly influential on later computational models, and many computational models of hippocampal function propose an explicitly spatial function for the hippocampus.

3.2.2 The Model of Burgess et al.

Burgess et al. (1994) describe a computational model of hippocampal involvement in rat navigation based on visual cues. At the heart of the model is a layered neural network containing cells corresponding to various areas of the hippocampal formation – the entorhinal cortex, hippocampus, subiculum, head-direction system, and another, unspecified area one synapse downstream from the subiculum which contains *goal cells* which code for the location of a goal relative to the position of the rat (see figure 3.1). There is also a population of "sensory cells," each of which codes for to the distance to an identified landmark in the environment. No sensory cell codes for more than one landmark, and no cell codes for the angular bearing to its associated landmark.

Activity is propagated from the sensory layer to the entorhinal layer by fixed connections, and from the entorhinal to the hippocampal and subicular layers by connections which are trained by competitive learning. An important feature of the model is that the firing of cells from the entorhinal cortex onwards is *phase coded* (O'Keefe and Recce 1993). The phase of firing of entorhinal cells (with respect to the theta rhythm) depends on the average angle from the rat's heading direction to the cues to which it responds, with cues ahead of the rat making the cell fire at a late phase, and cues behind the rat making it fire at an early phase. However, a neural mechanism by which this entorhinal phase phenomenon could be generated was not given.

The goal cells are trained by a type of reinforcement learning: when the rat reaches a goal, it looks around in 8 directions (north, northeast, east, etc.), and for each direction, the goal cell corresponding to that direction receives a reinforcement signal at the late phase of the theta rhythm, and all active input synapses are strengthened. After training, the output of the goal cells forms a population vector which points in the direction of the goal from any location.

A notable way in which this model differs from the cognitive map theory is that in this model, goals and obstacles are treated differently to other en-

population of goal cells, and other environmental objects possessing a population of sensory cells. Furthermore, goal cells have yet to be identified in the mammalian nervous system.

3.1.3 The hippocampus as the substrate for path-integration

McNaughton et al. (1996) have proposed that the hippocampus is the substrate for path-integration ability in mammals.

In their model, recurrent excitatory connections within the hippocampus (specifically, in the CA3 field) are pre-configured in a way that implicitly defines a set of two-dimensional surfaces. The two-dimensional surfaces do not map directly to the two-dimensional CA3 cell layer, but are an abstract construct that arises from the pattern of synaptic weights. The synaptic weight between two neurons is proposed to be an approximately Gaussian function of distance on the abstract surface. As a result, given appropriate global inhibition, each neuron defines the centre of an "attractor", i.e. a fixed firing pattern that is kept active by the recurrent connections.

McNaughton et al. (1996) also propose that connections between these abstract surfaces and another area, possibly the subiculum, allow the animal to perform a path integration function. The cells in this second area fire in relation to the rat's position and current heading direction, and project to the hippocampus with synapses biased towards the rat's current heading direction. In this way, the location of hippocampal activity pattern moves in the correct direction during linear motion. Modifiable connections between the hippocampus and neocortex allow the animal to record the positions of environmental landmarks, in the reference frame described by the abstract surfaces of the hippocampus. When the rat is then introduced to a familiar environment, these connections allow the rat to localise using visual information.

McNaughton et al. (1996) explain the results of many experiments in the framework of this model. However, at a conceptual level, many questions remain. For instance, the model requires a very intricate pattern of synap-

the more difficult because the abstract surfaces do not map onto the two-dimensional arrangement of the cells. No mechanisms are given by which the synapses might be trained, or formed during development.

Whishaw et al. (1995) show that rats with fimbria-fornix lesions are able to find a hidden platform in a water maze if they are originally (and extensively) trained to swim to a visible platform in the same location, but not if they are originally trained with a hidden platform. They take this as evidence that the hippocampus is responsible not for spatial memory, but for some motor function such as path integration. However, recent preliminary evidence suggesting that the hippocampus is not necessary for path integration (Alyan et al. 1997) would appear to deal path-integration theories a fatal blow.

3.1.4 The model of Brown and Sharp

Brown and Sharp (1995) describe a computational model of the involvement of hippocampus, postsubiculum, and nucleus accumbens in the performance of the water maze task.

Like Burgess *et. al.*, they assume a population of sensory cells which code for the egocentric location of identified stimuli, and a population of place cells which receive input from the sensory cells and are trained by competitive learning. However, they differ on the way in which the place cell output is used. In Brown and Sharp's model, the place cell output, along with head direction cell output, is the input to a reinforcement learning system located in the nucleus accumbens. The nucleus accumbens network contains two populations of cells, whose output produces left and right turns, and the input synapses of these cells are trained by a type of temporal-difference rule (Barto and Sutton 1981).

This model is essentially a stimulus-response system, and the fact that it can solve a water maze style task casts doubt on the commonly held belief that a locale system is necessary. A stimulus-response system would certainly not be able to perform instantaneous transfer, but as the authors point out,

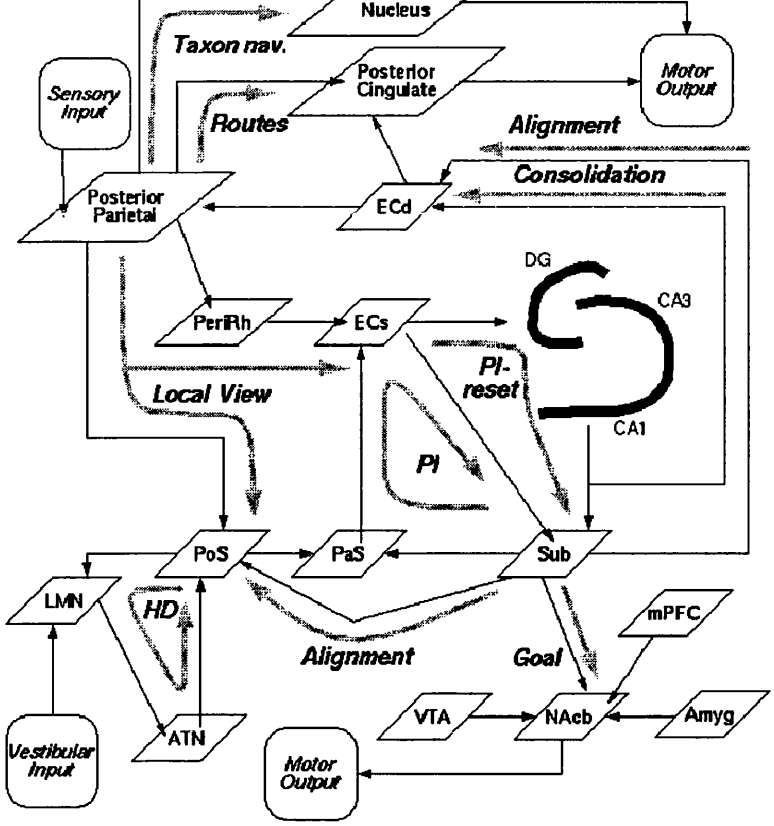

Figure 3.2: The models of Redish and Touretzky, reprinted from Redish (1997), figure 3.15. Specific roles are proposed for a total at least 16 separate brain areas, and different types of spatial function are processed by 11 different functional pathways (shown in grey).

the evidence for instantaneous transfer is not clear. However a model such as this one cannot explain the ability of animals to take short-cuts across unexplored regions, as this would require the animal to pass through areas of space in which no place cells had been formed, and therefore no output could be produced.

3.1.5 The models of Redish and Touretzky

In a series of papers, Redish and Touretzky (Touretzky and Redish 1995; Touretzky and Redish 1996; Redish and Touretzky 1997) propose a very detailed theory of rodent navigation including specific functions for the hippocampus, subiculum, postsubiculum, parasubiculum, basal ganglia, amygdala, lateral mammillary nuclei, anterior thalamic nuclei, and the entorhinal,

tices, as summarised in figure 3.2. They propose that the hippocampus has *three* functional modes: A storage mode, during which LTP occurs in the hippocampal formation but those synapses undergoing potentiation show little or no transmission; a recall mode in which no LTP occurs, and the dentate gyrus is silent, although direct connections from entorhinal cortex to CA3 and CA1 are active; and a replay mode in which all inputs to the hippocampus are silent, and the activity of hippocampal cells is due only to internal dynamics. A key feature of their model is the *path integrator* module, whose job is to update the position of the animal in an allocentric coordinate frame, and whose proposed anatomical substrate is a loop consisting of the subiculum, parasubiculum, and superficial entorhinal cortex. Notably, this loop does not include the hippocampus.

To make such a detailed model of neural function might be considered premature, given the very large number of brain structures modelled, and the level of experimental knowledge at the present time. However a model this detailed makes a great many experimental predictions which should be easily verifiable or falsifiable. For example, it predicts that while the superficial entorhinal cortex, parasubiculum, and subiculum are all necessary for path-integration, the hippocampus and the remainder of the neocortex (including the deep layers of the entorhinal cortex) are not.

3.2 Non-spatial theories

3.2.1 Episodic memory

Another school of thought, derived from the effects of hippocampal lesion in humans, does not give the hippocampus an explicitly spatial function, but instead a more general memory function.

This theory was first put forward in a computational form by Marr (1971). Marr proposed that the archicortex functions as a *simple memory*, that temporarily stores patterns of neocortical activity. Marr derived a neuronal ar-

memory), and mapped this architecture onto the anatomy of the hippocampal formation. One of the crucial mechanisms of this model, called the *collateral effect*, uses the activation of the recurrent collaterals between hippocampal pyramidal cells to produce a completed pattern from a partial or corrupted input. The auto-associative function that Marr proposed for the hippocampus can be contrasted with his earlier theory of the cerebellum (Marr 1969), which he proposed learns stimulus-response pairs (acts as a *hetero-associative memory*).

Subsequent papers have re-examined and refined Marr's ideas on the cellular basis and mechanisms of auto-associative memory in the hippocampus using modern simulation and analytical techniques (Gardner-Medwin 1976; Willshaw and Buckingham 1990; Treves and Rolls 1994).

3.2.2 Sequence memory

Marr's original theory was that hippocampus stored "snapshots" of cortical activity, i.e. memorised the cortical firing pattern active at one particular instant in time. Later theories, however, have proposed that the hippocampus actually stores sequences of cortical activity patterns (Buzsaki 1989; Wallenstein and Hasselmo 1997) – thus the hippocampus is not just recording a single event, but a whole sequence of events, like a cinema film as opposed to a still picture.

It has also been proposed that such a sequence memory, acting as a sequence predictor, could play a role in spatial function (Levy 1996). By memorising a sequence of places traversed along a route, the hippocampus could, on subsequent visits, recall the sequence of places travelled and thereby provide the animal with the ability to traverse the same route again. It is also claimed that the animal may be able to eliminate loops in previously travelled routes using such a mechanism – however a role for sequence memory in more general spatial functions has not yet been elucidated.

3.2.3 Conditioning models

Yet another school of thought for hippocampal function derives from classical conditioning experiments.

Although the hippocampus is not necessary for acquisition of the simple rabbit eyeblink conditioned response (Schmaltz and Theios 1972) there are other, more complicated types of classical conditioning which cannot be learned by hippocampal animals. Sutherland and Rudy (1989) proposed that the types of conditioning that are hippocampal-dependent are those that require the learning of a response to a configuration of stimuli, as opposed to a simple response to a single stimulus. They also claimed that the learning of such configural associations could also underlie many of the locale spatial tasks that are dependent on the hippocampus. However, the learning of a response to a set of stimuli in a spatial context is precisely what O'Keefe and Nadel said did *not* require a hippocampus! Also, the type of memory required is heteroassociative, which requires a different neuronal structure than the autoassociative architecture proposed by Marr (1971).

Gluck and Myers (1993) describe a high-level connectionist model of hippocampal involvement in classical conditioning, in which the function of the hippocampus is modelled by a *predictive autoencoder*. The predictive autoencoder is simulated by a 3 layer neural net that is presented with an input and trained by backpropagation on a vector composed of the desired output and a copy of the original input. The representations formed in the hidden layer of this net are then used to train the hidden layer of another three-layer net, which symbolises the cortex. The authors do not claim that backpropagation is actually used in the brain, but that the learning system that does exist in the hippocampus produces similar effects. This model, and its extensions, show good agreement with a wide variety of conditioning experiments under various manipulations (Myers and Gluck 1994; Myers et al. 1995; Myers et al. 1996), but appears to have less bearing on a role for the hippocampus in spatial function or episodic memory.

Ideally, a single hippocampal model would describe a single computational function that is necessary for the performance of all the tasks that the hippocampus seems to be necessary for, in all problem domains. An ideal model would also describe how this function is performed by the cells of the hippocampus, making reference to the anatomy of the hippocampus and the physiology of hippocampal cells. And finally, an ideal model would make predictions for new experiments, either behavioural of physiological, which could confirm or disprove the model.

Such a model is clearly a long way off. Most of the models described in this chapter restrict their attention to a single problem domain, such as spatial function, memory, or conditioning, and either propose that the experiments suggesting hippocampal involvement in other problem domains are wrong, or simply do not address the issue of hippocampal function in other domains.

Many of the qualitative models describe a type of task that the hippocampus is required for (e.g. "the hippocampus is required for spatial function") rather than describing a computational function (e.g. "the hippocampus functions as an autoassociative memory") and explaining how it may be performed at a cellular level. Furthermore, the roles proposed by the different models are very diverse, and at first sight seem incompatible.

Of the models described above, only Sutherland and Rudy attempt to provide a synthesis, proposing a single computational function which could unify the apparently different roles for the hippocampus. However, they do not say how this function could be performed a a cellular level, and the heteroassociative memory function they propose requires a different neuronal architecture to the autoassociative memory architecture which appears to match the cellular structure of the hippocampus better.

Those models that are described at a cellular level, on the other hand, are always concerned with the role of the hippocampus in a single problem domain, be it spatial function (Burgess et al. 1994), or memory (Marr 1971; Treves and

within their own problem domains, but make no reference to other domains. A model that could be described at a cellular level, and which provided a synthesis of the different problem domains, would be a step forward.

In chapter 5 we will propose a framework for performance of spatial functions in the brain in which the hippocampus functions as an autoassociative memory. Many modellers have already proposed that the hippocampus is an autoassociative memory, described how this function may be achieved at a cellular level, and described the role of autoassociative memory in episodic memory function. By giving evidence that autoassociative memory may also be involved in spatial function, we hope we have moved closer to a unified theory for the hippocampus.

Chapter 4

Spatial function in mobile robots

It is a working hypothesis of this thesis that the discovery of computational mechanisms underlying spatial function in animals will be useful for the construction of better robots, and that the discovery of computational mechanisms useful for robots may shed light on the mechanisms of spatial function in the brain.

The type of robots for which spatial function is necessary are those that are free to move around in space, or *mobile.* In this chapter we will describe the hardware from which mobile robots are constructed, give a brief history of mobile robot research from its beginnings after the second world war, and then describe some more modern techniques of mobile robot control, concentrating mainly on the problems of localisation and mapping, with which this thesis is concerned.

4.1 Mobile robot hardware

A mobile robot is a robot that can locomote, i.e. move from place to place. For mobile robots, there are two practical methods of locomotion – walking or wheels. While walking has several advantages over wheeled locomotion, such as the ability to deal with uneven terrain, there are many difficulties involved in the control of walking machines. For this reason, most mobile robots used in research have wheels.

Several types of sensor are used with mobile robots. Video cameras can, in theory, provide the robot with a great deal of useful spatial information. However, computer image processing is difficult, and extracting spatial information from video images is in itself a challenging research topic.

Rangefinders, sensors which measure the distance to obstacles, have proved popular with mobile robot researchers because the amount of preprocessing required is much less than with vision, and because their output is explicitly spatial, and thus directly useful for low-level spatial functions such as obstacle avoidance. Several types of rangefinder are in current use. Laser rangefinders are very accurate but also very expensive. Infra-red rangefinders are useful for detecting nearby objects, but due to limited power output, have a limited measurement range (60cm for the "Sensus 300" infra-red sensor found on the popular Nomad 200 robots). Also, because infra-red rangefinders work by measuring the amplitude of the returning radiation, their output depends on the nature of the reflecting surface as well as its range, and so cannot be directly interpreted as a distance to the object.

Sonar rangefinders are very popular because of their low cost, and because they are well characterised. Unlike infra-red rangefinders, sonar rangefinders measure distances by the time taken for reflection, and so, under appropriate circumstances, they can precisely measure the distance to an object, whatever its nature. However, these "appropriate circumstances" do not always occur. Briefly, there are two reasons why sonar does not always measure the distance to the nearest obstacle. The first is that some objects, particularly smooth walls, act as *specular reflectors*, which reflect the sonar beam like a mirror reflects light. If a sonar beam is obliquely incident on a specular surface, either no reflection is received, or a "multiple specular reflection" will occur, where the beam bounces off the specular wall, onto a secondary target, and back again.

The second reason is finite beam width. The beam coming out of the sonar transducer is approximately 30^o wide, and reflections may be caused by sound energy from any part of the beam. If a reflection is caused by energy from

the sensor to the target. Early work with sonar regarded finite beam width as a source of "inaccuracy". But in chapter 8 we will see that even for oblique reflection angles the output of the sonar sensor can be a reliable source of information, although it does not give the direct distance to the target.

4.2 Early approaches to mobile robotics

The history of robotics research reflects the scientific fashions that have come and gone since modern robot research started after world war II. At this time, the fashionable buzz-word was *cybernetics*, meaning the design of machines inspired by biology. The most famous robotic example of this trend was Gray Walter's "Tortoise" (Walter 1950), a simple machine that showed spatial ability that, in the language of chapter 1, we would give the level of chemotaxis. The tortoise was equipped with an optic sensor and a bump detector, allowing it to wander without mishap, and to recharge its own batteries by approaching a (brightly lit) charging hutch.

In the 1960s, the "artificial intelligence" movement was underway, characterised by a symbolic approach to computation, instead of an emphasis on trying to mimic living systems. The most famous robot of this time is Shakey, a five foot tall vision-equipped mobile robot (Nilsson 1984). Shakey's software was adapted from the MIT blocks world program, and as a result, all objects in the robots environment had to be uniformly coloured and have flat faces, and even the room in which the robot operated had to be specially constructed. A typical task for Shakey was to identify a given object, and push it to a given position. Due to the classical AI type search algorithms used in its control system, the robot had a limited ability to deal with unexpected events, and took about an hour to plan each move.

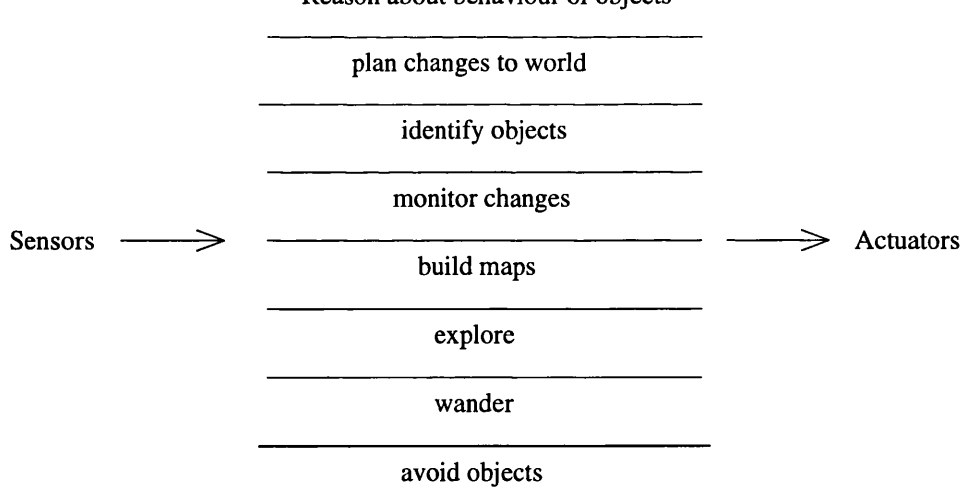

Figure 4.1: The layered control system for mobile robotics proposed by Brooks (1986). Each layer defines a behaviour, more complex than the last, with higher layers able to override the output of lower layers.

4.3 Behaviour-based robotics

By the 1980s scientific fashion was swinging away from artificial intelligence, and a new paradigm emerged in robotics. The limited success of rule-based approaches to robotics led Brooks (1986) to suggest a new, less ambitious, approach to robotics, in which the need for explicit symbolic models of the world was played down. However, Brooks was not simply proposing a return to the simple cybernetic ideas of the 1950s. The key difference was that Brooks proposed a specific computational scheme, called the *subsumption architecture*, which allows for the smooth integration of low-level behaviours such as obstacle avoidance with high-level behaviours such as path planning. In the subsumption architecture, each type of behaviour has its own computational module, with higher-level modules being able to override the output or input of lower-level modules. One of the advantages of the subsumption architecture is that it allows lower levels of behaviour to be developed and tested before higher ones. Indeed, Brooks' original paper described an 8 layered control system, of which only 3 levels were actually implemented (see figure 4.1).

Behaviour-based robotics is often interpreted as an approach in which robots simply react to their sensory inputs, and do not perform any kind

intention. To quote Brooks (1982, page 4):

> We want to build cheap robots which can wander around human inhabited space with no human intervention, advice or control and at the same time do useful work. Map making is therefore of crucial importance even when idealised blue prints of an environment are available.

Brooks was not proposing that robots do not need to build maps. Indeed, figure 4.1 contains a layer with precisely this function! He was merely saying that it is a better idea to concentrate on the more simple behaviours before attempting complex ones such as map making, or the more "intelligent" higher behaviours even further up on figure 4.1.

4.4 Modern techniques

Brooks' ideas have had a profound effect on thinking in robotics, and most robotics researchers today would agree with the idea of separate modules for lower-lower functions, even if they do not strictly follow the subsumption architecture as originally described. The emphasis on low-level sensory and motor processes, instead of high-level planning is also widely accepted.

Mobile robotics has been a very active field in the last 10 years, and it is again impossible to give a comprehensive review of the literature in this area. Instead I will just describe some work in the most relevant fields to this thesis: sonar-based mapping and localisation.

4.4.1 Sonar-based mapping

In sonar mapping, two methods of map representation have been in common use up to now.

In *grid-based representation*, the two dimensional space of the robot's environment is overlaid with a square grid. For each grid segment, information is

inal grid-based mapping system (Elfes 1987), a single number was stored for each segment, corresponding to a "probability" of occupancy of that segment. However, this number did not correspond to a probability in any strict sense, as it ranged from -1 to 1. In later work (Moravec 1988), actual probabilities were used, and Bayes' rule was used to derive the formula for updating these probabilities from new sensory data. However, there are two problems with such an approach, one practical and one mathematical.

The practical problem is that grid-based methods like these cannot detect specular walls. As such walls are common in the type of indoor office environments where mobile robots might be frequently used, this is clearly a problem. Because the reflection properties of specular walls depend on the orientation of the wall, a mapping system that can detect specular walls must also represent the orientation of the wall. Lim and Cho (1992) addressed the problem of specular walls by including an extra set of variables for each grid segment which code for the probability of different orientations of a wall passing through a grid segment. For more details of this approach, see chapter 7.

The mathematical problem with grid-based mapping systems is that they assume that the probability of occupancy of different cells is independent. This is clearly not true for indoor environments – if it was, the most likely configurations would be random scatters of occupied and unoccupied cells – but in reality you see much more structured patterns like walls.

The second type of map representation in common use is called *feature-based representation.* In feature-based representation, the map consists of a list of environmental features, such as walls and pillars, and the positions of these features are represented by Cartesian coordinates. Feature-based mapping has its roots in work describing the properties of sonar reflection from various types of reflector (Hallam 1986; Kuc and Siegel 1987; Bozma and Kuc 1994). The seminal work on feature-based mapping is that of Leonard and Durrant-Whyte (1992). In this system, a simplified version of Kuc and Siegel's sensor model is assumed. A crucial feature of the mapping system is the preprocessing

neighbouring sonar returns of approximately equal length. In a sense, the extraction of RCDs is required because of the simplified sonar model: when an RCD is detected, it indicates a strong specular reflection, a case to which the model applies. However, sonar returns which do not form part of an RCD often correspond to weak returns not covered by the simplified model, and so must be ignored.

A second issue in feature-based mapping is the *correspondence problem*, meaning the problem of assigning individual sonar readings to the features which caused them. The original work of Leonard and Durrant-Whyte used a set of heuristics to do this. Later work (Cox and Leonard 1994; Maskarov and Durrant-Whyte 1995) used a more rigorous approach called *multiple hypothesis tracking*. In this approach, probability theory and tree-searching techniques are used to find an optimal correspondence between measurements and features, and once the correspondence has been made, Kalman filters are used to track the positions of individual features.

In chapter 9, we will describe a new approach to feature-based mapping based on a more accurate model of sonar behaviour is used, in which all sonar readings can be used, rather than only those forming a region of constant depth. Also, because the model accounts for the interaction of several features in producing a sonar return, the correspondence problem is avoided.

4.4.2 Sonar-based localisation

For a sonar mapping system to be successful, it is necessary to integrate sonar information from several different viewpoints. It is therefore also necessary to have a good estimate of the robot's location at all times, which means that localisation is necessary for mapping. In this section we will look at previous research into localisation for a mobile robot using sonar.

Drumheller (1987) described a system of absolute localisation for a robot in an environment consisting only of rough walls, for which a map was given in advance. Drumheller's method could localise the robot from just one, very

dense, scan. The algorithm works by extracting straight line segments from the sonar beam endpoints, and using a tree-searching algorithm to find the best correspondence between walls detected by sonar and walls in the *a priori* map of the environment. The tree to be searched is very large, exponential in the number of walls extracted, but most correspondences are clearly impossible, and tree-pruning methods are used to reduce the size of search. This algorithm was shown to give good results – however it is not clear that it would work in the presence of specular walls, or without dense scans.

Elfes (1987) described a method for localisation based on matching of grid-based occupancy maps. Performing an exhaustive search on the space of all possible robot positions at high resolution is very computationally intensive. To overcome this problem, Elfes proposed to conduct the search on a hierarchy of reduced resolution versions of the map. A map is reduced in resolution by transferring it onto a grid of twice the spacing, in which the occupancy probability of each cell is equal to the maximum of the four corresponding cells on the higher-resolution map. Once a match is found at a low-resolution, a small search is needed to refine it to a match at the next highest resolution, and so on. It was claimed that a map of 3000 cells could be matched to 1 cell resolution in one second of (1987 speed) VAX time – but other than stating the execution speed of the matching process, no results were given.

Another approach to sonar-based localisation has been based on the use of extended Kalman filters (Leonard and Durrant-Whyte 1992). In this approach, a model is made of the mechanical response of the robot (the plant model), and of the sensory process (the measurement model). After each movement, the position of the robot is estimated from the plant model, and after each scan, the estimated position of the robot is used to predict the sonar readings. The difference between the predicted and observed readings is then used to correct the estimated position. The plant and measurement models may be non-linear; however in the EKF framework, plant and measurement noise must be modelled as additive and Gaussian. In practice, however, this restriction does not cause problems, and EKF based methods have been successfully

and Durrant-Whyte 1992), mixed indoor environments (Lee and Recce 1997), and even underground mines (Madhavan et al. 1998).

Absolute localisation, where the robot does not have an approximate position estimate, is necessary when a robot is first switched on (unless it is always switched on in the same place) and also to correct for very large odometry errors such as might be caused by collision with an obstacle. Absolute localisation has received little attention since Drumheller's work described above. Recently, however, and since the work of this thesis was underway, a new technique called *metric Markov localisation* has shown great success at both absolute and incremental localisation in a crowded museum environment (Burgard et al. 1998). This was implemented on a state-of-the-art robot with laser rangefinders, and required an accurate map of the environment to be specified in advance. Whether this technique will also work in the case of a simple robot with only cheap sonar rangefinders and a robot-constructed map is not yet clear – but the difficulty of the tasks successfully performed by this robot should not be underestimated.

4.5 The system of Lee and Recce

The system of Lee and Recce (1997) will be referred to very often in this thesis, and serves as a foundation for some of the work described here. We will therefore describe it here in detail.

The main direction of Lee and Recce's work was quantitative evaluation of exploration strategies for mobile robot map building. The relative merits of different exploration strategies is not a topic with which this thesis is particularly concerned. However, in order to evaluate the exploration strategies, they constructed a sonar-based mobile robot system that was capable of mapping, incremental localisation, and goal-directed movement. We will refer greatly to this system.

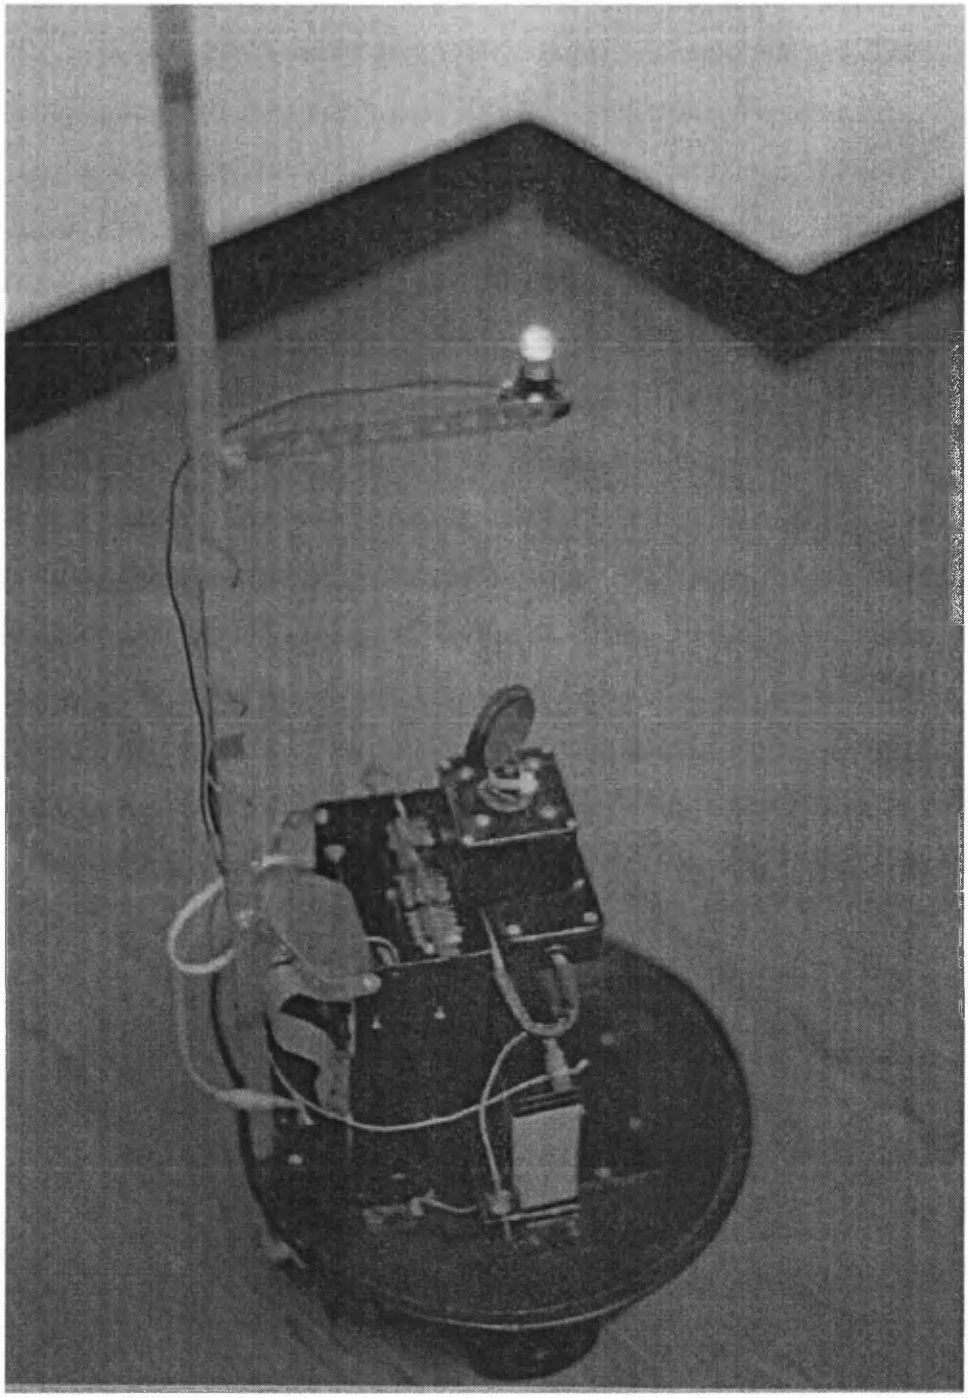

Figure 4.2: ARNE the robot. The light above ARNE is detected using an overhead camera in the water maze simulation of chapter 6. Directly below the light there is a single ultrasonic sensor on a pivot mount, that is used to find the distance to features in the environment. The robot has two wheels which are centred to allow ARNE to move forwards or turn on the spot.

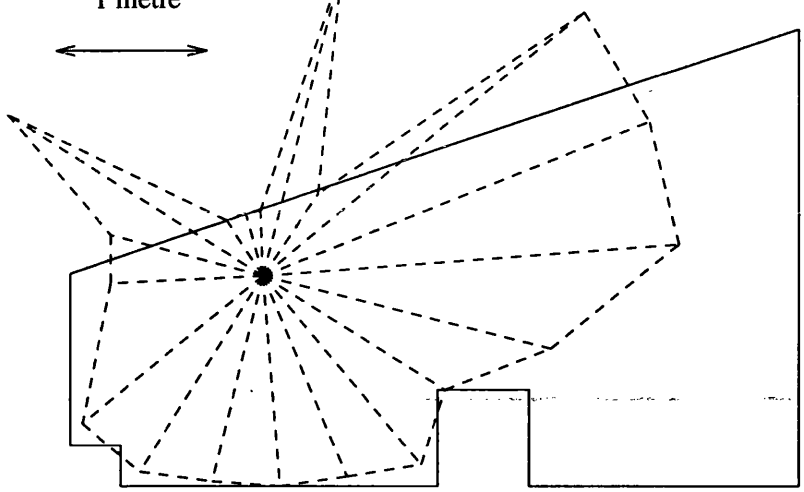

Figure 4.3: Sonar scan of one of the test environments, superimposed on a map of line feature locations. The sonar readings are taken with an angular spacing of 18^o. To guide the eye, a line has been drawn connecting the sonar readings. The reading pointing to the right of the robot is at the maximum detectable range of 253 cm, indicating that no return was detected.

4.5.1 ARNE the robot

A picture of the robot ARNE (Autonomous Robotic Navigator and Explorer) is shown in figure 4.2. ARNE has a 25 cm diameter circular base, and can make two types of movement: a specified distance forward or backward, or rotation by a specified angle. ARNE is equipped with shaft encoders to provide odometric feedback from its wheels, and provide a measure of the distance moved and the rotation angle of a turn. There is a single sonar sensor on a motorised pivot that is used to measure the positions of objects in its environment (range 253 cm, resolution 1 cm).

4.5.2 Sonar preprocessing

After every movement, the robot performs a sonar scan consisting of 20 sonar returns (we define a sonar *return* to be a single, simple range reading). Figure 4.3 shows the results of a single sonar scan within an environment. As discussed above, the sonar-measured distances do not always correspond to distance to the nearest obstacle due to the possibility of specular reflections, and the (approx. 30^o) beam divergence.

Following Leonard and Durrant-Whyte (1992), incoming sonar scans are first grouped to extract "regions of constant depth", groups of angularly neighbouring returns whose range measurements differ by no more than 3 cm. We refer to such a groups of sonar returns as sonar *readings*. Grouping the sonar returns in this way helps to reduce the uncertainty caused by the width of the sonar beam.

The RCDs extracted in this way are used to construct a map consisting of two parts, a feature-based map and a grid-based free-space map.

4.5.3 The feature map

The feature map contains two types of two-dimensional features: line features, which correspond to planar walls; and point features, which correspond to convex or concave corners. The data stored for each feature consists of a list of sonar reflection points that have been ascribed to the feature, and a geometric description of the feature's position. The feature's geometric position is calculated from the list of sonar reflection points as follows: for a point feature, the point position is the barycentre of the reflection points. For a line feature, the reflection points are fit to an infinite straight line by linear regression. The end points of the line feature are calculated by projecting the reflection points onto the infinite line, and choosing the extremal two.

After a movement, a sonar scan is made, and the raw sonar returns are grouped into readings as described above. Every sonar reading is then examined to see if any of the existing features in the map could have caused the reading. For every feature in the map, the hypothesis is tested that the feature could be responsible for the reading. A *contact point* is obtained, corresponding to the place of reflection of the sonar beam under this hypothesis. Because sonar range information can be trusted more than angular information, the distance from the robot to the contact point is the sonar measured distance, but the bearing from the robot's forward direction to the contact point is calculated from the position of the feature. For a line feature, the bearing to the contact point is the direction normal to the line, as it is assumed that the

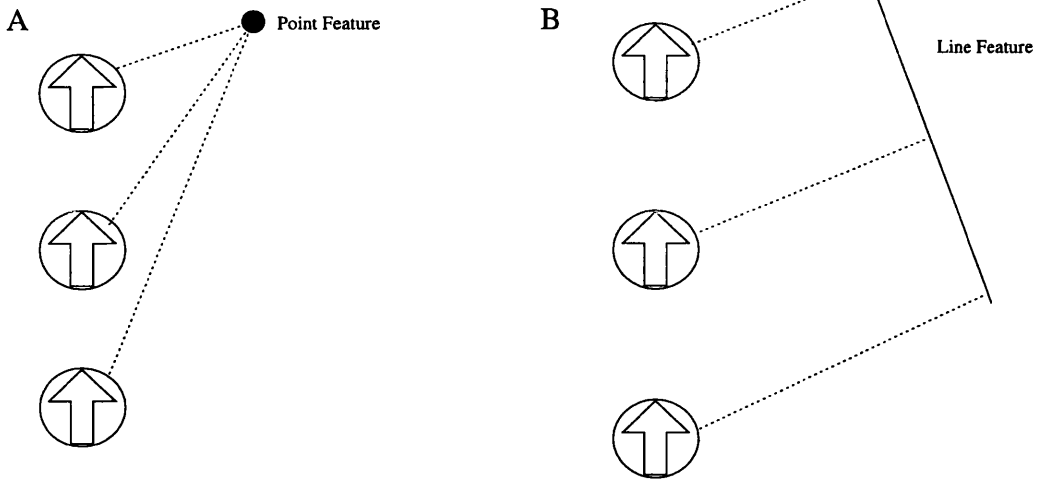

Figure 4.4: Construction of new features from sonar readings. When three sonar readings are taken from consecutive viewpoints which are all consistent with the presence of a new feature, the feature is added to the map. The consistency criterion is the "circle test" of Leonard and Durrant-Whyte. A) for a point feature, the sonar readings must all intersect at a common point. B) For a line feature, the readings must all be reflected perpendicularly to one line. The circle and arrow represents ARNE.

sonar beam was reflected perpendicularly off the wall. For a point feature, the bearing is the direction towards the point feature.

Four criteria are used to determine whether this reflection hypothesis is feasible: 1) Is the object observable from the robot's current position? 2) Is the object within the effective beam width of the reading? 3) Does range of the reading approximately match the distance to the object? 4) (For line features only) The hypothesis is rejected if the contact point lies too far from the end-points of the line feature.

For a given sonar reading, if only one feature exists in the map that satisfies the above criteria, the sonar reading is ascribed to the feature, and the feature's geometric description is recalculated. If more than one feature satisfies the criterion, the reading is ignored. If no features satisfy the criterion, the reading considered for the creation of a new feature.

A new feature is created when 3 sonar readings taken from consecutive viewpoints are all consistent with the existence of the new feature, as outlined in figure 4.4. The consistency criterion that is used is based on the *circle test*

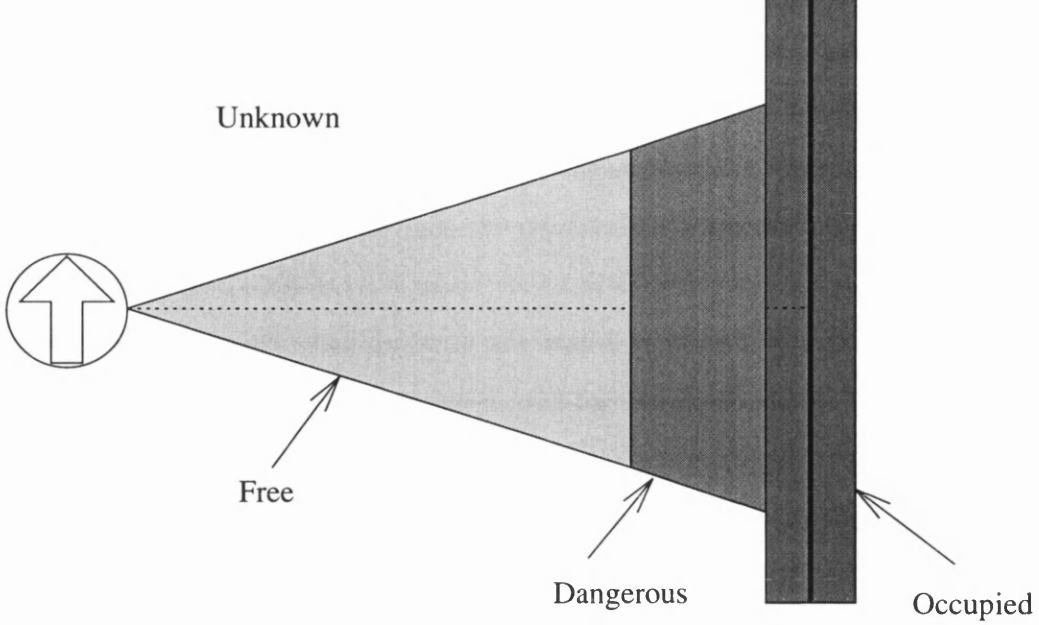

Figure 4.5: All cells on the free-space map are initially marked as unknown. When a feature is added to the map, all cells which the feature lies in are marked as occupied. When a sonar reading is taken that is ascribed to a new or existing feature, cells that the beam passed through are marked as free. Cells that would be marked as free, but are close enough to a feature that ARNE would risk a collision are marked as dangerous.

of Leonard and Durrant-Whyte (1992) (see Lee 1996 for further details).

4.5.4 The free-space map

The free-space map based on a 10 cm square grid. Each grid segment is marked as: *occupied*, if the segment contains a point or line feature; *free* if ARNE has travelled through the segment or a sonar beam has passed through the segment and been reflected from a known feature; *dangerous*, if the segment would be marked free but a point or line feature, but not goal, is close enough to the segment that ARNE would risk a collision by travelling through the cell; or *unknown* if no information is known about the segment. We shall say a segment is *marked* if it carries any mark other than unknown.

Initially every grid segment is labelled unknown. As ARNE passes through space, those cells which it has passed through are marked as free. When a feature is added to the map, the cells in which the feature falls are marked as

are marked as free. Only sonar beams that were used to detect the presence of a feature are used, to avoid the possibility of erroneously marking cells free due to specular reflections. When a new feature is added to the map, all free cells within one ARNE width of the feature are marked as dangerous. The marking of cells on the free-space map is summarised in figure 4.5.

When ARNE is required to go to a certain point in space, the shortest path through cells marked as free is calculated by a recursive search method first described by Jarvis and Byrne (see e.g. McKerrow 1991).

4.5.5 Localisation

The system of Lee and Recce also has an incremental localisation function which corrects for small odometry errors. After every movement, an extended Kalman filter is used to estimate the robots position and bearing. The filter uses two models. A "plant model" uses odometry feedback to predict the robots new position and bearing, together with associated uncertainties. A "measurement model" uses the locations of features in the robot's map to predict the results of a sonar scan, together with associated uncertainties. These two models, together with the actual sonar readings are used to derive an optimal estimate of the robots position and bearing. Full details of the models and Kalman filter technique can be found in Lee (1996).

4.5.6 Map quality metric

In order to evaluate the maps produced by the robot, Lee and Recce defined a quantitative metric. This metric compares the map derived by the robot to an ideal map that has been entered by the operator. The metric measures the fraction of a set of test journeys that the robot would complete successfully, if it used its current map to plan a trajectory.

The test journeys are derived from the ideal map as follows: A 300mm square grid is overlaid on the environment, and a path is planned between every pair of grid points using the ideal map. If a path can be successfully

followed between the two points, the journey is taken as a test journey.

To evaluate each test journey, a path is planned on the basis of the robots current map using the recursive search method mentioned in chapter 6. If such a path can be found, and this path would not lead the robot into a cell which, on the ideal map, is marked occupied, dangerous, or unknown, then the journey is marked as successful. The map quality is defined to be the fraction of all test journeys marked as successful.

4.6 Summary

Modern mobile robot science, inspired by the behaviour-based approach of Brooks, plays down the need for explicit world models, and places an emphasis on the solution of simple sensory-motor problems before moving onto more complex problems that might require "intelligence".

In the language of chapter 1, the lowest level problems of chemotaxis, stimulus-response behaviour, and path-integration have now been well solved by simple sensory-motor or behaviour-based techniques. Localisation and mapping, the next stages up in the hierarchy, have therefore received a great deal of attention recently.

In localisation, good progress has been made in incremental localisation, but for absolute localisation there is still much work to be done. In mapping (with sonar), the two current methods of choice both have problems. Grid-based mapping has difficulty detecting specular walls, and is based on incorrect mathematical assumptions. Feature-based mapping techniques can detect specular walls, but are based on arbitrary heuristics, and therefore contain many free and possibly non-optimal parameters. Furthermore, the heuristics used in current feature-based techniques require a large fraction of the sonar data to be thrown away, rather than being used for map construction.

The research of this thesis will describe work aimed at overcoming these problems. The following chapters will describe: a new approach to absolute localisation derived from a neurobiological theory for the function of the hippocampus; a new approach to grid-based mapping which will work in a

rect mathematical assumptions; and a new approach to feature-based mapping which uses all data, and is derived from mathematics rather than heuristics.

Part II

Chapter 5

A framework for spatial function in the mammalian brain

This chapter will describe a framework for the performance of spatial functions by the mammalian brain, paying particular attention to the role of the hippocampus. The following chapter will describe how a navigational system based on this framework was implemented on a mobile robot, and experiments conducted to test the implementation.

Why do we need another model of hippocampal involvement in spatial function? The cognitive map theory provides a well defined theory of this. What could be added to it?

The problem is that the cognitive map theory, as it stands, is incompatible with a role for the hippocampus in general memory. Unless we believe that the evidence for a memory function of the hippocampus is erroneous, there must be a single neuronal architecture that is capable of both a general memory function and a spatial function.

In this chapter we propose that an autoassociative memory is one such architecture, and we describe how the addition of an autoassociative memory to an egocentric map making system in the neocortex gives an animal absolute localisation ability. In the next chapter we will describe how this framework was instantiated on a mobile robot. In this chapter we will describe the framework at a qualitative level. We will discuss the results of previous biological

tions of the framework by which it may be distinguished from the original cognitive map theory and other models of spatial function in animals.

5.1 Egocentric and allocentric maps

Before fully describing the framework, it is worth clearly defining what we mean by a map. We will define a map to be a list, with each list entry describing the nature and location of an environmental feature. An egocentric map is one where the locations of the features are represented in egocentric coordinates, and an allocentric map is one where the locations of the features are represented in allocentric coordinates.

If an agent is to use a map to get to a goal location, it must produce a movement vector to the goal. The vector through which the agent must move to get to the goal will be equal to the goal's position vector in egocentric coordinates. If the agent has an egocentric map of space, it can read the required vector directly off the map. However, if the agent has an allocentric map, it will also need a representation of its own position in allocentric coordinates, and must derive the required movement vector by vector subtraction.

In order to produce useful behaviour in an environment, a complete allocentric map offers few advantages over a complete egocentric map. The only time an allocentric coordinate system provides extra functionality is when a goal is externally specified in allocentric coordinates. But if a goal location is determined by sensory input, which is the case we will consider, an allocentric coordinate system provides no advantage.

5.2 Role of the hippocampus: localisation as a form of pattern completion

Localisation, as it is normally defined, means using the observed egocentric locations of a few environmental features to produce an estimate of position in an allocentric coordinate system. This definition cannot be applied to an

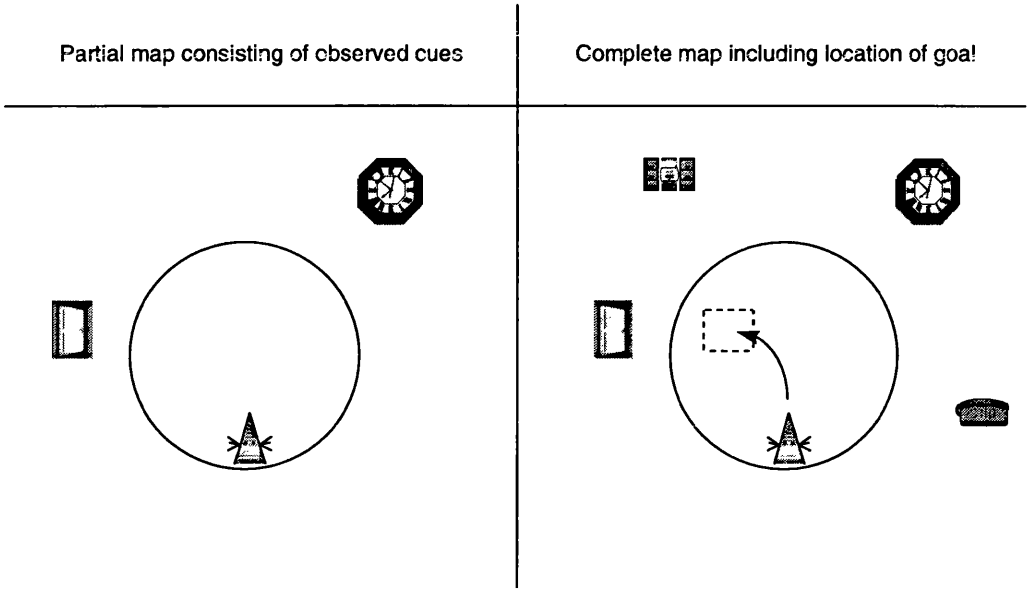

Figure 5.1: Localisation as pattern completion in the water maze task. The animal can observe a number of cues, which form a partial map (left). The hippocampus then completes the partial map to a full map (right), which also contains the location of the hidden platform.

agent that does not use allocentric coordinates. But, as we have argued above, to perform useful sensory-driven behaviour an allocentric coordinate system offers little advantage. We will therefore extend the definition of localisation to agents that do not use allocentric coordinates, and redefine it as the ability to produce a complete egocentric map from the observed egocentric positions of a small number of features.

To make this idea concrete, let us consider the classic test for localisation ability in animals, the Morris water maze (Morris 1981). In this test, the animal can perceive the locations of various cues but not the location of the submerged platform. To deduce the location of the platform, the animal must take the egocentric positions of the observable cues, and produce an egocentric vector to the hidden platform. In other words, the animal must take a partial egocentric map, which contains only the observable cues, and produce

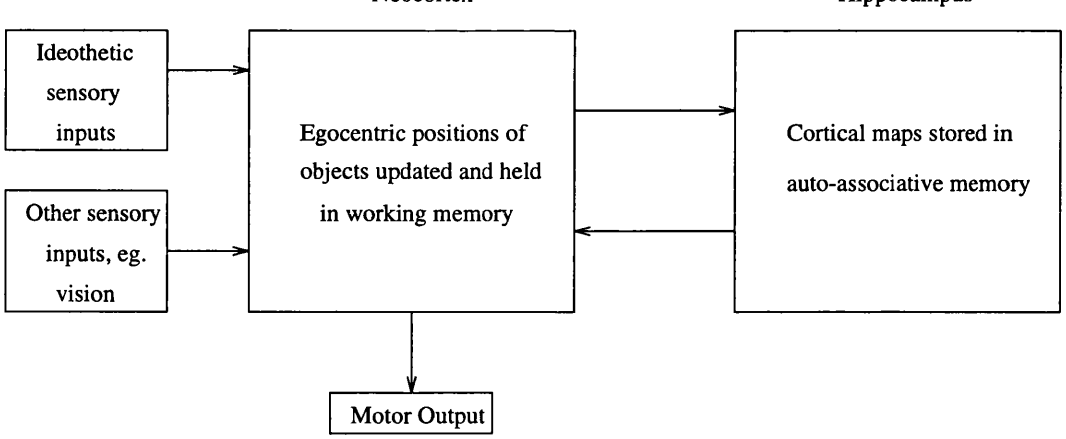

Figure 5.2: The neocortex processes sensory data to construct an egocentric map of the environment, and produces motor output using this map. The hippocampus is an autoassociative memory which performs pattern completion on partial maps constructed by the neocortex, and so performs a localisation function.

a complete map, which also contains the location of the hidden platform. This is illustrated in figure 5.1.

The problem of localisation can therefore be recast as one of pattern completion. Now we can make contact with neurobiology. According to Marr's theory, pattern completion is precisely the function of the hippocampus. We propose that the hippocampus functions as an autoassociative memory performing a pattern completion function, and that this plays a role in spatial function by giving the animal absolute localisation ability. Figure 5.2 summarises this model. In the following section, we will describe in detail the spatial functions required of neocortex, that, together with an autoassociative memory in the hippocampus, make up our framework for spatial function in the brain.

5.3 Spatial functions performed by the neocortex

In our model, the neocortex is responsible for five separate spatial functions (described below) which together allow it to represent an egocentric map of the environment.

of the neocortex is *redundancy compression*, that is, forming a compact neural representation of external stimuli by taking into account statistical regularities in sensory input. This idea is supported by cortical anatomy: in general, sensory input is processed through several successive cortical stages prior to reaching the hippocampus (Swanson 1983).

In a spatial context, redundancy compression could be achieved by classification of cues and determination of their location (*mapping*, in the language of chapter 1). In our framework, the neocortical activity pattern represents the locations of environmental features, including goals, in egocentric coordinates. We will refer to this set of cue locations as the *cortical egocentric map.*

Working memory. Working memory is often defined as a type of memory that is relevant only for a short period of time, usually on the scale of seconds. Working memory is also characterised as having a limited capacity, and being liable to interference if the subject is distracted (Baddeley 1986). It has been suggested that the prefrontal cortex is involved in working memory (see Goldman-Rakic 1990 for a review).

Our model requires the neocortex to store the nature and location of cues in working memory, so that they remain in the egocentric map even if the animal is not directly attending to them.

Path integration. We suggest that, in order to track the change in egocentric position of a cue due to locomotion, the neocortex uses motor efference copy and ideothetic input to update the positions of remembered cues in an egocentric coordinate frame. We propose that the positions of all the cues are updated simultaneously and in parallel. The animal is not updating its own position in an allocentric coordinate frame because it does not use an allocentric coordinate frame.

The idea that the egocentric positions of all cues are path integrated in parallel, rather than a single position for the animal, has been criticised re-

up independently in the stored positions of individual cues, leading to a loss of coherence of the map. This would be the case if the principal source of error was internal to the brain. However, it is much more likely that the principal source of path-integration errors is the process of determining the distance moved from ideothetic input and motor efference copy. In this case, the path-integration errors for different cues will not be independent but will be equal, so map coherency is not affected.

The proposed mechanism of simultaneous updating of egocentric cue position is similar to a model suggested for the monkey oculomotor system, where it has been proposed that cells coding for the retinal location of stimuli can be updated using efference copy of saccadic eye movement commands (Droulez and Berthoz 1991). In this model, an image is represented on a layer of cells in retinal coordinates. When an eye movement is made, the entire activation pattern of this cell layer is shifted to compensate for the eye movement, so the image stays in retinal coordinates. This proposal is supported by evidence that the receptive fields of some monkey parietal neurons shift in advance of eye movement (Duhamel et al. 1992).

Coordinate transformations. It has often been suggested that the parietal cortex is responsible for transforming sensory information into new coordinate systems. For example, recordings from neurons in the posterior parietal cortex of monkeys have identified neurons with retinotopic visual receptive fields, and with gain modulated by eye and head position (see Andersen 1995 for a review). It has been proposed that these neurons form a distributed representation for the spatial location of stimuli, from which location information can be extracted in a number of coordinate systems (Zipser and Andersen 1988). Different tasks presumably require particular associated coordinate systems, and changes in the coordinate system of sensory stimuli may be a key function of many regions of the brain.

Touretzky and Redish (1995) suggest that, upstream from the hippocam-

angle coded for by the head direction system. As the head direction system codes for the animals head direction relative to a fixed, but arbitrary, bearing, the map arriving at the hippocampal input synapses will always be oriented the same way with respect to this bearing, regardless of the animal's own head direction. As a result, place cell firing will be dependent on position, but not head direction, as seen in open field environments (Muller et al. 1994).

We have incorporated this role for the head-direction system in our model. We propose that the brain uses similar mechanisms to perform this rotation as it uses to perform other coordinate transformations, such as the conversion of visual information from retinal to head centred coordinates.

Planning and execution of movements. The final spatial function we require of the neocortex is the planning and execution of movements. Modelling of the neural substrates of (spatial or non-spatial) behaviour is an immensely difficult task, but one of the few things known for sure is that, in mammals, the neocortex plays an important role. Animals express a vast range of spatial behaviours including exploration, homing, foraging, hiding, and of course remaining stationary. To make a detailed model of these would be well beyond the scope of this thesis, and also unnecessary for evaluating the model presented here. We will make a rudimentary model of animal behaviour, to be described in the next chapter, but we are not proposing that it in any way approximates the complexity of behaviours seen in real animals.

5.4 Relation to previous experiments

We have described a framework for the performance of spatial functions by the brain in which the neocortex is responsible for processing sensory data to construct and maintain an egocentric map of the animal's environment, and the hippocampus is an autoassociative memory which performs pattern completion of partial egocentric maps, and thereby gives the animal absolute localisation ability.

We will now look at some previous experiments, considering how they might be explained under the current framework.

5.4.1 The water maze

The water maze (Morris 1981) is often cited as *the* spatial task which defines the spatial role of the hippocampus. In the language of chapter 1, the water maze is a test of absolute localisation ability.

We have already discussed performance of the water maze task in relation to the current framework in section 5.2. To reiterate: when the rat is introduced to the maze, it observes the positions of some distal cues, and this triggers the recall of a map stored during previous exploration. This recalled map contains the location of the hidden platform, as well as the locations of the distal cues, and the rat is then able to head directly to the hidden platform (see figure 5.1).

5.4.2 Short-cut ability

Tolman (1948) emphasised the ability of animals to cross parts of space that they have not previously explored in order to reach a goal location. Tolman's emphasis is best understood in a historical context: at the time this paper was published, ideas of behavioural psychology were dominant, with an strong emphasis on stimulus-response learning. The fact that rats can cross previously unexplored areas of space indicates that they are not using a simple stimulus-response system, and therefore dealt a serious blow to stimulus-response psychology.

Short-cut ability poses no problem for the current framework, because it is not a simple stimulus-response system. If the neocortex contains a representation of a goal location, whether or not that goal is currently visible to the animal, the animal will be able to head towards the goal, using the egocentric mechanisms of the neocortex.

In our framework, the hippocampus is *not* necessary for short-cut ability unless the task also involves localisation. In the experiment of Mittelstaedt

to a nest along the most direct route, even if they have not travelled this exact route before. This requires short-cut ability, but not localisation, and would therefore not be hippocampus-dependent under the current framework. However, a task that requires localisation as well as short-cut ability, such as the "sun-burst maze" of Tolman et al. (1946), would require the hippocampus.

5.4.3 Latent learning

Another piece of evidence used in support of cognitive mapping was the fact that rats can perform latent learning, i.e. they can learn about an environment simply by exploring it, without rewards or punishments.

Tolman (1948) describes an experiment using a Y-shaped maze, which had the end of one arm baited with food and the other baited with water. On the initial exploration, the rats were neither hungry or thirsty, so did not eat or drink, and consequently were not rewarded. On a subsequent probe trials, half the animals were made hungry and the other half made thirsty. The hungry animals went directly to the arm baited with food, and the thirsty ones to the arm baited with water. This was taken as evidence that animals can form maps, rather than merely learning to follow routes which lead to reward.

In our framework, this would be explained as follows. During the initial exploration session, the rats would construct a map of the whole environment, including the food and water. When a rat is then reintroduced to the environment, the features visible from the start location will trigger pattern completion of the full map constructed in previous exploration, and the rat's neocortical egocentric map will then contain a representation of the food and water locations. The rat can then use neocortical methods of path planning to move directly to the appropriate goal location.

The key point is that, in our framework, goals are not treated differently to other types of environmental feature. In the stimulus-response models that Tolman was arguing against, goals acted as reinforcers of behaviour, and so had a very different representation to other environmental features. The idea of

involvement in spatial function (Brown and Sharp 1995), and it is not clear that these models alone can explain the latent learning phenomenon.

5.4.4 Entorhinal place cells

Quirk et al. (1992) showed that place cells exist in medial entorhinal cortex. However, the firing patterns of entorhinal place cells are in some ways different to those of hippocampal place cells.

The most striking findings came when a rat was moved from a circular environment to a visually similar square environment. Both environments had plain grey walls and a single white cue card. The firing patterns of hippocampal place cells in the two environments showed very little relation: a cell with a firing field in one of the environments may be virtually silent in the other, and if a cell does have place fields in both environments, their locations are unrelated. Entorhinal place cells, on the other hand, showed similar firing properties in the two environments. The mean firing rate of a given cell was approximately equal in both environments, and the firing rate maps of a given cell were "topologically related", meaning that preferred firing locations of the cell in the two environments were in approximately the same position relative to the cue card.

Quirk *et al.* concluded that the positional firing of entorhinal cells was more "sensory bound" than that of hippocampal cells, and that the ability to discriminate environments, while present in the hippocampus, is not yet present in medial entorhinal cortex.

In order to consider these results in the framework of this chapter, we must think about the way the autoassociative memory function is performed by the hippocampus. The current theory does not attempt to explain this, but previous authors have provided detailed analyses of how it might be done (Marr 1971; Treves and Rolls 1994). A feature common to most hippocampal models of autoassociative memory is *orthogonalisation.* Autoassociative memories perform poorly when required to store patterns with a high degree of overlap,

thogonal. Marr proposes that the first stage of hippocampal processing (which occurs in the dentate gyrus) is devoted to setting up a "simple representation" of the input, in which the patterns corresponding to different inputs are close to orthogonal.

Now let us consider the results of Quirk *et. al.*. The sensory input at a location in the square environment will be similar, but not identical, to the sensory input at the topologically related location of the cylindrical environment. If the firing of entorhinal cells codes for sensory features, then the firing patterns seen in entorhinal cortex for the two locations will have a high, but not complete degree of overlap, and so entorhinal place cell firing will be similar in the two locations. However, the hippocampus will transform this overlapping representation to a nearly orthogonal representation. Therefore hippocampal firing patterns will be very different in the two environments, and topological transformation will not be seen.

5.4.5 The effect of environmental manipulations of place cell firing

Many experiments have investigated the effect of environmental manipulations on the place cell firing pattern. It is an unsurprising fact that, if all environmental cues are moved coherently, then the firing fields of place cells translate and rotate to match the environmental manipulation (Muller and Kubie 1987; O'Keefe and Speakman 1987; Knierim et al. 1995). In the current framework this can be explained simply: the firing of neocortical cells is based upon sensory input, and represents the egocentric configuration of landmarks. It will therefore follow any coherent movement of environmental cues. The hippocampal activity pattern is a function of the neocortical activity pattern, and will therefore follow suit.

If the majority of landmarks are left in the same place, but one or two landmarks are removed, the hippocampal place code is largely unaffected (Muller and Kubie 1987; O'Keefe and Speakman 1987). This may be interpreted in

egocentric map, without the missing cue, will trigger the recall of the full map from the hippocampus, and therefore the same hippocampal firing pattern will be activated. The hippocampal output will, of course, be in conflict with the cortical map, because the missing feature will be in the recalled map but will not be observed. In such a case, we expect the neocortex to give preference to recent sensory input above hippocampal recall based on previous exploration. We shall see in the next chapter that such a mechanism is necessary to successfully instantiate this framework on a robot.

In other experiments (Quirk et al. 1990), a rat was placed in a familiar environment and, after a delay, the lights were switched out. Hippocampal place cells continued to fire as normal, although their place fields drifted slowly with time. This may be explained within our framework using the working memory function of the neocortex. When the lights are switched out, the neocortex still maintains an estimate of the egocentric locations of cues in working memory. This estimate is updated by ideothetic information, and so can remain relatively accurate in the dark. Because the hippocampal firing pattern is a function of the neocortical map, the hippocampal place code will continue to reflect an estimate of the animals current position. In the dark, though, localisation is not possible, and the cortical egocentric map may become inaccurate due to path integration errors. These path integration errors will affect the estimated positions of all cues by the same amount, and so the entire egocentric map will shift coherently. Thus, hippocampal place cells will continue to fire, but their firing positions will also drift coherently with the egocentric map.

Further experiments have examined the effects of moving environmental features while the rat is actually performing a task. Gothard et al. (1996) trained rats to run between a movable box and a fixed goal on a linear track. On initial recording sessions, the box remained fixed in one position, but in later probe trials, the box was moved during the rats outward run (when it could not see the box being moved). During the initial part of a run, both

firing was determined by the distance from the starting point of the run. But towards the end of a run, place cell firing was determined by distance to the destination. This may be interpreted in the current framework as follows. In the initial part of a journey, the rat leaves the start location with a neocortical representation of the start location in egocentric coordinates. This causes place cells to fire that were trained in the initial sessions when the rat was at the same distance from the start location. As the rat nears its destination, the goal comes into view (remember that rats have very limited visual acuity). At this point, the neocortex contains a representation of the goal position in egocentric coordinates, which activates place cells trained at the corresponding distance from the destination in the initial sessions. The recall produced by these place cells contradicts the path-integrated egocentric representation of the start location. However, as the start position has not been observed for a long time, and so might anyway be afflicted with path integration errors, the recall produced by more recent observations will take precedence.

5.4.6 Non-spatial determinants of place cell firing

Finally, several experiments show that, in certain situations, "place" cell firing is modulated by non-spatial factors, such as odour (Eichenbaum et al. 1986), current behaviour (Wiener et al. 1989; Markus et al. 1995), and the presence of local cues (Young et al. 1994). These experiments have been interpreted as evidence against the cognitive map theory, and other theories of exclusively spatial hippocampal function.

In the current framework, the hippocampus functions as an autoassociative memory which stores snapshots of the current cortical firing pattern. This pattern includes, but is not limited to, the cortical egocentric map. We therefore expect hippocampal cell firing to be affected by any factor that influences the firing pattern of those cortical cells that project to hippocampus. Such factors could certainly include odours, current behaviour, and the presence of local cues.

The framework makes two strong experimental predictions by which it can be distinguished from the original cognitive map theory.

5.5.1 The effect of lesions on path integration ability

In our framework, it is the egocentric position of landmarks that is updated by ideothetic input, rather than the allocentric position of the animal. We propose that the neocortex is responsible for this path integration. This implies that path integration ability should be more affected by lesions of the appropriate parts of neocortex, than by lesions of the hippocampus. Since this prediction was originally published (Recce and Harris 1996), preliminary evidence confirming it has been presented (Alyan et al. 1997).

5.5.2 Instantaneous transfer

Our framework makes predictions for the behaviour of animals in unexplored parts of familiar environments. According to the model, if the animal arrives at a novel location by its own motion, or is taken there in such a way that it can carry out passive path integration, it will still have a valid egocentric map, and will be able to follow an accurate trajectory to a goal. By contrast, if the animal is disoriented before being placed in the novel location, it will not have a valid egocentric map, and therefore would not be able to follow an accurate trajectory to a goal. The model therefore predicts that animals are not capable of instantaneous transfer.

5.6 Summary

In this chapter we have proposed, at a qualitative level, a framework for spatial function in the brain. In this proposed scheme, the neocortex is responsible for building and representing a map of space in egocentric coordinates, and the hippocampus is responsible for performing pattern completion on partial maps produced by the neocortex, giving the animal absolute localisation abil-

autoassociative memory network.

Many authors have previously argued (see e.g. Marr 1971; Treves and Rolls 1994) that the hippocampus must function as an autoassociative memory. But these arguments were based on the evidence for a general memory role for the hippocampus, not a spatial role. The framework described here describes how an autoassociative memory may also play a role in spatial function, and therefore allows these models to be unified with the evidence for a spatial role for the hippocampus.

The framework described here is consistent with previous biological experiments, and also makes two predictions for future experiments. One of these predictions has since been confirmed by preliminary evidence.

Chapter 6

Instantiation of the framework on a mobile robot

The last chapter described a framework for spatial function in the brain. This chapter will describe how the theory was implemented and tested with a mobile robot.

The principal test we used was to see if the robot could perform a task analogous to the Morris (1981) water maze. In order to test the theory we had to implement software that performs the functions that, in the last chapter, we ascribed to the neocortex, software that performs the autoassociative memory function ascribed to the hippocampus, and a rudimentary system for controlling the robot's behaviour.

In order to mimic the invisible goal of the water maze task, we placed a video camera above a certain place in the environment. An incandescent light bulb was put on top of the robot's central axis (as seen in figure 4.2). The robot's presence under the camera was automatically detected by a simple thresholding of the maximum brightness of the camera reading. Because it was mounted high on a wall, the camera was not detectable by sonar, and the location of the goal was only available to the robot when it passed under the camera.

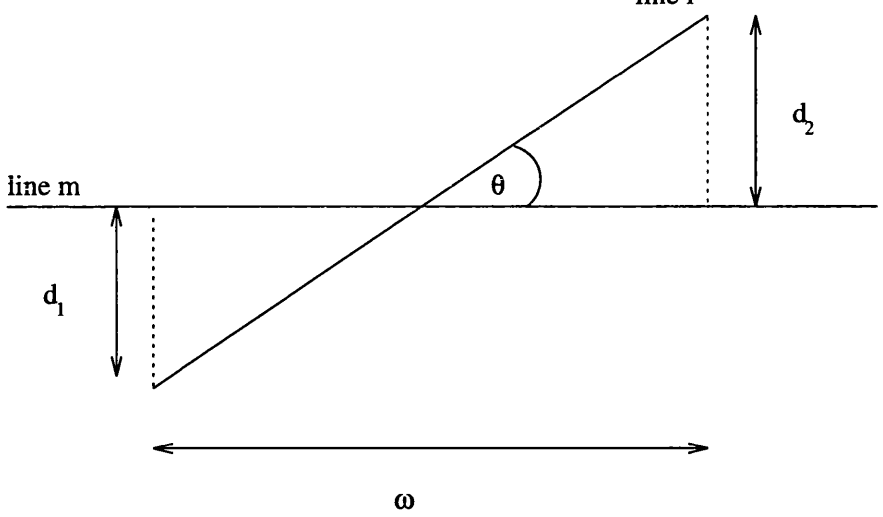

Figure 6.1: Geometrical definition of the function `overlap(l, m)`. ω is the fraction of line m which is overlapped by the projection of line l. d_1 and d_2 are the lengths of the orthogonal lines of projection on the endpoints of line l onto line m. θ is the angle between the lines. `overlap(l, m)` = $\omega e^{-\theta^2/C^2-(d_1+d_2)^2/D^2}$. For the experiments described here, C was 10^o, and D was 200 mm.

6.1 The simulated neocortex

This section describes software that performs the spatial functions that were ascribed in the last chapter to the neocortex. The software is largely adapted from code previously written by David Lee, which is described in section 4.5. As the function of this software is to build a map from sonar data and odometry, we will also refer to it as the *map-maker*. The old system performs nearly all the functions we require of the modelled neocortex. However, a few modifications were required, which are described below.

6.1.1 Line merging

When a new line feature is created, or an old one is modified, all other lines are checked to see if they could correspond to the same wall in the robots environment. If so, the two lines are merged to create a new line feature whose sonar reading list contains the sonar readings of both the original lines. The merging criterion uses the function `overlap`, defined in figure 6.1, which measures the fraction of the old line which is overlapped by the projection of

lines or a large separation of their endpoints. If this function exceeds a fixed threshold (0.3 for the simulations described below), the two lines are merged. The merging process is recursive; when a new line is created by merging two existing lines, the same check is carried out on the newly created line.

6.1.2 Goal representation

A goal is represented in the map by a simple pair of Cartesian coordinates. The goal position corresponds to the position of the overhead camera described above. The goal is only detected when ARNE passes under the camera. When this happens, ARNE's presence is automatically detected, and a signal is sent to the map-making software which sets the goal position to ARNE's current location.

6.1.3 Coordinate conversion

In the last chapter, we proposed that the neocortex represents maps in egocentric coordinates. Unfortunately, the software developed by David Lee functions in allocentric coordinates. Instead of rewriting this software to function in egocentric coordinates, we simply added a module that translates the map it generates into egocentric coordinates before passing to the simulated hippocampus. This results in an equivalent system, but a large saving in development time.

6.2 The simulated hippocampus

This section describes the addition of an autoassociative memory unit, simulating the hippocampus, to the simulated neocortex of the previous section. A diagram of the full system is shown in figure 6.2.

At every time step, the current map stored in the map-maker is translated to robot-centred coordinates, rotated by an angle corresponding to the robot's heading direction relative to a fixed allocentric bearing (how this angle is calculated is described below), and fed to the simulated hippocampus. When

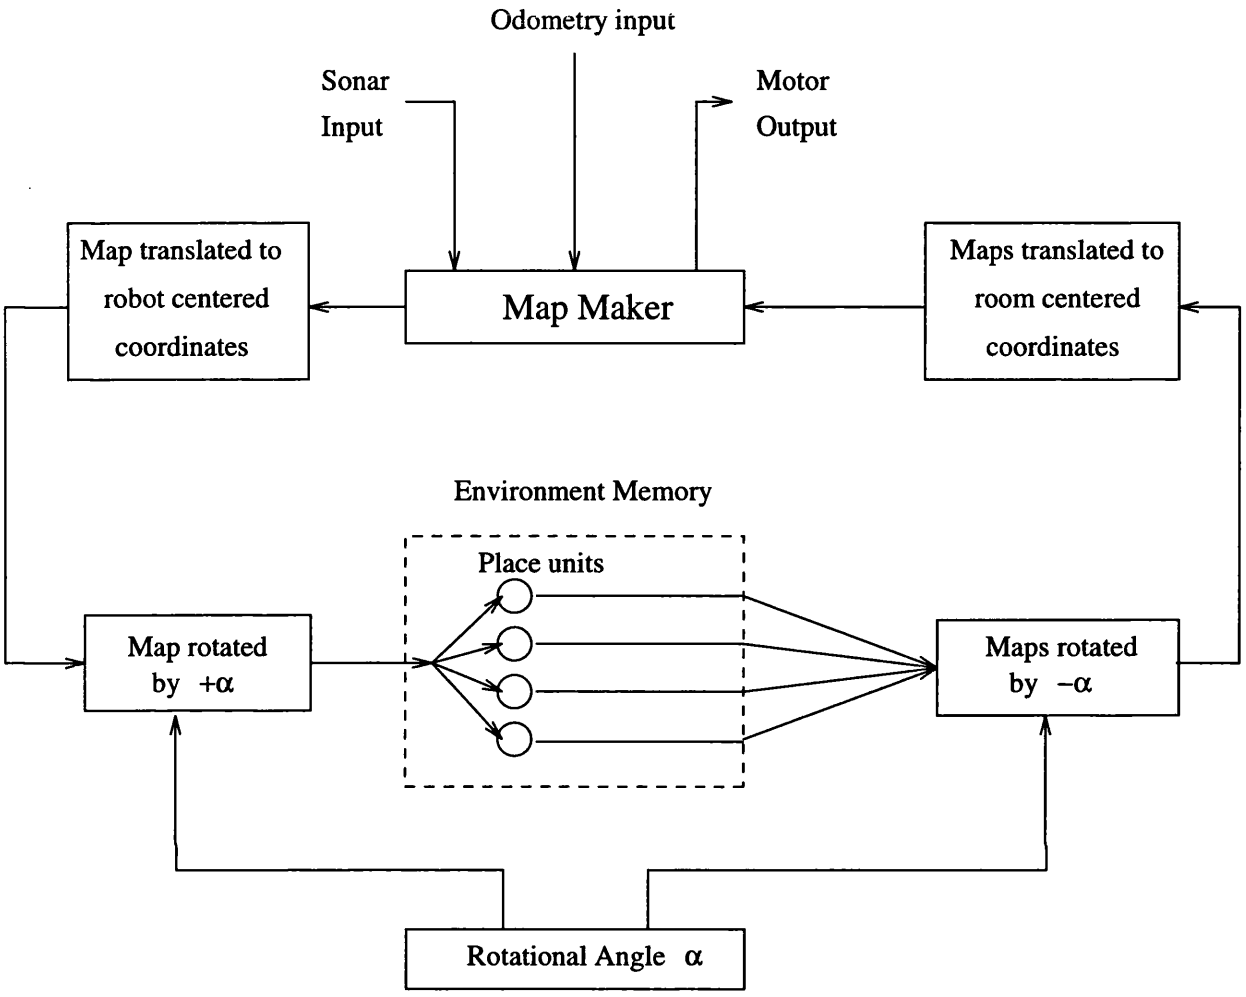

Figure 6.2: ARNE's control system. The map-maker processes sonar and odometry input to build up a map of the environment. Partial maps are translated to robot centred coordinates, rotated by a "head-direction" angle (described in section 6.3), and fed to the simulated hippocampus. In a previously explored environment, a fragmentary map passed to the simulated hippocampus will trigger the recall of a full map stored in the previous exploration session. This full map is then merged with the partial map currently stored by the map-maker.

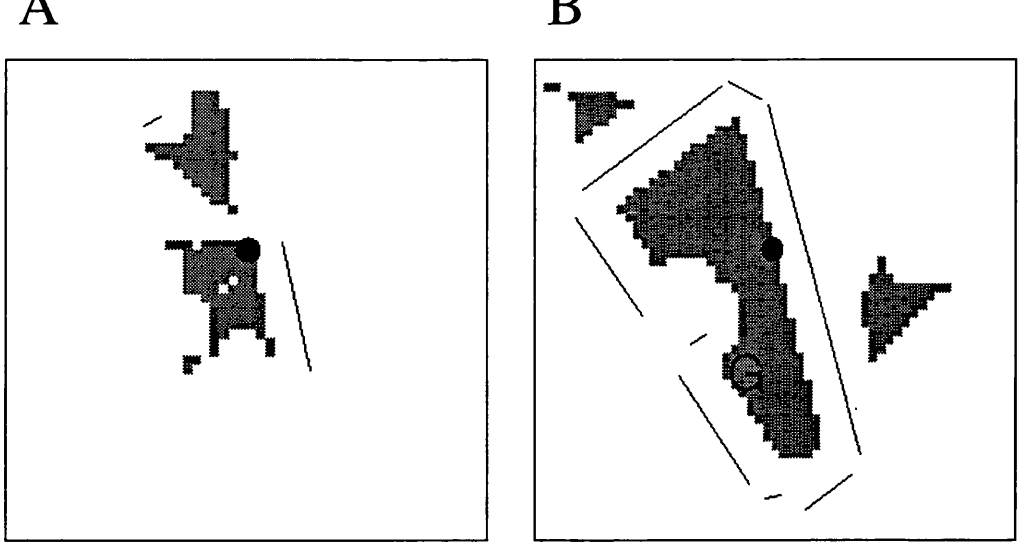

Figure 6.3: An example of spatial pattern completion. The filled circle represents ARNE, solid lines represent line features, the grey area represents free-space, and the letter "G" represents the hidden goal. A) A fragmentary map built up after 5 time steps in a previously explored environment. B) A map returned by the simulated hippocampus in response to the partial map shown in A. This map includes a representation of the goal, and so ARNE can now move directly to the goal.

the robot is exploring a new environment, the simulated hippocampus stores maps it receives as input. When the robot is replaced in a previously explored environment, a fragmentary input map triggers the recall of a full map stored during the previous exploration, and the simulated hippocampus returns the completed map to the map-maker.

As described in the last chapter, the coupling of an autoassociative memory to an egocentric mapping system leads to absolute localisation ability. When the robot is placed in an environment, it starts with an empty map. Due to the limited range of the sonar sensor and the fact that multiple observations of a feature are required to overcome sensory inaccuracy, it takes several movements of the robot before a map can begin to be built. Figure 6.3a shows a fragmentary map built by the simulated neocortex five timesteps after ARNE was introduced to a previously explored environment. Figure 6.3b shows a map returned by the simulated hippocampus in response to this fragmentary input. Since the completed map contains the location of the hidden goal,

At no point does the simulated hippocampus produce a pair of Cartesian coordinates that represent the absolute position of the robot. However, in terms of producing behaviour, all that is required is that the robot know the location of features in the environment relative to itself. As argued in the last chapter, production of a pair of absolute Cartesian coordinates is unnecessary.

6.2.1 Simulated place cells

The simulated hippocampus consists of several independent *place units.* A place unit is a virtual neuron which either fires or is silent at every time step. Each place unit receives as input a map in robot-centred coordinates, and can store a map in robot-centred coordinates. This map is stored in the same symbolic representation used by the map-maker. Each place unit can store one map, which consists of the input free-space map and line features. Point features are not stored as they did not improve performance. Place units exist in two states: *naive* or *trained.* Initially all place units are naive, and their stored maps are empty. When a map is stored in the environment memory, it is copied into one of the naive place units, and that place unit is marked as trained. Once a place unit has been trained, its map is never altered.

At every time step, for every place unit, a score is calculated to determine the closeness of fit of the place unit's input map to the place unit's stored map. This score corresponds roughly to the fraction of the input map which is also a subset of the stored map. More precisely, the score consists of an average of two scores, one for the free-space map and one for the feature map. The free-space score is the fraction of those marked grid segments in the input map, for which the stored map carries a the same mark (free, dangerous, or occupied). This score can be at most one, with equality when the input map is an exact subset of the stored map. The feature map score uses the function `overlap` defined in figure 6.1. The function `overlap(l, m)` is very close to zero unless the lines are approximately collinear, in which case it is the fraction of line `m` that is "overlapped" by line `l`. The feature map score is the sum for all lines

the number of lines in the input map. This score is then approximately the fraction of the total length of lines in the input map that are overlapped by lines in the stored map. It is approximately bounded above by 1, with equality when all of the lines in the input map are subsegments of lines in the stored map.

A trained place unit fires if its score exceeds a fixed threshold, which was 0.6 for the simulations described below. There is no limit on the number of place units that fire, and the place units do not inhibit (or activate) each other. When a place unit fires, the map it stores is rotated and translated back to room-centred coordinates, and returned to the map-maker. The elements of the returned maps are merged into the current map. In this merging process, a general principle is employed that the robot should trust map elements observed directly by sonar more than elements recalled from the simulated hippocampus. In practice this is accomplished by having every map element (grid segment mark or line feature) in the map-maker carry a further mark of *confirmed* for those observed directly or *provisional* for those recalled from the simulated hippocampus. If an element is returned by the simulated hippocampus for which there is already a confirmed element in the current map (i.e. for a returned grid segment mark, if the corresponding grid segment in the current map already carries a confirmed mark; or for a line feature if there is a confirmed line feature that satisfies the line merging criterion), the returned element is ignored. Also, if a sonar reading is taken which contradicts the existence of a provisional element (i.e. a new mark is allocated to a grid segment which carries a provisional mark; or a sonar reading cuts through a provisional line feature), the provisional element is deleted. The map passed as input to the simulated hippocampus consists only of the confirmed elements in the map-maker's map; provisional elements are not passed.

A few more issues must be dealt with to make a system that can solve a water maze style task:

1) In order to efficiently solve the task, it is necessary for the robot to perform different behaviours in different contexts. When the robot is introduced to the environment, it should first find out if the environment has been previously explored. If it has, the robot should head directly to the goal. If it has not, the robot should explore. Similar behaviour is observed in rats, who will actually ignore a food reward in order to explore a novel environment.

2) If a map is input to a place unit which matches the stored map well, but has a different absolute orientation, the similarity score will be small. However, the map-maker always constructs maps oriented relative to the robots forward direction when first placed in the environment. In order that previously explored environments be recognised for any initial orientation, the map from the map-maker must be rotated before being passed to the simulated hippocampus, so that all maps of one environment have the same absolute orientation. It does not matter what this absolute orientation is, and it may differ from environment to environment, as long as it is constant within one environment. The robot should determine the rotational angle automatically, without user input.

3) In early stages of exploration, the number of elements in the map will be quite small. A small map consisting of, say, one line feature will be a fragment of many stored maps, not just those of the place unit which corresponds to the robot's actual position. This therefore leads to many erroneously recalled maps. Furthermore, if such small maps are stored in place units, they will only hinder the system later. The problem can be solved by inhibiting the simulated hippocampus until the current map is large enough.

These three issues are all tackled together by introducing *behavioural modes.* At any time, the robot is in one of three behavioural modes. Before describing the modes, we must describe two measures of map quality that are used to

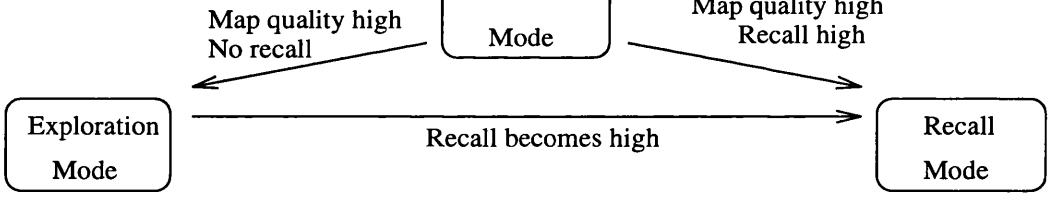

Figure 6.4: Transitions between ARNE's three behavioural modes make use of the measures of map quality and recall quality defined in the text. When ARNE is first placed in an environment, it is in orientation mode. If map quality becomes high but recall quality is low, this indicates an unexplored environment, and ARNE goes into exploration mode. On the other hand, if recall quality is high, this indicates a previously explored environment, and ARNE goes into recall mode. If ARNE is in exploration mode and recall quality becomes high, this indicates that the environment has been sufficiently explored, and ARNE switches to recall mode.

determine which mode the robot is in.

The *recall quality* is a running average of the number of place cells firing. It is therefore high in an environment which the robot has previously explored. The *map quality* is a running average of the fraction of sonar readings which are ascribed to existing features (see section 4.5.3). This measure will be high when most of the walls in the environment are represented by features in the robot's map. The running average a_t of a series x_t is defined here by $a_t = 0.7a_{t-1} + 0.3x_t$.

When the robot is first placed in an environment, it is in *orientation mode.* In this mode, its objective is to determine whether the environment is new or has been previously explored. In this mode, the robot always moves according to a simple wall-following exploration strategy (Lee and Recce 1997). Recall and storage in the simulated hippocampus is inhibited in this mode, to avoid problem 3 above. However, the number of place cells firing is still calculated, as this is needed to determine the recall quality measure. In orientation mode, the angle by which maps should be rotated before being passed to the map-maker is unknown. At every time step, all angles are searched (step size 1/32 of a revolution), and the angle is chosen for which the maximum number of place units fire. In the event that the environment had been explored before, this

initial orientations between the first exploration of the environment and this exposure; in the event that the environment had not been explored before, this would produce an arbitrary angle.

When the map quality is high enough, the robot leaves orientation mode and transfers to *exploration mode* if recall quality is low, or *recall mode* if recall quality is high. Precisely, if the map quality is at least 0.4, there are at least 2 line features in the current map, and recall quality is at most 1.0, the robot goes to exploration mode. If map quality is at least 0.2, there are at least 2 line features in the current map, and recall quality is at least 1.0, the robot goes to recall mode.

In exploration mode, the robot again uses a simple wall-following exploration strategy, and trains a new place unit at each new position. In this mode, the rotational angle is fixed, in order that the maps stored in place units all have the same orientation. If recall quality becomes high (greater than 3.0), the robot transfers to recall mode.

In recall mode, the robot's behaviour depends on the status of the goal representation in its current map. If the map contains a goal which is marked as provisional (i.e. it has been recalled by the environment memory, but has not been directly observed by the robot passing under the camera), it moves directly to the goal, following the shortest path through free-space as calculated by recursive search. If the map contains no goal, or if the robot has already passed under the camera, the robot follows a wall-following strategy. In recall mode, the rotational angle is again fixed.

The conditions for transitions between the different modes are shown in figure 6.4.

6.4 Experimental results

Figure 6.5 shows ARNE's performance at the water maze style task. The initial exploration of ARNE in the environment lasted for 120 time steps, and followed the trajectory shown in figure 6.5A. ARNE was then replaced at six

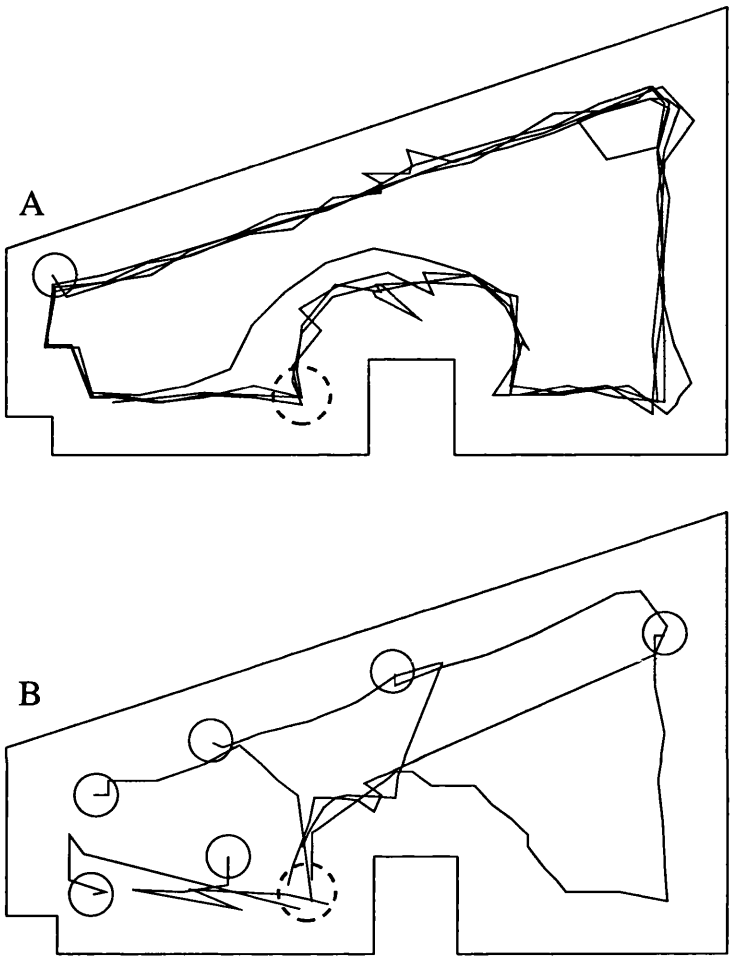

Figure 6.5: The trajectories followed by ARNE in the water maze style task. A) Trajectory in the training phase. B) Trajectories in the testing phase. Unbroken circles represent ARNE's starting positions; the dotted circle represents the goal location.

different test locations. The trajectories followed by the robot are shown in figure 6.5B. In each case, ARNE successfully reached the goal. The mean number of steps between introduction and reaching the goal was 14.7, with standard deviation 7.0.

In order to asses the increased functionality the simulated hippocampus gives the robot in an environment without a goal, we used a previously defined measure of map quality (Lee and Recce 1997). This measure is the fraction of all possible journeys that could be made in an environment which would be made successfully on the basis of the robot's current map. The measure is a real number in the range 0 to 1, with 1 denoting a perfect map. To calculate this measure requires a user specified map of the environment. As such maps had already been compiled in previous work, we assessed the robots performance using trace files of robot movements and sonar readings previously collected during 10 runs in the "columns" environment of Lee and Recce (1997). These trace files consist of all the data passing from the robot to the map-maker during the experiment, and are enough to completely reproduce the original run. We fed the trace files into the full robot software, and a version without the simulated hippocampus unit. Figure 6.6a shows map quality as a function of time for the last run in the environment. The map quality is consistently higher with the simulated hippocampus present. Figure 6.6b shows the mean and standard deviation difference in map quality between the two systems versus time, averaged over all runs. In the early stages of exploration, the map quality of the full system is significantly higher. This indicates that the simulated hippocampus had successfully recalled the environment, before the map-maker had a chance to build up a complete map. In later stages of exploration, there is no significant difference between the quality of the two systems.

In order to compare the firing properties of the place units of our system to real hippocampal neurons, and in order to see if the robot could distinguish different environments, we performed another test, in which we directly examined the firing properties of simulated place units. For this test, we used

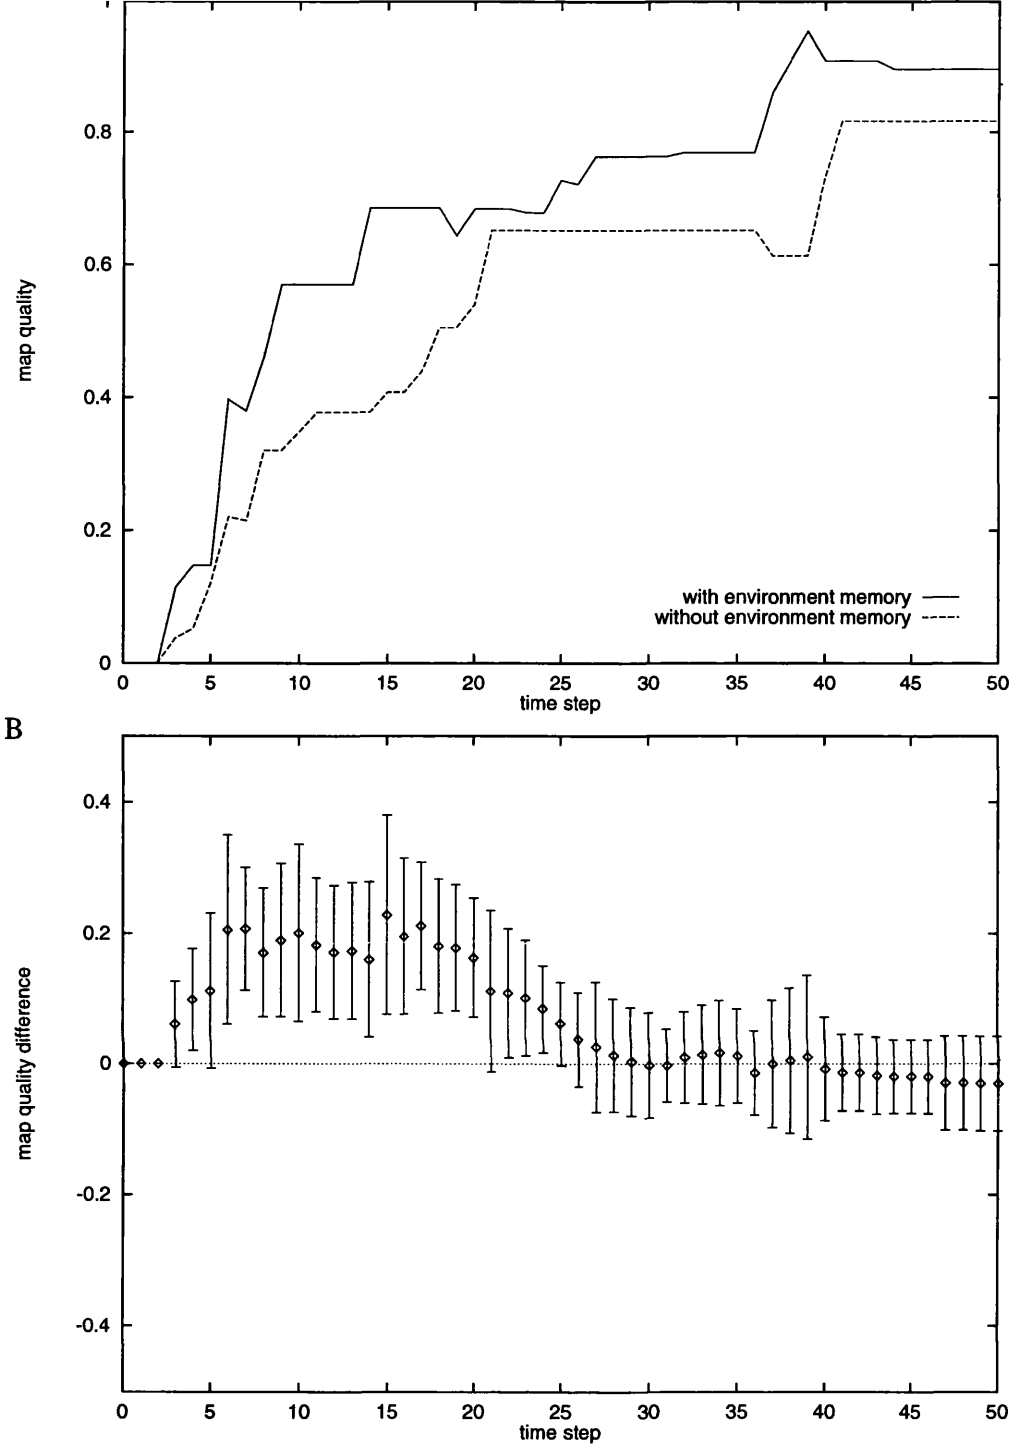

B

Figure 6.6: Assessments of map quality with and without absolute localisation. Map quality is assessed by the map quality metric of Lee and Recce (1997), for trace files of 10 consecutive runs in an environment without a goal. The map-maker was reset before each run, and each run began with a new starting position and orientation. A) Map quality on the last run in the environment. The system with the simulated hippocampus present shows consistently better performance. B) Mean and standard deviation of difference between map qualities with and without the map-maker. The difference is significant in the early stages of exploration.

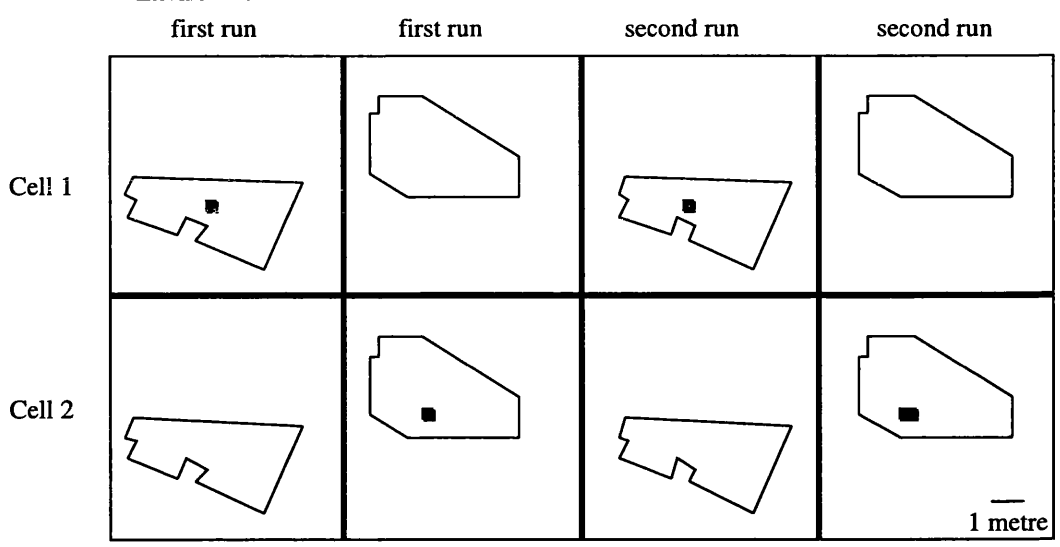

Figure 6.7: Firing rate maps for two simulated place units. Each row shows the firing rate maps of a single place unit for each of the two runs in the two environments. The first run in each environment lasted 120 steps and the second run lasted 80 steps. The firing rate maps are computed on a grid as the ratio of the number of times the neuron fired in one grid location over the time spent in that location.

two environments, neither of which contained a goal. ARNE was placed in environment A and was allowed to explore for 120 timesteps. ARNE was then removed and placed in environment B for 120 timesteps. ARNE was then reintroduced to environment A for 80 timesteps, and then to environment B for 80 timesteps. Figure 6.7 shows firing rate maps for two example place units. A total of 186 place units were trained during the four runs, of which 94 fired only in the first environment, 90 fired only in the second environment, and 2 fired in both environments. 107 of the place units were active in both visits to their preferred environment. All but 2 of the place units shown coded for a single location in an environment.

6.5 Discussion

The system described here has enabled a robot with a simple sonar sensor to recognise its location in an environment after a small number of exposures to the environment. This has been demonstrated by: 1) requiring ARNE to

early stages of exploration, and 3) directly observing the activity patterns of simulated place units.

The localisation system described here is different to other localisation systems described so far, in that a separate map is stored for every place the robot has visited, and the robot's location is found by comparing its current map to the maps stored for each of these places. In other approaches, for example (Drumheller 1987), a single map is stored, and the robots position on the map is determined using a search procedure. Obviously, storing many maps consumes more memory than storing a single map, but for the results described here, map storage still required less than 1 megabyte.

The activity pattern of the modelled place units has many of the reported properties of place cells. For example the activity pattern is determined by the entire constellation of cues in an environment, rather that one single cue. Also, both real and simulated place cells have environment-specific activity.

In the system of Lee and Recce, the map-maker represents a map in allocentric coordinates. In the current system, this map must then be translated to egocentric coordinates before being passed to the simulated hippocampus. In the brain however, we believe that the cerebral cortex performs the functions of the map-maker in egocentric coordinates. This avoids the need to translate coordinate systems between the cortex and the hippocampus, and also allows egocentric sensory information to be added directly to the map, and egocentric motor output to be directly computed.

The model describes the calculation of a rotation angle, between the direction that the robot is moving, and the orientation of the stored maps. Since all of the stored maps for an environment have a common orientation, once found, this rotation angle mimics the instantaneous orientation direction of the robot. This rotation angle therefore has similar properties to the head-direction cells of the post-subiculum.

Even though the system was tested with a robot that used only a single sonar sensor, the place unit approach is applicable to any sensory system. All

that is needed is a measure of similarity for egocentric maps. In the case of a vision system, having a separate map stored for every location potentially avoids recognition complications that occur with objects which look different from different viewing angles, but which must be held in the same way in a single map. The inclusion of additional sensory systems should lead to faster map construction, and more rapid recognition of a familiar environment.

6.5.1 Implications for neurobiology

In the previous chapter we described a theory for performance of spatial functions in the mammalian brain. In this chapter we described a robotic system that was based on this theory. What has the process of implementation taught us about the original biological theory?

The main thing that the simulation has shown is that the theory is plausible – i.e. that the system described is actually capable of performing the task. It is very easy to underestimate the difficulty of tasks that animals perform apparently effortlessly. Computer simulation helps avoid this problem by forcing the modeller to clearly define the task and model. The use of a robot further restricts the model by making it deal with noisy, real-world data.

Many models of spatial function in animals, including some computer simulations, assume that cues can be identified and localised with complete accuracy. By contrast, with the sonar data available to ARNE, features are often badly localised or missed altogether. This has constrained the model in a number of ways. For example, some models of spatial function in animals assume that the cue location input to the hippocampus consists only of distances to landmarks, and not bearings. In these models, non-directional place fields can therefore be formed by a simple competitive learning process. We originally tried such an approach but found that, due to the unreliable localisation of features, it was also necessary to use angular information. However, this meant that directional place-fields would be formed in open-field environments, which is not the case in rats. So a solution was found using the rotational angle mechanism described in section 6.3. This rotational angle,

to the angle coded for by the post-subicular head-direction cells.

Another feature that was not originally planned, but was forced upon us by robot implementation was the necessity of treating features which have been directly observed from sensory data differently to those that have been recalled from hippocampus. In early implementations in which all features were treated equally, "positive feedback" cycles would occur, in which an erroneous map recall early in exploration would lead to non-existent features being present in the cortical map, which would in turn lead to further recalls and more erroneous features. This was avoided by marking those features recalled from the hippocampus as "provisional", and not sending them back to the hippocampus as input. It is tempting to speculate that a similar system of differentiating directly observed features from recalled features exists in the brain.

One feature of the current model we do not believe corresponds to neurobiology is the simulation of the neocortex. To simulate the spatial functions performed by the neocortex, we used some previously developed software which was not originally intended as a biological model, but nonetheless needed only minor modifications to perform the functions we required. But although it performed the functions we ascribed to the neocortex, the mechanisms by which these functions were performed are very different to any we believe may be employed by the brain. For example, the system uses a symbolic representation for the positions of features, something that is very unlikely to be the case in a biological system.

Another feature of the current implementation is that, because maps are built up incrementally in the neocortex, the set of maps stored in the hippocampus will be nested, i.e. maps stored earlier will be submaps of maps stored later. This is not an optimal use of storage capacity, and furthermore storage of small map fragments can actually impair later performance. While the behavioural modes described above reduce the number of small map fragments stored by preventing map storage during orientation mode, we

subpatterns. One possibility is that this occurs during sleep.

Part III

Chapter 7

A neural network system for grid-based sonar mapping

From now on, we will turn our attention to the problem of map construction. This chapter describes a technique of grid-based map construction in which the grid segments are modelled by integrate-and-fire neurons, and have lateral connections styled after models of visual cortex.

7.1 Introduction

For a robot with a sonar sensor, two methods of map representation have been in common use up to now. In *grid-based representation* (Elfes 1987), the two-dimensional space of the robot's environment is overlaid with a square grid. For each grid segment, a set of numbers is stored representing sensory information about the contents of that grid segment in physical space. In the simplest cases, known as *certainty grid* methods (Moravec 1988; Cho 1990), a single number is stored representing the probability that the segment is occupied. In *feature-based representation* (Leonard and Durrant-Whyte 1992) the map consists of a list of physical features, such as walls and pillars, and the positions of these features are represented by Cartesian coordinates.

Grid-based representation has the advantage that a grid map of free space is useful for path planning. Furthermore, the probabilistic nature of grid maps allows them to represent uncertain information. In a feature-based map, on the other hand, features are only added to the map when their existence is

it takes many sonar readings to confirm the existence of a feature, but every sonar reading can be used to update the probabilities stored in a grid-based map.

A severe limitation of most certainty grid methods is their inability to detect specular walls. One might guess that the reason for this is the type of map representation used. With specular walls, the range returned from the wall depends heavily on the angle between the sonar beam and the wall orientation. However, in most grid-based systems, the orientations of walls are not represented.

Lim and Cho (1992) have proposed an extension of the certainty grid technique for use in specular environments. The central feature of this method is a change in map representation. In addition to storing an occupancy probability for each grid segment, the also store a set of numbers corresponding to the probability that a wall passing through the grid segment has each of a fixed number of orientations. They then describe a modified sensor model in which the probability of a direct return from a grid segment is equal to the probability that the segment contains a wall oriented normal to the sonar beam. A second key feature of Lim and Cho's sensor model is the use of a *range confidence factor*, which incorporates the fact that long range readings are more likely to be specular than short ones. Using the modified sensor model, Bayes' theorem is employed to derive an update rule for the estimated occupancy and orientation probabilities for each grid segment given new sensor data.

7.2 Certainty grids and neural networks

Cells of a certainty grid function similarly to neurons in an artificial neural network. In the certainty grid method, a single number is stored for each grid segment, corresponding to the probability that the segment is occupied. When a sonar reading suggests the segment is occupied, this number is increased, and when a sonar reading suggests the segment is free, this number is decreased. If the number exceeds a threshold, the segment is estimated to be occupied,

as a formal integrate-and-fire neuron, with membrane potential corresponding to the probability of occupancy, and firing occurring when the segment is estimated to be occupied. Making the analogy to neural networks allows us to use concepts such as receptive fields and lateral connections, that are familiar from the neural network and neuroscience literature.

The derivation of the rules for updating the probabilities stored in a certainty grid makes the assumption that the occupancy probabilities of all cells are independent. This assumption, also made in Lim and Cho's system, is in general not justified. Walls and empty spaces in the environment are usually continuous, and, in man-made environments, walls are usually straight. A simple grid-based approach does not make use of this prior information about the environment. In this chapter we present an extension of grid-based mapping in which we do not assume the occupancy of grid segments is independent. Lateral connections between the neurons for neighbouring segments reinforce patterns likely to be found in the environment, such as straight walls, and inhibit patterns that are unlikely to be found, such as random scatter.

7.3 The network

7.3.1 Choice of coordinate system

A robot can represent a map of space in either egocentric coordinates, meaning that the positions of features are expressed relative to the robot, or in allocentric coordinates, meaning that the positions of features and of the robot are expressed in a fixed, world-centred frame. In this chapter we use egocentric coordinates, inspired by the neurobiological theory of chapter 5.

7.3.2 Network architecture

The space surrounding the robot is divided into a square grid of spacing 10cm, and for each grid segment there are several neurons, coding for different types of features in the environment. The neurons are therefore topographically

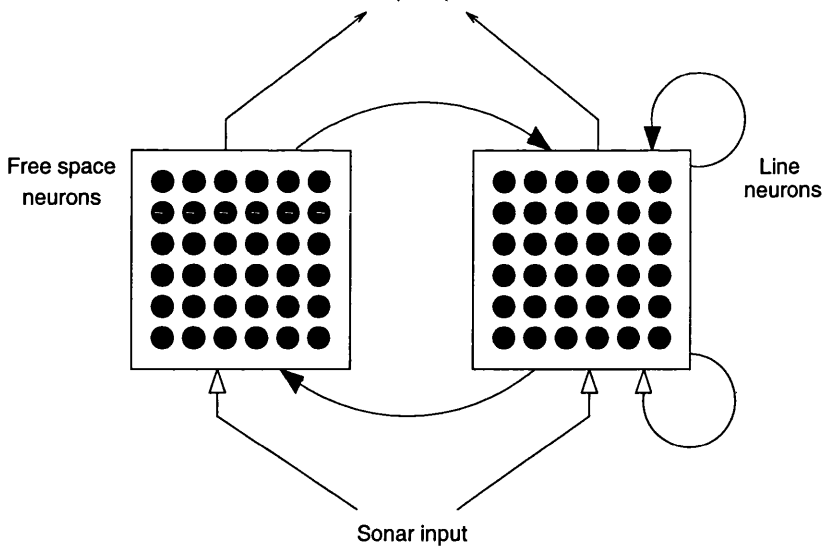

Figure 7.1: The network architecture. Open arrows indicate excitatory connections, filled arrows represent inhibitory connections. The network consists of two sets of topographically arranged neurons, the free-space neurons and the line neurons. Both sets receive excitatory feed-forward input from sonar, and there are excitatory and inhibitory recurrent connections between the neurons.

arranged, with each neuron representing the presence of a preset feature type in a preset region of egocentric space. When we refer to the "position" of a neuron, we mean the position of its underlying grid segment in egocentric space.

Figure 7.1 shows a schematic diagram of the network described here. The network consists of two types of neurons. For each grid segment there is a free-space neuron, and a set of line neurons. The firing of a free-space neuron signals that the corresponding region of egocentric space is unoccupied. For each grid segment there are also 16 line neurons, parametrised by a preferred orientation which ranges from 0^o to 180^o. The firing of a line neuron signals that a straight wall of the correct orientation passes through the grid segment. The encoding of the orientation of features in the environment by these neurons is similar to edge coding cells that have been found in the visual system of animals, except that in this model they are encoding the egocentric location of the walls in the environment, instead of the retinotopic location of edges in a visual scene.

All neurons are of integrate and fire type, with a membrane potential x_i, binary output y_i, and feed-forward input z_i obeying the following equations:

$$x_i(t+1) = x_i(t) + z_i(t) + \sum_j w_{ij} y_j(t)$$
$$y_i(t) = \Theta(x_i(t) - 1.0)$$

The weights w_{ij} of the recurrent connections are fixed, and do not learn (their values are given in section 7.3.5). Map information is stored in the activation pattern of the neurons, rather than in the synaptic weights. The integrate and fire type neurons allow sonar input to be accumulated over several time steps, and therefore from several different robot positions. Note that in this model the integrated membrane potential does not passively decay. All neurons have a fixed threshold of 1.0 and no bias. This threshold is sufficiently high so that on average several inputs are required to make the neuron active.

7.3.4 Sensory input

The external input z_i of a neuron is calculated as follows. The neuron receives an excitatory contribution for every sonar return, and the contribution from all the sonar returns are added to produce the total input for the neuron. The contribution for a given sonar return is a function of the "sonar reflection point" of the return, defined to be the point the measured distance from the robot at an angle equal to the beam centre direction. As we shall see in chapter 8 this point may *not* be the actual physical location where the reflection occurred – but we will continue to use the term in this chapter.

Example graphs of input contribution against reflection point position are shown in figure 7.2 for various neurons. A given neuron's input is close to zero except in a certain region of space. Following the neuroscience literature, we will call this region the *receptive field* of the neuron.

For a line neuron, the receptive field is elongated, directed along the neurons preferred orientation. The exact formula for the input is a product:

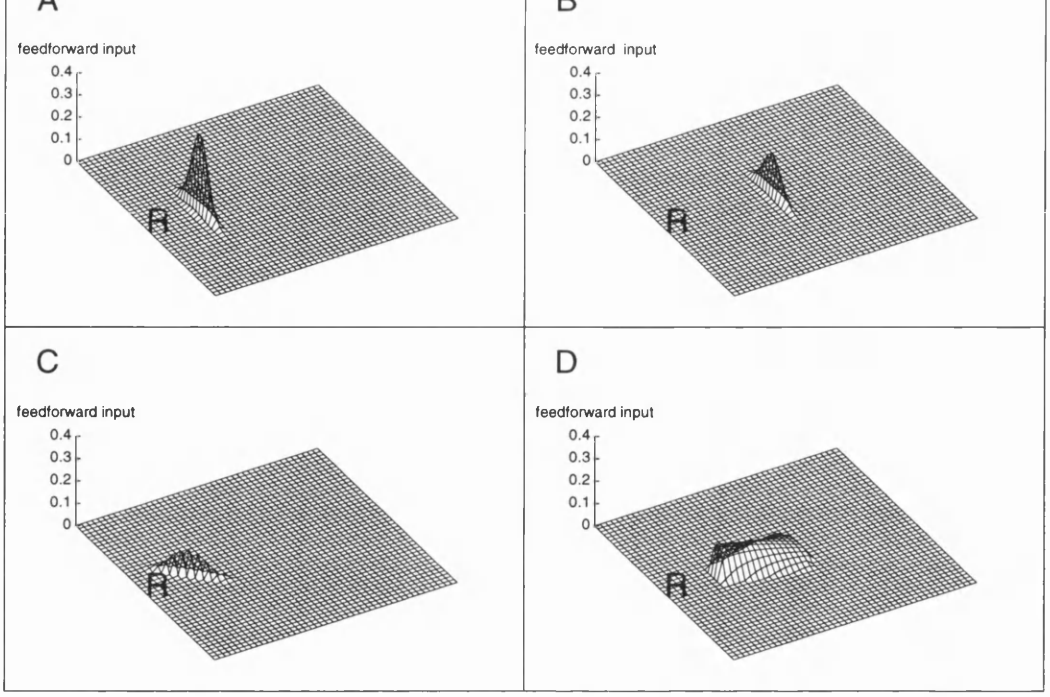

Figure 7.2: Receptive fields of various neurons. The figures show the excitatory input to the neuron due to a single sonar return, as a function of sonar beam reflection position. The letter "R" denotes the position of the robot. The square mesh has a $10cm$ spacing. A) Receptive field of the line neuron for the grid segment $1m$ east of the robot with north-south orientation. B) Receptive field of the line neuron for the grid segment $1.5m$ east of the robot with north-south orientation. C) Receptive field of the line neuron for the grid segment $1m$ east of the robot with northwest-southeast orientation. D) Receptive field of the free-space neuron for the grid segment $1m$ east of the robot.

$$z_{line} = C_{line} \times F_{pos} \times F_{ang} \times F_{RCF}$$

C_{line} is a constant, equal to 0.5 for the simulations described below. F_{pos} is a positional factor which gives the receptive fields their oblong shape. It is given by a double Gaussian:

$$F_{pos} = e^{-\left(d_{par}^2/2\sigma_{par}^2 + d_{perp}^2/2\sigma_{perp}^2\right)}$$

where d_{par} is the distance of the sonar reflection point from the line neuron position in the direction parallel with the neuron's orientation, d_{perp} is the distance of the sonar reflection point from the line neuron position in the

direction perpendicular to the neuron's orientation, and σ_{par} and σ_{perp} are constant distances equal to $300mm$ and $40mm$ respectively.

F_{ang} is an angular factor, to take into account the fact that sonar is preferentially reflected perpendicularly from specular walls by increasing the input of those line segments whose orientation is perpendicular to the reflected sonar beam. It is given by a Gaussian:

$$F_{ang} = e^{-\Delta\theta^2/2\omega^2}$$

where $\Delta\theta$ is the angular difference between the sonar beam angle and a perpendicular to the neurons orientation, and ω is the sonar beam half-width, 15^o.

F_{RCF} is a *range confidence factor* (Lim and Cho 1992), whose purpose is to reduce the contribution of long range readings, which are more likely to be specular. It has the form

$$F_{RCF} = \left(1 - \frac{returning\_range}{max\_detect\_range \times range\_weight}\right)^k$$

where $returning\_range$ is the measured length of the sonar beam, max_detect_range is the maximum measurable length ($2.53m$ for our sensor), and $range\_weight$ and k are constants equal to 1.1 and 0.8 respectively (values given in Lim and Cho 1992).

Figure 7.2a shows the input strength as a function of sonar reflection position, to a neuron from a grid segment $1m$ to the east of the robot with orientation in the north-south direction. Figure 7.2b shows the input to a neuron $1.5m$ to the east of the robot, with the same orientation. The overall strength of input to this neuron is lower, due to the range confidence factor. Figure 7.2c shows the input strength to a neuron from the same grid segment as that in figure 7.2a, but oriented at a 45^o angle. This neuron also has a lower overall input strength than the neuron of figure 7.2a, due to the angular factor.

The length of a line neuron receptive field is much greater than the grid size. The line neurons therefore perform line detection in a similar manner to

reflected from many points along the wall, all lying in a straight line, and will therefore give a large input to a line neuron whose orientation is collinear with the wall.

For a free-space neuron, the receptive field consists the region of space from which the sonar beam would not have been reflected if the the part of space the neuron corresponds to was occupied. This region consists of all points which are further away from the robot than the neuron, and whose bearing from the robot is less than one sonar beam width different from the neurons bearing. The strength of the input is given by:

$$z_{fs} = C_{fs} \times F_{pos} \times F_{RCF}$$

where C_{fs} is a constant equal to 0.2, F_{RCF} is as above, and F_{pos} is a positional factor given by

$$F_{pos} = 1 - \Delta\phi^2/\omega^2$$

where $\Delta\phi$ is the angular difference between the sonar beam and the bearing to the neuron, and ω is the beam half-width as above. An example of a free-space neuron receptive field is shown in figure 7.2d.

7.3.5 Recurrent connections

The recurrent connections of the network are summarised in figure 7.3. A line neuron is connected to all other line neurons lying in a rectangular box the size of its receptive field (size $2\sigma_{par} \times 2\sigma_{perp}$). If the second neuron has the same orientation, the connection is excitatory, with weight 0.01. If the second neuron has a different orientation, it is inhibitory, with weight -0.02. The aim of these connections is to enforce collinearity of line segments corresponding to a single wall in the environment. A line neuron also inhibits with weight -0.3 all free-space neurons lying in its receptive field, and receives an inhibitory connection from them of weight -0.2. The aim of these connections is ensure that specular

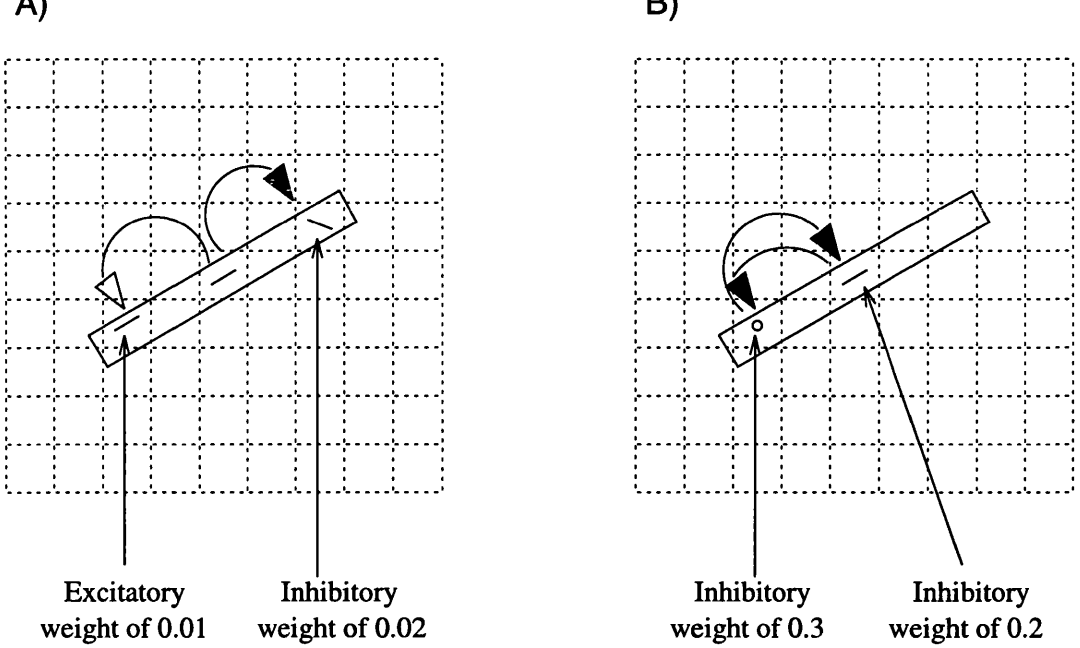

Figure 7.3: This figure illustrates the recurrent connections between neurons. A) A line neuron is connected to all line neurons for grid segments within its receptive field. If the two line neurons have the same orientation, the connection is excitatory, otherwise it is inhibitory. B) A line neuron inhibits all free-space neurons in its receptive field, and is reciprocally inhibited by them.

reflections do not cause line neurons to fire erroneously in unoccupied regions, or free-space neurons to fire erroneously in occupied regions.

7.3.6 Execution cycle

At each time-step, a sonar scan is taken, the feed-forward inputs are calculated, and the membrane potentials of all the neurons are updated. Any neurons which exceed their threshold will fire. The recurrent connections are then activated once, and a map is output on the basis of the neural firing pattern. The robot then makes a movement. After movement, the entire neural activation pattern, including sub-threshold potentials, is shifted through the distance the robot has moved, so it remains in an egocentric coordinate system. A new time-step then begins.

This allows the robot to integrate information from several viewpoints as follows: When the robot is initially placed in an environment all of the neurons

inputs to several neurons in the network. When the robot moves, this potential is transferred to the neurons corresponding to the new egocentric position of the features that caused the initial sub-threshold potential, by the shifting mechanism just described. A second scan receiving reflections from the same features will activate these same neurons, allowing sonar information from several viewpoints to be integrated.

7.4 Experimental Verification

In order to asses the performance of the above mapping system in comparison to other methods, we also implemented a Bayesian grid-based mapping system. The grid-based methods of Elfes (1987), Moravec (1988), and Cho (1990) performed poorly on sonar data collected from our robot due to specular reflection from smooth wood walls in the environment. We therefore implemented the method of Lim and Cho (1992), which was designed to overcome the problems of specular reflection. For the parameter C (unspecified by Lim and Cho), we used the value 0.001, as we found this gave best results. In order to produce output that is visually comparable to that produced by our model, we thresholded the occupancy probabilities. If a segment has occupancy probability less than 0.3, it is judged to be empty, if it has occupancy probability greater that 0.7, it is judged occupied, otherwise it is marked unknown. Occupied cells were drawn on the output maps as line segments, with orientation given by the maximum of the orientation probability stored for the segment. We also compared the performance of our model to that of the model of Lee and Recce (Lee and Recce 1997; Lee 1996).

Figure 7.5 shows the maps produced 10 and 30 time-steps into a run in the environment shown in figure 7.4.

By 10 time-steps, all 3 systems produced a partial map of the environment, and by 30 time-steps, they produced a fairly complete map. The neural and Bayesian methods produce broadly similar maps. However, the effect of the lateral connections can be seen in the neural method. For example there is a

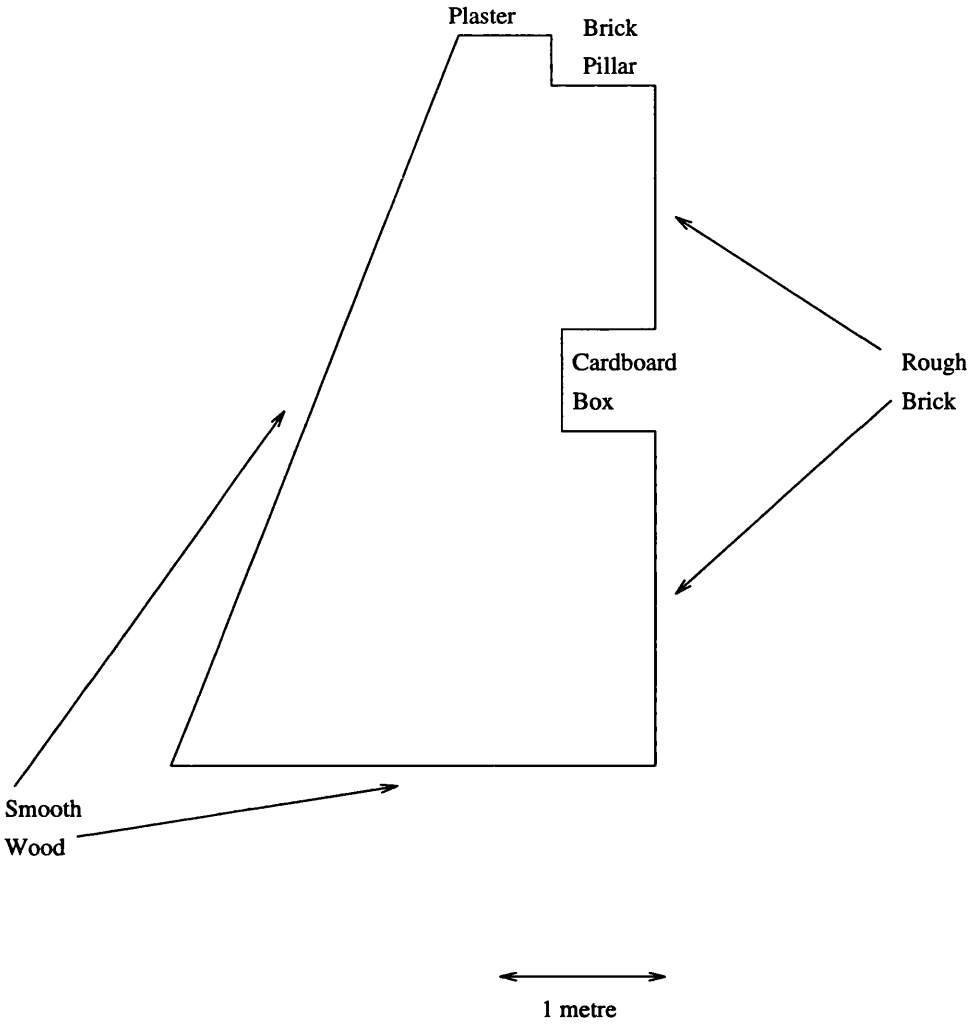

Figure 7.4: The environment used to collect sonar data. The environment contains a mix of specular and non-specular surfaces.

greater coherence and collinearity of line segments that correspond to walls. Also in the Bayesian method there are incorrectly identified line segments in the middle of free-space regions (see figure 7.5), that would cause the robot to make unneeded detours, and could make free space regions unreachable.

The feature-based method is slightly slower than the other two at accumulating a map, but is the most accurate at representing the positions of the walls. This is because of its non-probabilistic nature, whereby it must be sure of the existence of a feature before adding it to the map. For example, even by time 30, this method has not detected the lines corresponding to the cardboard box or smooth plaster wall. Note that the smaller size of the free-space

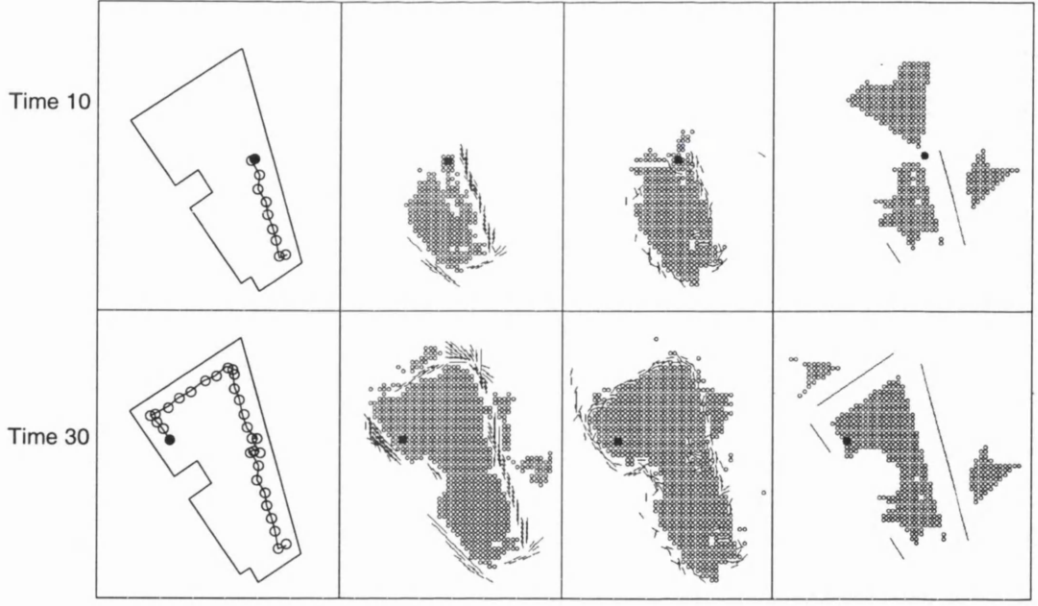

Figure 7.5: The maps produced by the neural method described in this paper, the Bayesian method of (Lim and Cho 1992), and the feature-based method of (Lee and Recce 1997), from the same sonar data, after 10 and 30 time-steps. Open circles represent areas of space marked as free, and line segments or lines represent walls. The filled circle represents the robot position.

area with this method is due to a "safety zone" of one robot diameter placed around the walls by the mapping software, rather than a failure to detect free space.

7.5 Discussion

In this chapter, we described functional similarities between grid-based maps for robot sonar and artificial neural networks. Making use of this analogy, we proposed a new grid-based mapping system making use of the neural concepts of receptive fields and lateral connections. Neurons with elongated receptive fields detect straight walls in a similar manner to a Hough transform. The lateral connections in the network are designed to reinforce maps that represent configurations likely to occur in the world, like straight walls, and inhibit configurations unlikely to occur, such as random scatter.

The receptive fields of the line neurons used in the model have properties in common with the experimentally observed receptive fields of some neurons

cortex which maintain the memory of an object through sustained activity (Miyashita 1988), similarly to the way cells in our network continue to code for features even when they are further from the robot than the maximum sonar range.

Our network functions in egocentric coordinates, and therefore requires activation patterns to be shifted after each robot movement. Droulez and Berthoz (1991) and Zhang (1996) have described neural mechanisms by which an entire neural activation pattern can be shifted to remain in egocentric coordinates after movement. This shift of activity pattern is the equivalent of dead-reckoning in egocentric coordinates, as discussed in the last two chapters. In the current implementation we did not use a neural mechanism to perform the shift, but simply shifted the array of membrane potentials with a loop. Implementing a neural mechanism for this shift would make the network more biologically realistic but it would not contribute to its performance as an engineering method for guiding the navigation of a robot.

We compared the new network to an established grid-based mapping system, and to a feature-based mapping system, and found that the representation of walls by line segments is more coherent in the new neural method than in the Bayesian method, and that both the neural and Bayesian methods were faster at building a map than the feature-based method.

The new method is based on the use of a single layer of neurons, and if implemented on parallel neural hardware, the map would updated one computational time-step after the presentation of the new sensory input.

The current model represents only the first step in neural-based sonar mapping. The real strength of neural networks is their ability to learn using real world data. The net we described here, however, has fixed weights, and a large number of parameters whose values were not systematically optimised or derived from a mathematical principle. A possibility for future work would be to incorporate learning into neural sonar mapping algorithms, and thereby reduce the number of free parameters.

system it is possible to derive mapping rules from a mathematical principle, while keeping the benefits of the new system described in this chapter.

Chapter 8

Probabilistic modelling of the sonar sensor

In the last section we described a new method of grid-based sonar mapping based on a neural network with recurrent connections. This approach showed promising results but, because it was based based on heuristics rather than derived from an optimality principle, it contained many free parameters whose values could not be guaranteed to be optimal.

In the next chapter, we will describe a new feature-based sonar mapping system which is derived from an optimality principle, Bayes' theorem. This mapping system requires an accurate probabilistic model of the sonar sensor. In particular, it requires a model which gives a probability distribution for sonar returns for every range and incidence angle from the 3 types of reflector used: specular walls, rough walls, and sharp edges. In this section, we will describe how such a model was built by collection and analysis of experimental data.

8.1 Introduction

One approach to sensor modelling is to start from a physical model of the sensing process. Kuc and Siegel (1987) describe a model of ultrasonic rangefinders that is based on principles of acoustics and a knowledge of the detection circuitry of the Polaroid ultrasonic ranging system. Their model predicts the range read from any sensor position in an environment consisting of specular

walls, corners, and sharp edges. However, it is not a probabilistic model, for a given sensor position, the model predicts a single reading, rather than specifying a probability distribution, and therefore can take no account of sensory noise. Furthermore, Kuc and Siegel do not simulate rangefinder output in the case of rough walls. A later paper deals with rough walls (Bozma and Kuc 1994); however, it describes a more complicated sensor than a simple time-of-flight rangefinder.

A different approach to sensor modelling is to start by collecting data with a real sensor, and then construct a model to fit the observed data. Leonard and Durrant-Whyte (1992) and Lee and Recce (1997) have discussed the properties of sonar returns from various reflectors with reference to actual scans and derived probabilistic models from their observations. However, these models are too restricted for our purposes. While they describe direct reflections well, they are inaccurate for oblique beam incidence angles. This is not a problem for mapping systems that start by extracting regions of constant depth, as RCDs correspond to strong returns. But restriction to RCDs means that many sonar returns are simply not used in the map construction process.

The mapping system of the next chapter will require a probabilistic model of reflection at all incidence angles and distances from specular walls, rough walls, and sharp edges. No such model has yet been given, and so we must construct our own. Instead of starting from a physical model, we follow an empirical approach. We collect a large number of sonar readings from rough walls, specular walls, and sharp edge reflectors, and use this to construct a modelled probability distribution for sensor output as a function of the distance d from the sensor to the wall, and the beam incidence angle θ_b. The models we derive give a probability distribution for the returned range, the parameters of which are functions of the distance to the object and the beam incidence angle.

8.2 Collection of Experimental Data

The sonar data presented here was collected with the Polaroid 6500 series ultrasonic ranging system used by the robot ARNE. The transducer was mounted on a motorised pivot on top of the robot, allowing multiple scans to be automatically taken from the same position with an angular spacing of 1.8^o. Due to a pole on the back of the robot, the 17 readings directly behind the robot could not be taken. The output of the rangefinding system was timed using an 8-bit digital counter, allowing a resolution of 1cm and a maximum measurable length of 253cm.

8.2.1 Reflections from rough and specular walls

All scans were collected in the testing environment shown in figure 8.1a. The robot was positioned by hand at each of the testing positions in the figure, and its position was checked with a tape measure held on a level beam of wood, from which a plumb-bob was hung to measure the distance of the sonar sensor's centre from the north (smooth wood) and west (rough brick) walls. The sensor was directed straight at the north wall, and its angle was checked for every robot position by placing a mirror flush on the north wall, and another mirror flush on the sonar sensor. The sensor was correctly aligned when multiple reflections could be seen in the two mirrors.

From each position, 50 complete scans were taken, and the raw data saved to trace files. The trace files were then processed to extract those readings corresponding to reflections from the north and west walls. In order to be sure that only the appropriate wall played a part in generating the reflection, readings were rejected if any other walls or obstacles fell in a beam of half angle 18^o from the sensor (see figure 8.2a).

8.2.2 Reflections from sharp edges

The robot was was placed by hand at the position shown by the dotted circle in figure 8.1b, with the sensor facing the south wall, and remained there

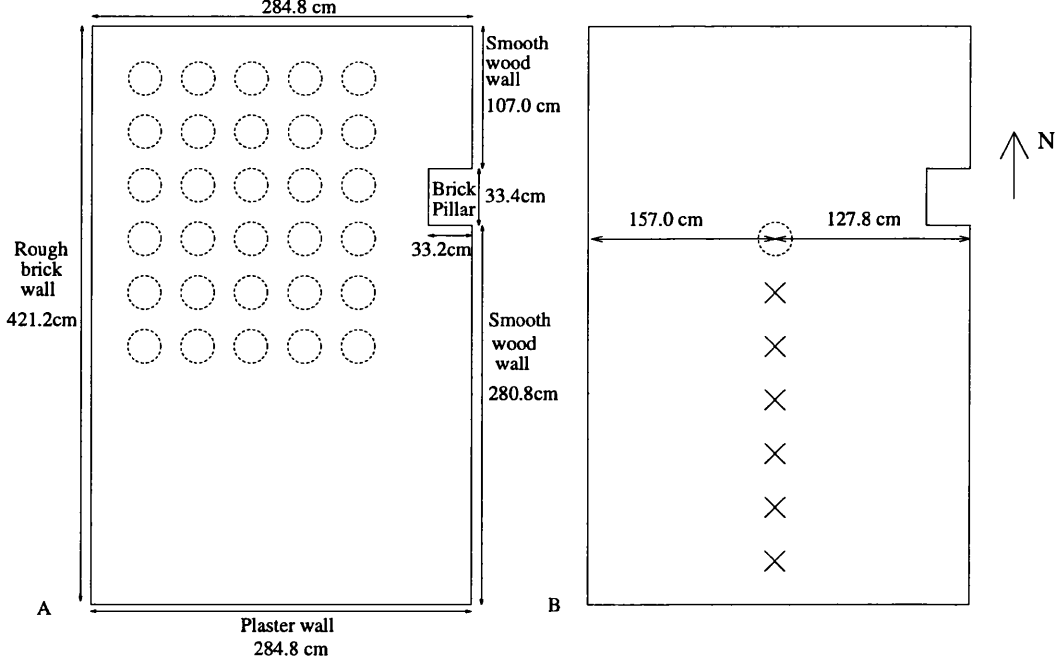

Figure 8.1: The environment in which sonar data was collected. A: Collection of data from rough and specular walls. Each circle represents a robot position. From each position, 50 complete angular scans were taken, and those readings corresponding to reflections from the north and west walls were kept. The measuring positions are on a 40cm grid, and are set relative to the west and north walls. The robots sensor zero angle was set relative to the north wall. B: Collection of data from sharp edge reflectors. The circle represents the fixed position of the robot. Each cross marks a location of the edge reflector. The edge locations are spaced 40cm apart, on the north-south line from the robot. Two different reflectors were used: the edge of a metal metre rule, and the corner of a small metal filing cabinet. 50 scans were taken for both reflectors in each position. The robots sensor zero angle was set relative to the south wall.

throughout the experiment. The sensor alignment was checked before each scan using the mirror technique described in the previous section. The perpendicular line to from the wall to the sensor was marked on the floor with tape. Two types of edge reflectors (the edge of a metal metre-ruler, and the corner of a small metal filing cabinet) were placed at 40, 80, 120, 160, 200, and 240 cm distances from the sensor along this line, and for each position, 50 scans were made and saved to a trace file.

In order to avoid the possibility of mistaking reflections from the side walls for reflections from the test object, readings were rejected if either of the side

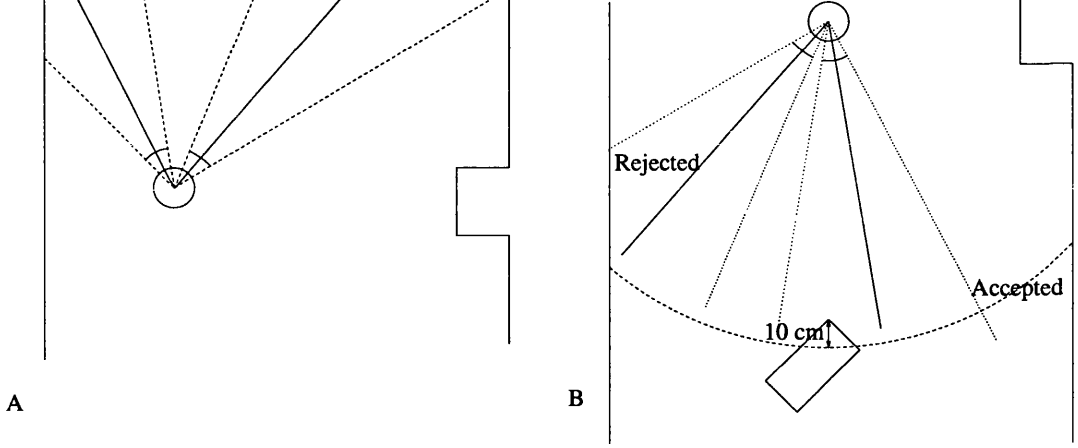

Figure 8.2: A: The criterion used for selecting sonar readings that have only been influenced by one wall. The circle represents the robot, and the solid lines emanating from it represent sensor directions. For each sensor angle, the reading is accepted if the two lines 18 degrees either side of the sensor direction (shown dotted) intersect the same wall as the beam centre, and do not hit any other walls or obstacles. B: The criterion used for selecting readings that have only been influenced by the edge reflector. If either of the two lines 18 degrees either side of the sensor direction intersect the side walls at a distance from the sensor of less that the distance to the edge reflector plus 10cm, the reading is rejected.

walls intersects a beam of half angle 18^o at a distance from the sensor of less than the distance to the edge plus 10cm (see figure 8.2b).

8.3 Results and Models

8.3.1 General form of probabilistic models

Range readings for reflections from all objects are modelled by the same family of probability distributions, characterised by three real-valued parameters, p, μ, and σ. Under the modelled distribution, the sensor will detect a direct reflection for the object in question with probability p. If a direct reflection is detected, the modelled sensor output is normally distributed with mean μ and standard deviation σ (see figure 8.3). The remainder of this section will deal with finding values of the three parameters as closed-form functions of the location of the reflecting object, in order to produce a good fit to the observed data.

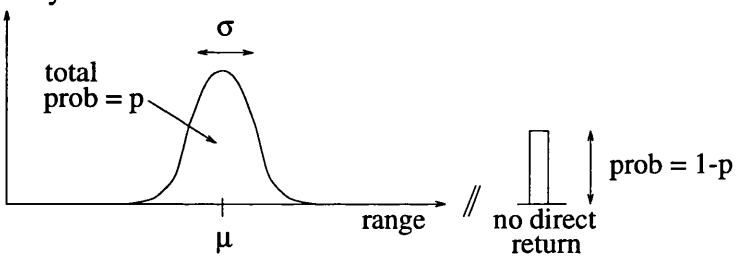

Figure 8.3: The family of probability distributions with which we model the range measured by sonar in response to a single feature. There is a probability p that a direct return will be detected from this feature. If a direct return is detected, the measured range is normally distributed with mean μ and standard deviation σ.

8.3.2 Rough walls

Figure 8.4a shows the superimposed readings from all robot positions, plotted against sensor incidence angle. Five concave bands are visible, corresponding to the five possible distances from the rough wall. All readings taken fit into one of these bands, and we therefore model $p = 1$ for any incidence angle to a rough wall.

The simplest model of reflection from a rough wall would be the "simple beam model," in which the distance measured is equal to the distance to the wall along the beam centre direction. This provides an approximate fit, as seen in figure 8.4a. However, for large angles, it is a consistent overestimate. The overestimation is almost certainly caused by finite beam width: it is likely that sound from the fringes of the beam is causing a return to be detected before the sound from the centre of the beam has reached the sensor. We used the distance d from the wall and the returned range r to compute the angle of the beam portion causing the return according to the formula $\theta_r = cos^{-1}(d/r)$ (see figure 8.5). A plot of θ_r against beam angle θ_b is shown in figure 8.6, along with two fit lines. The fit lines were chosen by hand subject to the constraint that their slope have magnitude 1.

The fit line for $\theta_b < 0$ has an intercept of -2.7^o, and that for $\theta_b > 0$ has an intercept of 4.1^o, indicating that returns are indeed caused by the fringes of the beam. The difference between the two intercepts is most likely due to

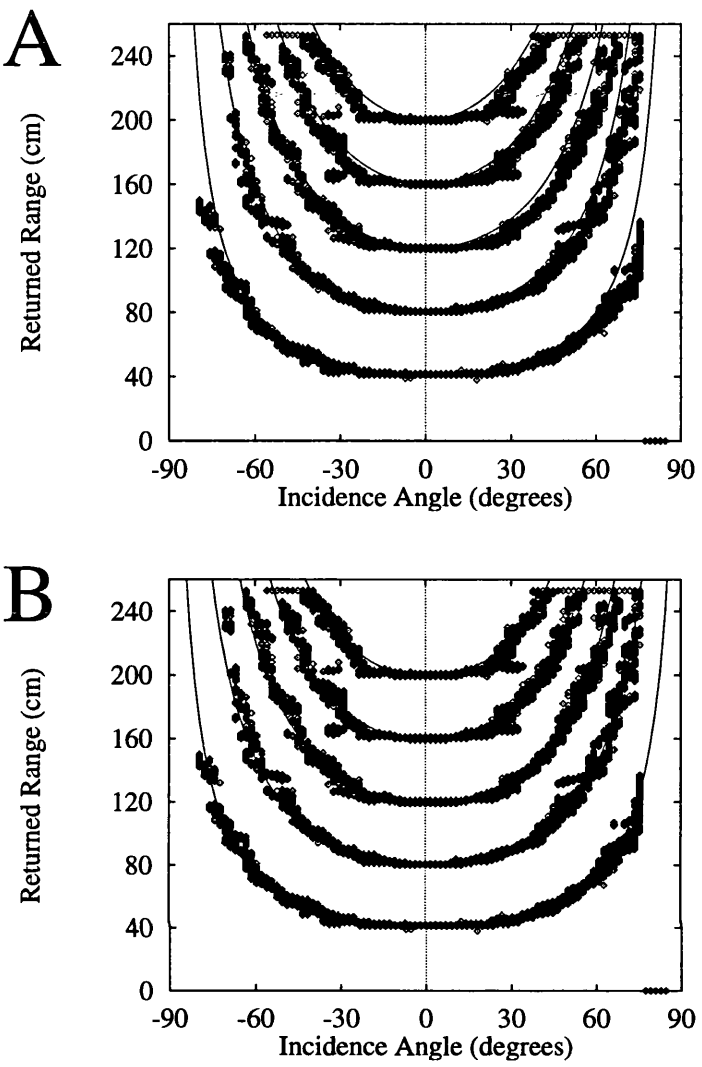

Figure 8.4: A scatter plot of range reading against sensor angle for the brick wall. A point is plotted for every reading that passes the beam test, and the plots for all sensor positions are superimposed. Five concave bands can be distinguished, corresponding to the five possible distances from the wall. A: The scatter plot has been superimposed with the prediction of the simple beam model, in which the sensor measures the distance to the wall in the direction of the beam centre-line. This provides an approximate fit but is an overestimate for oblique angles. B: The scatter plot has been superimposed with the readings predicted by the corrected model. The problem of overestimation is gone.

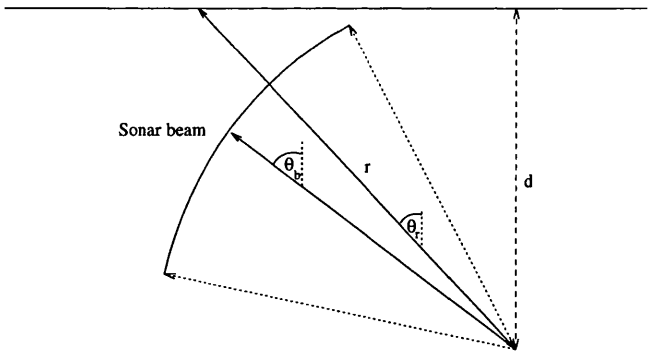

Figure 8.5: Due to non-zero beam width, reflections at oblique angles from the brick wall may come from the inside of the sonar beam. The incidence angle is θ_b, but the range reading r corresponds to a smaller angle, $\theta_r = cos^{-1}(d/r)$.

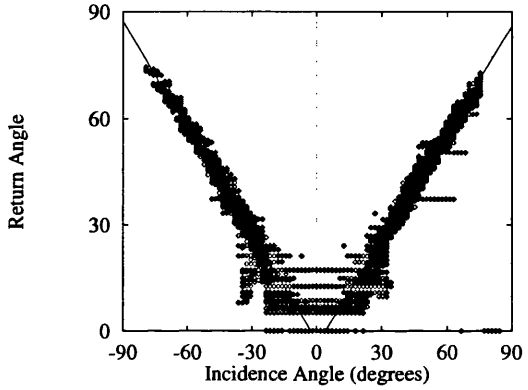

Figure 8.6: A plot of return angle θ_r against beam incidence angle θ_b, along with two fit lines. The fit lines were chosen by hand with the constraint of slope 1. Their intercepts are -2.7 and 4.1 degrees.

a small systematic error in the angular positioning of the sensor relative to the wall. This probably arises because the the sensor angle was set relative to the north (wood) wall, and the angle between the north and west walls is not quite 90^o. For the purposes of model building, we will assume that the sensor's behaviour is symmetrical, and therefore that the systematic error is 0.7^o

Using these results, we model the mean returned range from the brick wall by $\mu = d/cos(\theta_r)$, where θ_r is the angle of the beam portion causing the return, given by

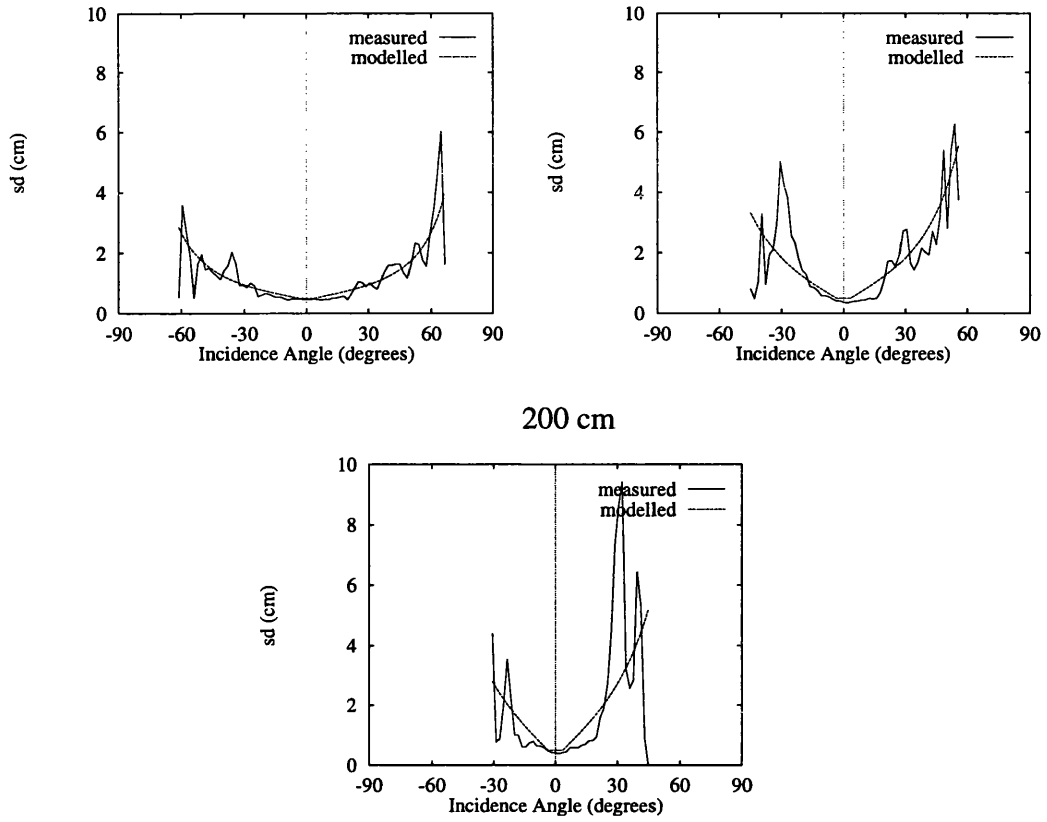

Figure 8.7: Example fits for the modelled standard deviation from the rough wall. The number above each plot is d, the perpendicular distance of the sensor from the wall.

$$\theta_r = \begin{cases} \theta_b + \phi & \theta_b \leq -\phi \\ 0 & -\phi < \theta_b < \phi \\ \theta_b - \phi & \theta_b \geq \phi \end{cases} \tag{8.1}$$

where $\phi = 3.4^o$. The corrected model can be seen in figure 8.4b.

To model the standard deviation σ of returned ranges, we assume that the return angle θ_r has a constant variability $\Delta\theta_r$. Then, since the measured range is $r = \frac{d}{cos(\theta_r)}$, the variability of range measurements will be $\Delta r = \Delta\theta_r \times \frac{d}{d\theta_r}\left(\frac{d}{cos(\theta_r)}\right)$ which is equal to $\Delta\theta_r \times d \times \frac{sin(\theta_r)}{cos^2(\theta_r)}$. As this formula predicts a zero error for $\theta_b = 0$, we also add a constant term Δr_0, equal to the standard deviation when $\theta_b = 0$. Figure 8.7 shows the modelled standard deviation, plotted against angle, with $\Delta\theta_r = 0.02$ rad and $\Delta r_0 = 0.5$cm, superimposed on graphs of measured standard deviation. Examination of the figure shows that, although the measured standard deviation is a highly convoluted function, the

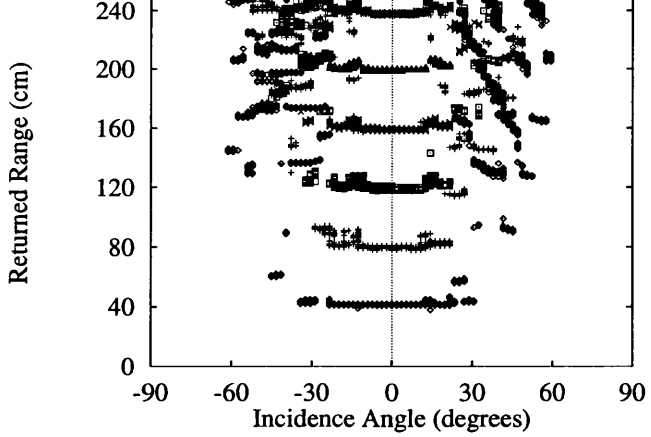

Figure 8.8: A scatter plot of range reading against sensor angle for the smooth wood wall. The plots for all sensor positions are superimposed. For small sensor angles, the readings fall into six horizontal bands corresponding to the six different sensor distances. For large sensor angles, the readings are unpredictable, due to multiple specular reflection.

model captures the broad shape of the standard deviation curve.

8.3.3 Specular walls

Figure 8.8 shows the superimposed range readings from the smooth wood wall, plotted against incidence angle for all sensor positions. Unlike the case of the rough wall, the readings do not fall into neat bands, except for small incidence angles where the measured range is approximately equal to the perpendicular distance to the wall. At oblique incidence angles, the range returned is always greater than the perpendicular distance, but its exact value depends erratically on the beam angle. The erratic behaviour is almost certainly caused by multiple specular reflection.

In order to compute empirical estimates of p, μ, and σ, we classify a reading as a *direct return* if the range measured lies within 10cm of the perpendicular distance to the wall (this value is chosen to be large enough to allow readings where the distance is misjudged to be included, but small enough to exclude multiple specular reflections). An estimate for p is then calculated as the fraction of all returns that are direct, and estimates for μ are given σ by the mean and standard deviation of direct returns. We define $\delta\mu$ to be the

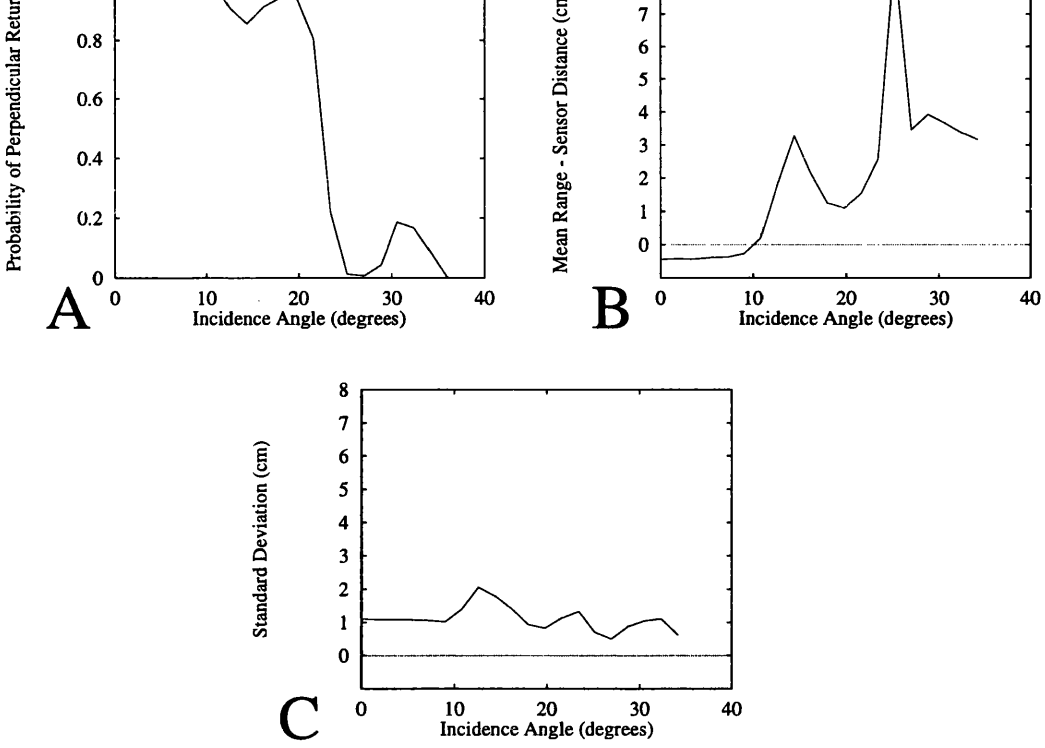

Figure 8.9: A: Probability of direct return plotted against incidence angle for the smooth wood wall. The positions of peaks and troughs in this graph correspond to maxima and minima of the transducer emitted power function. B and C: Mean and standard deviation of the measured range minus perpendicular distance, plotted against incidence angle

difference between the empirical μ and the perpendicular distance to the wall (i.e. the mean amount by which the sensor overestimates the range). Note that $\delta\mu$ is defined with respect to the perpendicular distance d to the wall, rather than the distance $d/cos(\theta_b)$ along the beam centre line.

The empirical values of p, $\delta\mu$, and σ were found not to depend on perpendicular distance in any systematic way. We therefore computed grand average estimates of the parameters as a function of absolute incidence angle. These are shown in figure 8.9.

The positions of the peaks and troughs in figure 8.9 are in correspondence with the peaks and troughs of the sonar sensor angular power function (Polaroid 1991). Leonard and Durrant-Whyte (1992, appendix A.2) suggest modelling the transducer as a circular piston in an infinite baffle. In this case, the emitted acoustic power depends on angle as

$$P(\theta) = \left[\frac{2J_1(kasin(\theta))}{kasin(\theta)}\right]^2$$

Where J_1 is a Bessel function of the first kind, $k = 2\pi/\lambda$ is the wave number, and a is the transducer radius. We have taken a value of 15.5 for ka, as this gives the best fit with the positions of the peaks and troughs in the graph of figure 8.9a. According to the circular piston model, the acoustic power at a point in space is also modulated by a factor of $1/r$ where r is the distance from the transducer. However, the designers of the Polaroid sensing system included a time-variable amplifier whose gain increases linearly with time to ensure that reflections from objects at all distances are detected with the same amplitude. This explains the above noted fact that p, $\delta\mu$, and σ do not vary systematically with d.

We now seek to relate the empirical values of p, $\delta\mu$, and σ to the radiated power function $P(\theta)$. Figure 8.10a shows a graph of empirical return probability against log emitted power. The empirically estimated probabilities show a roughly sigmoidal dependence on $logP$, asymptoting to $p = 1$ for $logP > -4$, and $p = 0$ for $logP < -6$. There is an outlier at $logP = -10.7$, corresponding to an anomalously high return probability for scans with beam incidence angle 14.4°. According to the circular piston approximation, the radiated power is very low at this angle ($P(\theta) = 2 \times 10^{-5}$) but increases sharply on either side. The most likely explanation for the presence of the outlier is that in reality the transducer power at this angle is greater than that predicted by the circular piston approximation. We have ignored scans from this angle when model-fitting.

The standard statistical technique for modelling the dependence of a probability of an event on the value of a continuous independent variable is *logistic regression* (see e.g. Agresti 1990). In logistic regression, the probability π of an event is modelled as a logistic function of an independent parameter x:

$$\pi(x) = \frac{1}{1 + exp(-\alpha + -\beta x)}$$

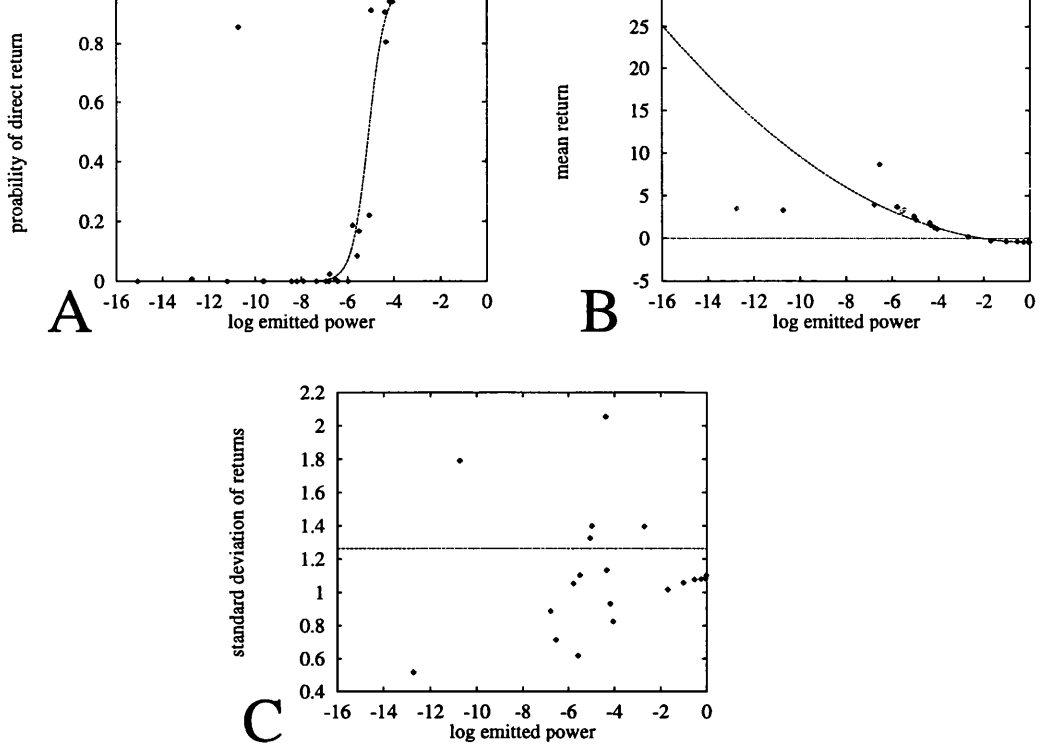

Figure 8.10: A: Probability of direct return plotted against log emitted power for the smooth wood wall. The dotted line shows the fit generated by logistic regression. The outlier at $logP$ = -10.7 was ignored in generating the fit. B: Mean range overestimate plotted against log power. The dotted line shows the fit generated by linear regression in $(logP)^2$. C: Standard deviation of returned ranges plotted against log power. The dotted line shows the modelled constant standard deviation equal to the residual s.d. after the fit shown in B.

The values of α and β are found by maximum-likelihood estimation. In our case, taking $\log P(\theta)$ as independent parameter, the values of α and β were found to be 14.38 and 2.83, using the statistical package SPSS.

Figure 8.10b shows the empirical value of $\delta\mu$ plotted against log emitted power. Again there is an outlier at $logP(\theta) = -10.7$, which is ignored for the purposes of model fitting. The outlier at $logP = -13$ was *not* ignored, but it has a small leverage due to the very low probability of detecting returns of this power. The plots shows a roughly parabolic shape, and we therefore use a quadratic model:

$$\delta\mu = \eta_s + \zeta_s logP(\theta)^2$$

The best fit parameters were found by linear regression to be $\eta_s = -0.12$ and $\zeta_s = 0.10$.

Figure 8.10c shows the standard deviation of returned ranges plotted against log power. Unlike the probability and mean range, the standard deviation shows little systematic dependence on log power, and is therefore modelled by a constant. The appropriate value is the residual standard deviation of the regression model, i.e. the standard deviation of the difference between the observed readings and those predicted by the regression model. This gives a value of $\sigma = 1.26$.

8.4 Sharp Edges

We would like to relate the probability distribution of returns from an edge reflector to received power, as we did for specular walls. Since a sharp edge produces a diffracted signal, the acoustic power received by the sensor is modulated by a factor of λ/d, where λ is the sonic wavelength, in addition to the previous factor of $1/d$ that was compensated by the time-dependent gain mechanism of the detection circuitry. Figure 8.11 shows plots of p, μ, and σ against log power. We have taken a value of $0.695cm$ for λ, assuming a frequency of $49.4kHz$ and speed of sound $343.2ms^{-1}$. The resulting graphs shown in figure 8.11 are similar to the case of the specular wall, and we perform the same analysis of logistic regression for p and quadratic regression for μ. The regression coefficients are found to be $\alpha = 21.24$, $\beta = 3.33$, $\eta = -0.023$, $\zeta = 0.056$, and $\sigma = 0.74$.

8.5 Discussion

8.5.1 Summary of Models

In this chapter, we have given closed-form probabilistic models of time-of-flight sonar returns from rough walls, specular walls, and sharp edges. In all cases, the modelled sensor output comes from the same family of probability distributions, shown in figure 8.3. A probability distribution from this fam-

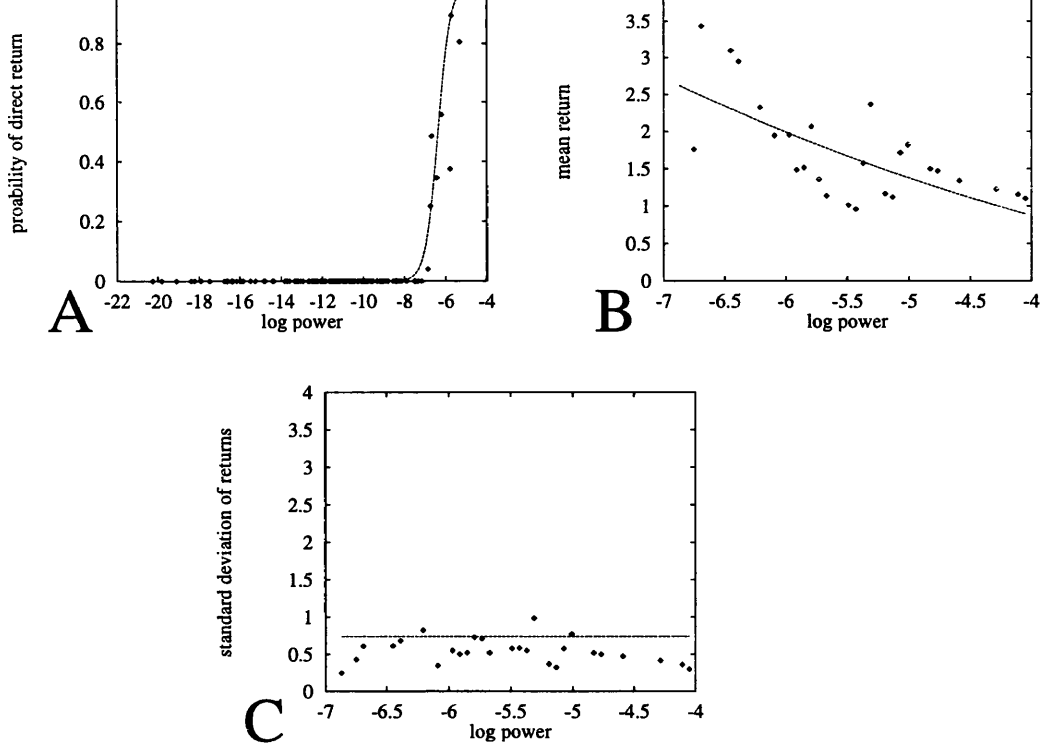

Figure 8.11: A: Probability of direct return plotted against log received power for the two sharp edge reflectors. The dotted line shows the fit generated by logistic regression. Note that for edge reflectors, the log received power is modelled by the log emitted power $P(\theta)$ multiplied by a factor of λ/d where λ is wavelength and d is distance to the edge reflector. B: Mean range overestimate plotted against log power. The dotted line shows the fit generated by linear regression in $(logP)^2$. C: Standard deviation of returned ranges plotted against log power. The dotted line shows the modelled constant standard deviation equal to the residual s.d. after the fit shown in B.

ily is specified by three parameters, p, μ, and σ, where p is the probability that the sensor will detect a reflection from the feature in question, μ is the mean measured range of returns caused by the feature, and σ is the standard deviation.

Rough Walls

In the case of rough walls, the distribution parameters depend on the perpendicular distance d from the sensor to the wall and the beam incidence angle θ_b through an intermediate parameter, θ_r. θ_r represents the angle of the beam portion causing the return, and is given by:

$$\theta_r = \begin{cases} \theta_b + \phi & \theta_b \leq -\phi \\ 0 & -\phi < \theta_b < \phi \\ \theta_b - \phi & \theta_b \geq \phi \end{cases}$$

Where $\phi = 3.4^o$. The model parameters are given by:

$$\begin{aligned} p &= 1 \\ \mu &= d/cos(\theta_r) \\ \sigma &= \Delta r_0 + \Delta\theta_r \times d \times \frac{sin(\theta_r)}{cos^2(\theta_r)} \end{aligned}$$

Where $\Delta r_0 = 0.5cm$ and $\Delta\theta_r = 0.02$.

Specular walls

In the case of specular walls, the distribution parameters depend on perpendicular distance d and incidence angle θ_b through the modelled emitted power of the sensor, $P(\theta)$, which is given by:

$$P(\theta) = \left[\frac{2J_1(kasin(\theta))}{kasin(\theta)}\right]^2$$

where $ka = 15.5$. The distribution parameters take the values

$$\begin{aligned} p &= \frac{1}{1 + e^{-\alpha_s - \beta_s logP(\theta_b)}} \\ \mu &= d + \eta_s + \zeta_s logP(\theta_b)^2 \\ \sigma &= \sigma_s \end{aligned} \tag{8.2}$$

where $\alpha_s = 14.38$, $\beta_s = 2.83$,$\eta_s = -4.22$, $\zeta_s = 0.10$, and $\sigma_s = 1.26$.

Sharp edges

In the case of sharp edges, the distribution parameters depend on the distance d from the sensor to the edge, and incidence angle θ through a different power function $P(\theta, d)$ defined by

$$P(\theta, d) = \frac{\lambda}{d}\left[\frac{2J_1(kasin(\theta))}{kasin(\theta)}\right]^2$$

The distribution parameters take the values

$$\begin{aligned} p &= \frac{1}{1+e^{-\alpha_e-\beta_e logP(\theta_b,d)}} \\ \mu &= d+\eta_e+\zeta_e logP(\theta_b)^2 \\ \sigma &= \sigma_e \end{aligned} \tag{8.3}$$

where $\alpha_e = 21.24$, $\beta_e = 3.33$, $\eta_e = 0.023$, $\zeta_e = 0.056$, and $\sigma_e = 0.73$.

8.5.2 Inter-system variability

The models and parameters presented above were derived from experimental data collected with a single ultrasonic rangefinder. A natural concern is that results derived with this particular sensor might not be applicable to other sensors.

We believe that the functional form of the models presented above will generalise to other ultrasonic rangefinding systems, although probably not to other types of rangefinder (such as infra-red or laser rangefinders). It is still possible, however, that the parameters of the model will be different for different ultrasonic sensors. The parameter ka, for example, is the product of wave-number with sensor radius. This parameter may vary due to variability in the manufacturing process, and it is worth noting that our estimated value of 15.5 for ka differs from Leonard and Durrant-Whyte's 17.2. Other parameters, such as the regression coefficients η and ζ might also be expected to vary due to variability in the thresholding mechanisms of the rangefinding system. This issue can only be resolved by repeating the experiments described here with other sensors.

The sensor models of this chapter were constructed because they are required for the mapping system of the next chapter. However, these models could also be useful for the construction of sonar sensor simulations.

Simulation of a robot's interaction with its environment has obvious advantages for the development and testing of robot control systems. If robot performance can be quantified in a simulated world, then considerable time can be saved in the development process. Furthermore, learning techniques such as genetic algorithms require such a large number of performance evaluations that simulation is essential if development time is not to be become an insurmountable obstacle.

If simulation is used, however, accuracy is essential. If an inaccurate simulation is used in development, the performance of the resulting control system will be optimal for the inaccurate simulation, rather than the real world. In particular it is essential to have a good model of the inaccuracy and noise inherent in the robots sensors. Otherwise the resulting system will only perform optimally in an unrealistic noise-free environment.

The models of this chapter, unlike the models of Kuc and Siegel (1987), are probabilistic, so they take into account sensory noise. One thing that they do not take account of, however, is multiple specular reflection. If sonar is obliquely incident on a specular wall, the models of this chapter predict that "no direct return is produced," but they do not specify if a multiple specular reflection is produced, or if no return is received at all. In order to construct a simulator, it is necessary to specify this. One way to do it might be to place "virtual features" at the position of actual features reflected in the specular walls (see figure 8.12).

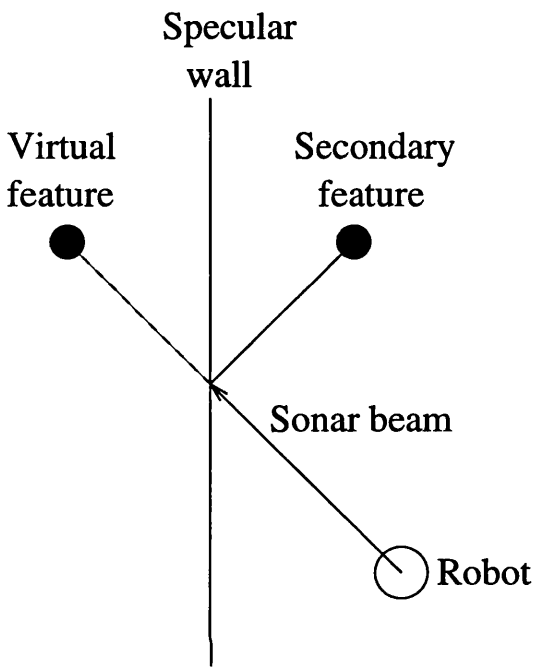

Figure 8.12: A possible method for simulating multiple specular reflections. Where the sonar beam is incident on a specular wall, a "virtual feature" is constructed for every feature that may act as a secondary target, in a position corresponding to the reflection of the original feature in the specular wall. The probabilistic models of the above section can then be used for this virtual feature.

Chapter 9

A Bayesian system for feature-based sonar mapping

9.1 Introduction

In chapter 7, we described a grid-based sonar mapping system that was inspired by models of primary visual cortex, and sought to overcome the incorrect assumption of statistical independence of different grid segment occupancies. This method produced good results when tested with real sonar data. However, it was based on heuristics, rather than derived from mathematical principles. This meant that we could not be sure that its results were optimal in any way.

Previous grid-based mapping systems have been derived from Bayes' theorem, but making the assumption of independence of occupancy of different grid segments. They have to make this assumption, because they are *incremental* methods. This means that an *a posteriori* probability distribution is stored in memory, and as new sonar data arrives this probability distribution is updated to take account of the new data. By making the independence assumption, these incremental grid-based systems can represent the probability distribution by storing only a single probability for each segment. A truly Bayesian incremental mapping system, which did not make the independence assumption, would need to store a probability for every single possible map configuration. This is clearly impractical as the number of possible configurations is exponential in the number of grid segments.

probabilities for all possible configurations – a single most probable map will suffice. This suggests a new non-incremental, Bayesian mapping strategy, in which a record of all sensory readings is kept, and these are used to construct a probability function on the space of all maps, which is then searched for the map of maximum *a posteriori* probability. The space of all maps is very large, but with a suitable choice of feature-based map representation, it has a structure that makes an approximate global search computationally tractable. In this chapter we describe such a system.

In section 9.2 we lay out a general formalism, and construct a probability function on the space of all possible maps. In section 9.3 we describe the map representation and details of the search strategy for the high-dimensional space of all possible maps. In section 9.4, we describe a quantitative map quality metric used to test the algorithm. In section 9.5 we describe the results of the algorithm. In section 9.6 we draw conclusions, and in section 9.7 make proposals for future research. The mathematical notation used in this chapter is summarised in appendix B.

9.2 Probabilistic Modelling

9.2.1 General Formalism

The mapping system described in this chapter is *not* an incremental system. Instead, it is a prescription for generating a map given a history of measurements. Each measurement α is characterised by transducer location $\xi_\alpha = (x_\alpha, y_\alpha, \theta_\alpha)$, and measured range R_α. We consider the transducer locations $\{\xi_\alpha\}$ to be fixed, but the ranges $\mathcal{R} = \{R_\alpha\}$ to be random variables. We assume that the output of the rangefinder is discrete (which will always be the case for systems connected to a digital computer). We model the robot's environment as consisting of a set of features $\mathcal{F} = \{F_i : i = 1 \ldots N_f\}$, each of which is characterised by a feature type and location. In this work we restrict the F_i to describe walls and edges, but the analysis of this section also

By Bayes' theorem, the probability of a map $\{f_i\}$ given a set of readings $\{r_\alpha\}$ is given by

$$P(\mathcal{F} = \{f_i\} | \mathcal{R} = \{r_\alpha\}) = \frac{P(\mathcal{R} = \{r_\alpha\} | \mathcal{F} = \{f_i\}) P(\mathcal{F} = \{f_i\})}{P(\mathcal{R} = \{r_\alpha\})} \quad (9.1)$$

To produce a map for a given measurement history, we must find the feature configuration $\{f_i\}$ that maximises this expression. Since the denominator does not depend on $\{f_i\}$, we may ignore it. The second term of the numerator is the *a priori* probability for the map $\{f_i\}$. The form of this function will depend on the type of environment the robot is typically found in – for example in a regular indoor environment, maps consisting of straight walls intersecting at right angles would have a higher probability than maps consisting of an unorganised collection of walls intersecting at acute angles. In the remainder of this paper we will model this term by a prior $P_f(f)$ on the position of each feature, and a probability distribution on the number of features, N_f. We use a Poisson distribution, where the mean λ is a parameter of the model (equal to 10 for the experiments below):

$$P(\mathcal{F} = \{f_i\}) = \frac{e^{-\lambda}\lambda^{N_f}}{N_f!} \prod_{i=1}^{N_f} P_f(f_i) \quad (9.2)$$

$P_f(f_i)$ is the prior probability density for the position of a single feature. We want our choice of P_f to have minimal effect on the results of the mapping system. In Bayesian theory, such a prior is called a *reference prior* (Antelman 1997), and will be approximated by a flat distribution. A flat distribution over a continuous variable is *improper*, i.e. cannot be normalised to have total probability 1, but may be taken as a limiting case.

In order to model $P(\mathcal{R} = \{r_\alpha\} | \mathcal{F} = \{f_i\})$, we need a model of the way environmental features interact to produce range readings. Our model is illustrated in figure 9.1. For each range reading α, each feature F_i may produce an echo, which arrives back at the transducer with time delay corresponding to a distance of $X_{i\alpha}(F_i, \xi_\alpha)$ (If the feature does not produce a reflection, $X_{i\alpha}$ is

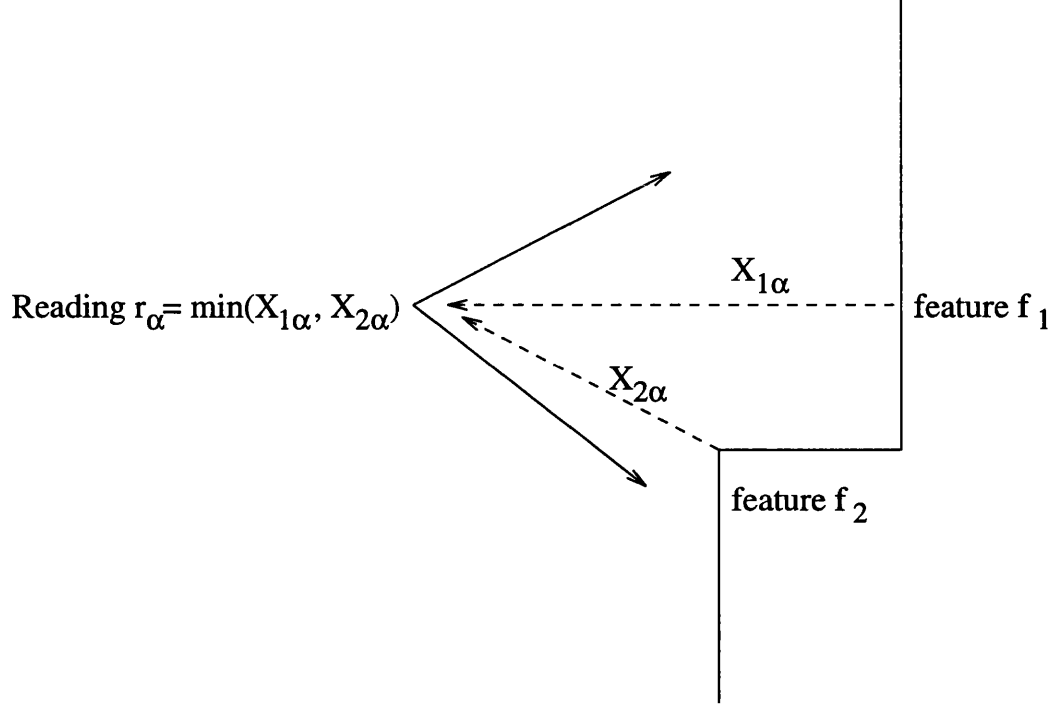

Figure 9.1: The way environmental features interact to produce a sonar reading. Each feature f_i produces an echo $X_{i\alpha}$ in response to sonar beam α. The range measured by the sensor is $min_i(X_{i\alpha})$, the smallest of the echos. In the situation shown in the figure, the echo from the corner f_2 is shorter than the echo from the wall f_1, so the sensor measures the distance to the corner, $X_{2\alpha}$.

defined to be infinity). We will often abbreviate $X_{i\alpha}(F_i, \xi_\alpha)$ to $X_{i\alpha}$, with the functional dependence on ξ_α and F_i being understood. The range reported by the sensor corresponds to the first echo to arrive at the sensor, and therefore will be $min_i\,(X_{i\alpha})$, the smallest of all the ranges produced by the features in the environment (see figure 9.1). In order to account for sensory noise, the $X_{i\alpha}$ are modelled by stochastic functions of feature position. We assume that sensory noise is a random process that affects separate readings independently, and therefore that the $X_{i\alpha}$ are independent for a given feature configuration $\{F_i\}$. The set of R_α are therefore also independent for a given feature configuration.

Now, because $R_\alpha = min_i\,(X_{i\alpha})$, R_α will take the value r when all the $X_{i\alpha}$ are greater than or equal to r, but they are not all strictly greater than r:

$$P\,(R_\alpha = r) = P\,(X_{i\alpha} \geq r, \forall i) - P\,(X_{i\alpha} > r, \forall i)$$

$$P(R_\alpha = r) = \prod_i P(X_{i\alpha} \geq r) - \prod_i P(X_{i\alpha} > r) \tag{9.3}$$

We define the *map score* of a map $\{f_i\}$ given a set of readings $\{r_\alpha\}$ to be the log of the *a posteriori* probability of a feature configuration:

$$SC_{map}(\{f_i\};\{r_\alpha\}) = logP(\mathcal{F} = \{f_i\} | \mathcal{R} = \{r_\alpha\})$$

Using equations 9.1, 9.2, and 9.3, the independence of the $\{R_\alpha\}$, and the flat approximation to $P_f(f)$, this may be written as:

$$SC_{map}(\{f_i\};\{r_\alpha\}) = const + \sum_\alpha SC_{reading}(r_\alpha; f_i) + logP(N_f) \tag{9.4}$$

Where $SC_{reading}(r_\alpha;\{f_i\})$ is the *reading score* of a reading r_α, given a map $\{f_i\}$:

$$SC_{reading}(r_\alpha;\{f_i\}) = log\left[\prod_i P(X_{i\alpha} \geq r_\alpha) - \prod_i P(X_{i\alpha} > r_\alpha)\right] \tag{9.5}$$

The quantity $SC_{reading}(r_\alpha;\{f_i\})$ can be interpreted as describing how well the observation r_α is explained by the features $\{f_i\}$ through their reflections $X_{i\alpha}$. To interpret this, let us consider the effect of adding a new feature to an existing map, illustrated in figure 9.2. If the feature predicts a shorter reading than r_α, both the left and right products will decrease, lowering the overall reading score. If the feature predicts exactly r_α, then the left product will not decrease, but the right one will, increasing the reading score. If the feature predicts a longer reading than the observed one (or no return at all), neither term will be affected and the reading score will be unchanged.

Note that if just one reading has *a posteriori* probability of zero, then the map score is $-\infty$. This is an undesirable situation, because it means that the scores of *all* other readings will have no effect on the total score, and an incorrect map will be built. Furthermore, if there are inaccuracies in the

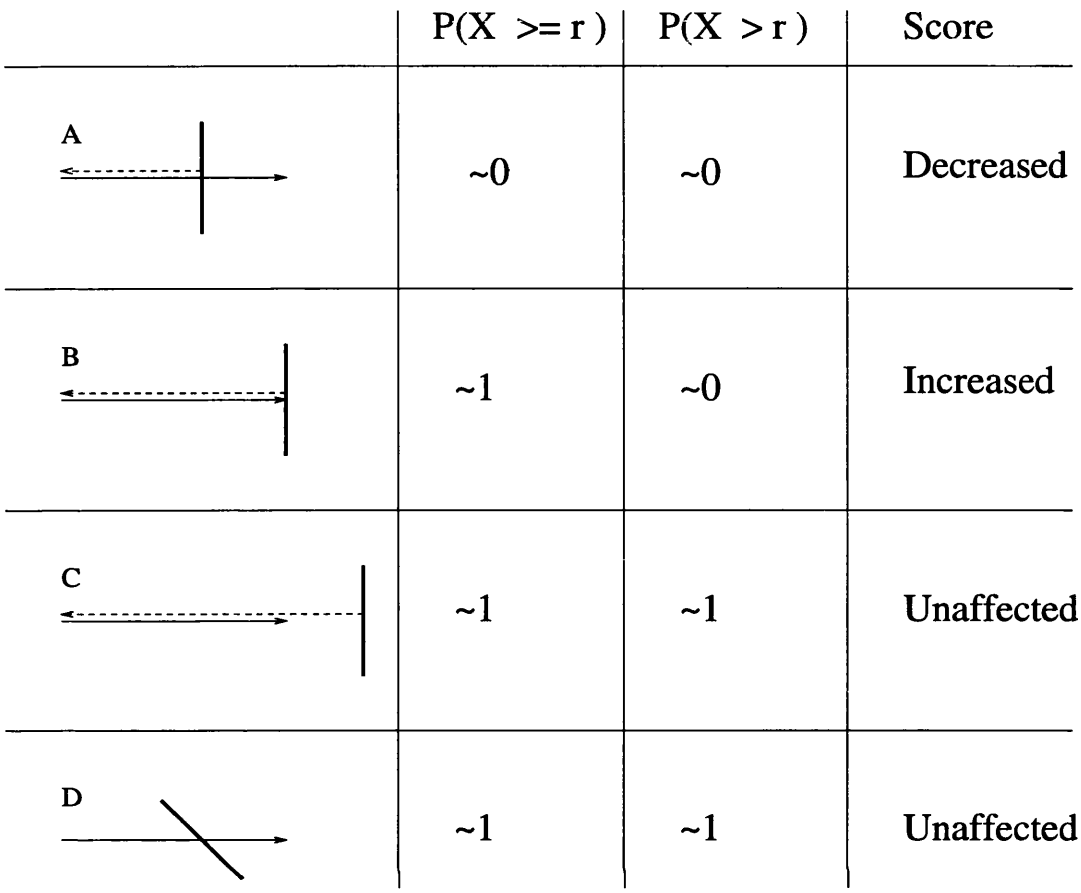

Figure 9.2: The way the position of an individual feature affects $SC_{reading}(r_\alpha; \{f_i\})$. The solid arrow represents an observed sonar reading with measured range r_α, from sensor position ξ_α. The thick line represents a hypothesised feature f_i. The dotted arrow represents $X_{i\alpha}(f_i, \xi_\alpha)$, the predicted reflection off feature f_i for sensor position ξ_α. A) If the feature predicts a reading shorter than r_α, then $P(X_{i\alpha} \geq r_\alpha)$ and $P(X_{i\alpha} > r_\alpha)$ will both be small, and the overall score will decrease. B) If the feature predicts a reading equal to r_α, $P(X_{i\alpha} \geq r_\alpha)$ will be close to 1, but $P(X_{i\alpha} > r_\alpha)$ will be small, and the overall score will increase. C) If the feature predicts a reading longer than r_α the score is unaffected. D) A specular wall at an oblique angle predicts no reflection, so $X_{i\alpha} = \infty$, $P(X_{i\alpha} \geq r_\alpha)$ and $P(X_{i\alpha} > r_\alpha)$ are both 1, and the score is unchanged.

disproportionately large effect on the overall score. We therefore introduce a probability floor limit p_{min} into the system, and any predicted probabilities below p_{min} are replaced with p_{min}. For the experiments described below, p_{min} was 10^{-4}.

9.2.2 Similarity to statistical model fitting

The mapping system described here is, in effect, a type of statistical model fitting. A map $\{f_i\}$ defines a probability distribution on the set of all possible readings, and we can therefore think of the positions of features in the map as being the parameters of the probability distribution. We are given a set of observed data (the sonar readings), and we need to find the values of the parameters that fit the data best. In this case, finding the best values of the parameters means finding the best map.

The mapping system must determine the number of features in the map, as well as the optimal positions for those features. This situation is similar to fitting a mixture of Gaussians to a set of data, where one must determine the number of components to fit, as well as their means and variances. Even in the simple case of a mixture of Gaussians, this is a difficult ongoing research problem (see e.g. Richardson and Green 1997).

A common problem in statistical model fitting is *overfitting*, when the model follows small irregularities of the sample data, rather than capturing the real-world phenomenon being modelled. In the current situation, overfitting will be manifested by the appearance of a large number of small map features which match the positions of individual sonar reflection points.

9.3 Implementational details

In the last section, we derived a formula (eqns 9.4 and 9.5) for the *a posteriori* probability of a map $\{f_i\}$ given a set of sonar readings $\{r_\alpha\}$.

In this section, we will describe a mapping algorithm based on this formalism. Section 9.3.1 describes how the probabilistic models of chapter 8 are used

to calculate the terms $P(r_{i\alpha} \leq r_\alpha)$ and $P(r_{i\alpha} > r_\alpha)$ that appear in equation 9.5. Section 9.3.2 describes a "mean-field" approximation that reduces the intractable search over the very large space of maps to a tractable search over the space of a single map feature. Section 9.3.3 describes how this search is performed, and section 9.3.4 describes how the features found by this search are combined to produce a map. Finally, section 9.3.5 describes a technique that further reduces computation time by filtering the set of readings used to build the probability function.

9.3.1 Physical Model

In order to evaluate $SC_{reading}(r_\alpha; \{f_i\})$, we need to know the probability distribution of $X_{i\alpha}(f_i, \xi_\alpha)$ for a given feature f_i and sensor position ξ_α. This is where we need the empirical model of chapter 8.

In this model, there are three types of feature: specular walls, rough walls, and sharp edges. A feature f_i is described by a variable $t = 1 \ldots 3$ describing its type, and in the case of a wall feature two pairs of Cartesian coordinates describing the wall's endpoints, or in the case of an edge feature a single pair of coordinates describing the edge position.

The probability distributions $p(X_{i\alpha}(f_i, \xi_\alpha) = r)$ are given by the equations in section 8.5.1. These equations give continuous probability distributions. However, the actual readings made by the sonar sensor are discrete, with 1cm bin resolution. We calculate the probability that $X_{i\alpha}$ will lie in a given range bin by the integral of the probability density over that bin. Because the modelled distribution is Gaussian, this integral may be computed analytically.

Effect of wall endpoints

In chapter 8, we modelled the probability distribution of returns from a wall as a function of perpendicular distance and beam incidence angle. We did not model the dependence on lateral displacement along the wall. One would not expect the lateral displacement to have an effect unless the beam is close to either end of the wall – and we deliberately excluded these readings from the

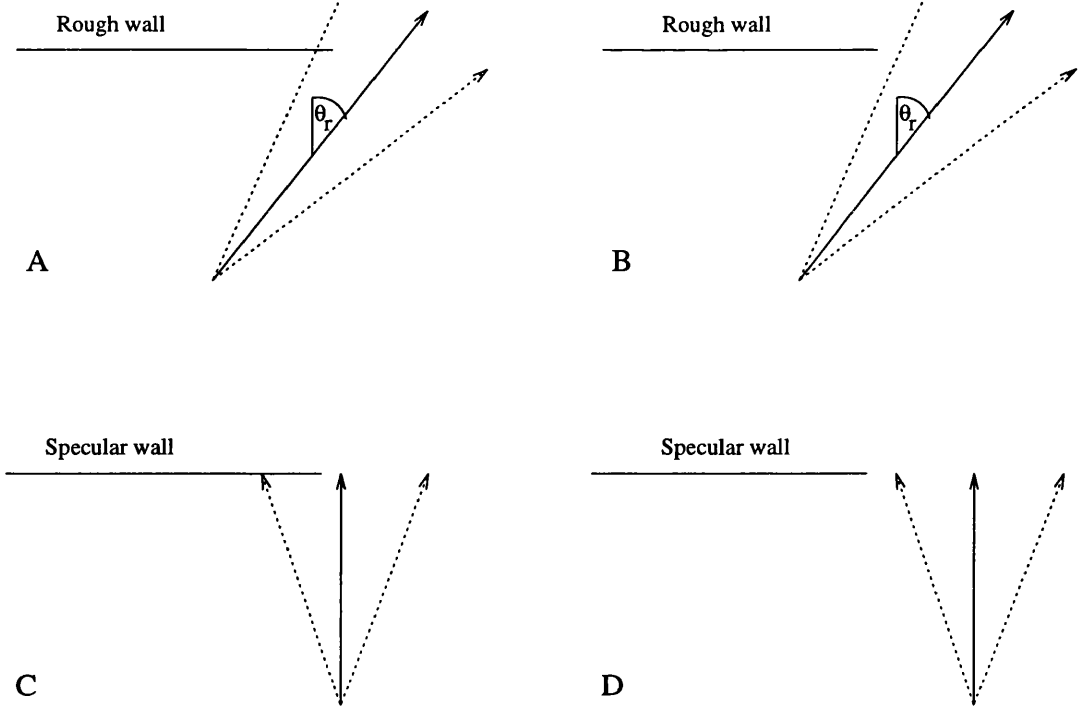

Figure 9.3: The way we model the effect of wall endpoint positions. For rough walls, a reflection is produced if any part of the wall intersects any part of a sonar beam, of half-width 0.25 radians, centred on the returning radiation direction θ_r; so a reflection will be produced in (A) but not in (B). For specular walls, a reflection is produced if any part of the wall intersects any part of a sonar beam, of half-width 0.25 radians, centred on the perpendicular line from the sensor position to the wall; so a reflection is produced in (C) but not (D). Note that for specular walls, the beam centre direction does not play a role in this criterion. But it still affects the return probability, because the modelled probability of reflection is very low if the beam direction is not close to perpendicular.

analysis of chapter 8.

If the beam misses the wall by a large distance, one would not expect a reflection. But if the beam only just misses the wall, one might still expect a reflection to be produced. We will see in section 9.3.3 that making a certain simple choice of model for this phenomenon will vastly speed up computation time for the mapping algorithm. The simple model is illustrated in figure 9.3, and works as follows: We consider a beam of half-width 0.25 radians, centred on the returning sonar radiation at angle θ_r. For rough walls, θ_r is given by equation 8.1, and for specular walls, all reflections are perpendicular, so θ_r is

with probability p given by the models of chapter 8; if not, no reflection is produced, and $p = 0$ (see figure 9.3).

9.3.2 Mean-field approximation

We have recast the problem of map construction as one of a global optimisation of $SC_{map}(\{f_i\};\{r_\alpha\})$ over the space of all maps $\{f_i\}$. Unfortunately, global optimisations of this size are very difficult problems, with no general solution. Even advanced global optimisation techniques such as simulated annealing and genetic algorithms cannot thoroughly search a space of this size. This section will describe a way to make this global optimisation problem computationally tractable.

To make the search tractable, we will use a mean-field approximation to find the positions of plausible features in the environment, and then choose a subset of these features that gives a high map score.

We will define the *mean field* score of a feature f_0 to be the log likelihood of the observed readings given the presence of the *single* feature f_0 in the map $\mathcal{F}$:

$$SC_{mf}(f_0;\{r_\alpha\}) = logP(\mathcal{R} = \{r_\alpha\}|f_0 \in \mathcal{F})$$

Every f_0 that locally maximises this function will correspond to a plausible single feature.

It is possible to explicitly compute this score by performing a sum over all maps that contain the feature f_0. The details of the calculation are given in appendix C. The result is:

$$SC_{mf}(f_0;\{r_\alpha\}) = \sum_\alpha log\left[A_\alpha P(X_{0\alpha} \geq r_\alpha) - B_\alpha P(X_{0\alpha} > r_\alpha)\right] \qquad (9.6)$$

The numbers A_α and B_α depend only on the observed range r_α, and can be interpreted as the probability that an "average" feature causes a reading

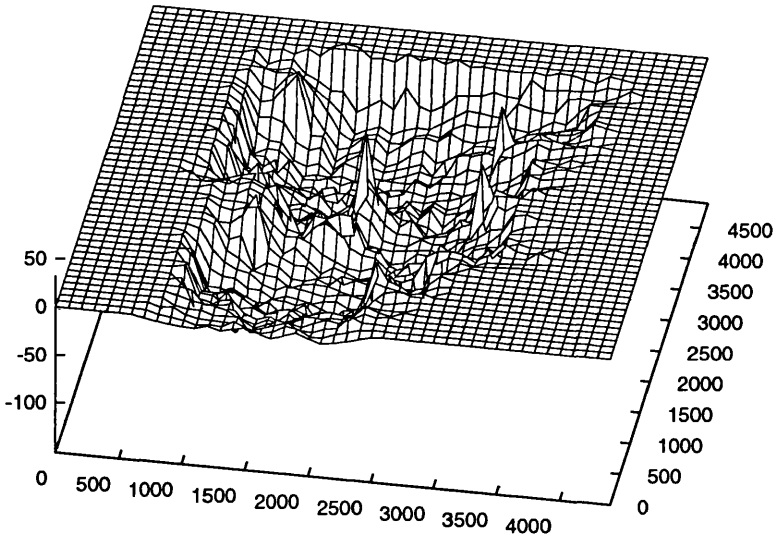

Figure 9.4: The mean-field approximated score for a single edge feature, calculated after 50 timesteps of exploration in the walls environment. The grid spacing is 10cm. At this resolution, the peaks of this function are very sharp, often only spanning 1 grid point. However, every peak still has a local maximum on the grid.

greater than or equal to r_α, or greater than r_α, respectively. The calculation of these numbers is given by a simple formula, derived in appendix C.

9.3.3 Searching for a single feature

We now have a formula which depends on the position of a single feature, and whose local maxima correspond to plausible features. We now need to find these maxima.

Edge features

In the case of edge features, we must find all local maxima of a function on a two dimensional space. To illustrate the difficulty of the problem, figure 9.4 shows a 3d plot of the mean field score against edge position, given a set of readings collected from 50 different positions within the walls environment. Figures 9.5, and 9.6 show close-ups of individual local maxima. The peaks of

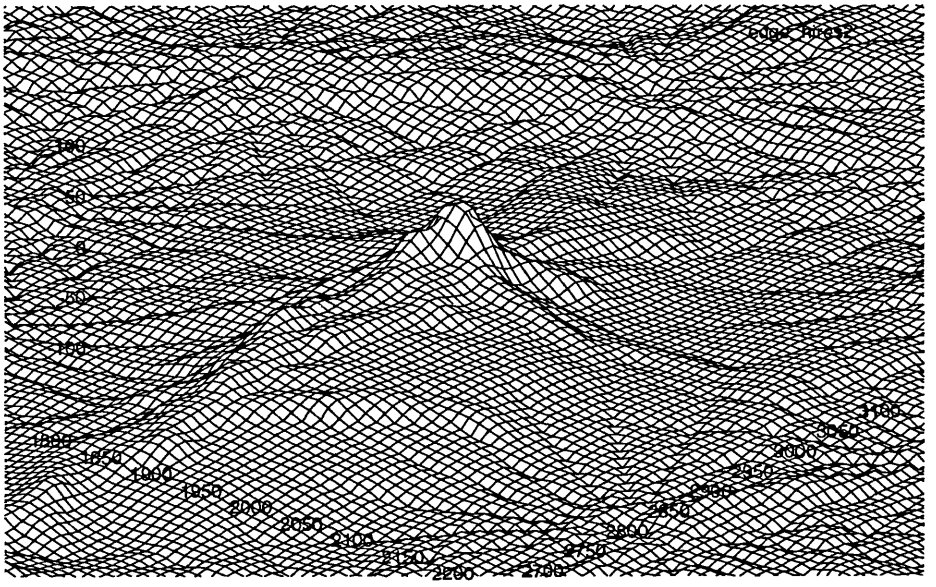

Figure 9.5: A close-up of the mean-field score landscape, showing the peak found for the edge reflector corresponding to the end of the cardboard wall in the middle of the room. The grid spacing here is 1 cm, 10 times higher than the exhaustive search resolution used in the program. But the program can find the location of this peak accurately by starting a local search from the position of the local maximum found on the 10cm resolution exhaustive search.

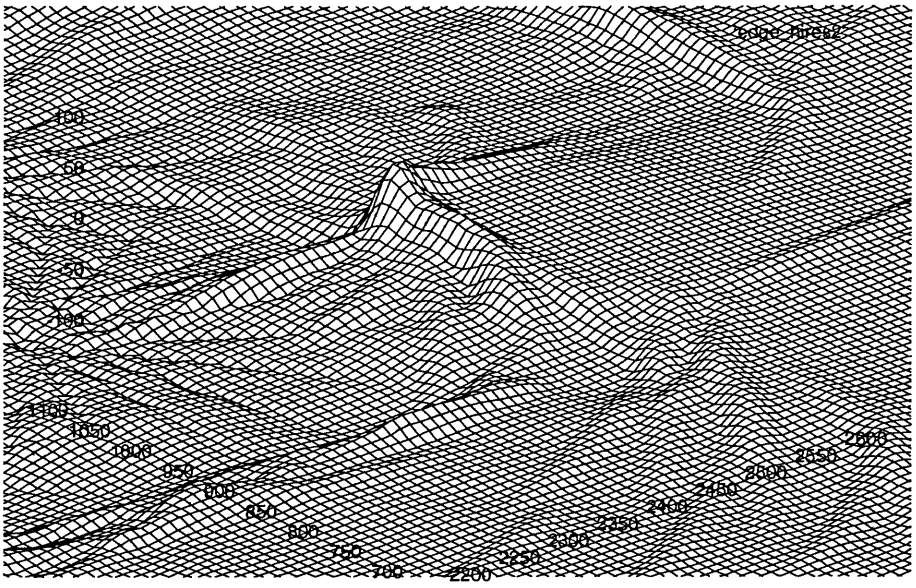

Figure 9.6: Another close-up, this time centred on a corner. The ridges correspond to walls. The lowlands to the left of the peak corresponds to unoccupied areas of the environment, where any features would have a low score because they would contradict a lot of readings that passed through that area.

this function are very sharp. In a way this is a good thing, because it means the peaks of this function give a very precise estimate of the position of the edge reflectors in the environment. But when it comes to searching for the positions of these peaks, it causes problems.

If we were to use a simple grid-based search to find the local maxima of this function, we would have two problems. The first is that we may miss peaks altogether, if the peak is so sharp that it falls in between the grid points. The second is that, even if we find the peak, the grid point that has the highest value may be only half way up the peak. In other words, if we want to accurately estimate the peak position, we need an extremely high search resolution. To have a search resolution this high is not practical in terms of the computer power available at the moment, and we *do* need this accuracy, in order for the method of combining features described in the next section to work.

We solve the problem as follows: first we do an exhaustive search on a grid

detect the exact position of the peak tip. Then we find all grid points that are local maxima of the function on this grid. And from each one of these points, we start a local search to find the position of the peak to the accuracy we require. For the local search, because we don't have gradient information, we used the *simplex method* (Press et al. 1993, section 10.4). We took a fractional tolerance of 0.05, and the search was abandoned if it went on for more than 50 iterations. Most searches took 20-30 iterations to find the maximum. This is not the fastest local search algorithm by any means, but it is reliable and easy to implement where no gradient information is available.

Wall features

To find the location of plausible wall features, we must find the local maxima of function on a four dimensional space. This is a worrying prospect, because to exhaustively search a four-dimensional space, even at low resolution, would take a very long time. However, because of the simple model we took for the effect of wall endpoint positions (section 9.3.1), it is possible to reduce the dimension of the space that must be exhaustively searched to 2.

To reduce the search dimension, we use a new coordinate system to describe the position of a wall. This coordinate system is shown in figure 9.7. Readers familiar with the Hough transform, a common tool in image processing, will recognise this way of parametrising lines. We use two variables, r_h and θ_h, to describe the slope and intercept of the line, and a two more variables, l_1 and l_2, describe the positions of the endpoints of the line.

The function $SC_{mf}(f_0; \{r_\alpha\})$ depends on the position of the feature f_0 through the parameters p, μ, and σ of the probability distributions of $X_{0\alpha}$, for each reading α. These parameters depend on the position of f_0 through d, the perpendicular distance from the sensor position to the wall, and θ_b, the beam incidence angle. The crucial point is that d and θ_b depend only on the slope and intercept of line f_0, as described by the coordinates r_h and θ_h, but do not depend on the endpoint positions l_1 and l_2. So when we are doing the search

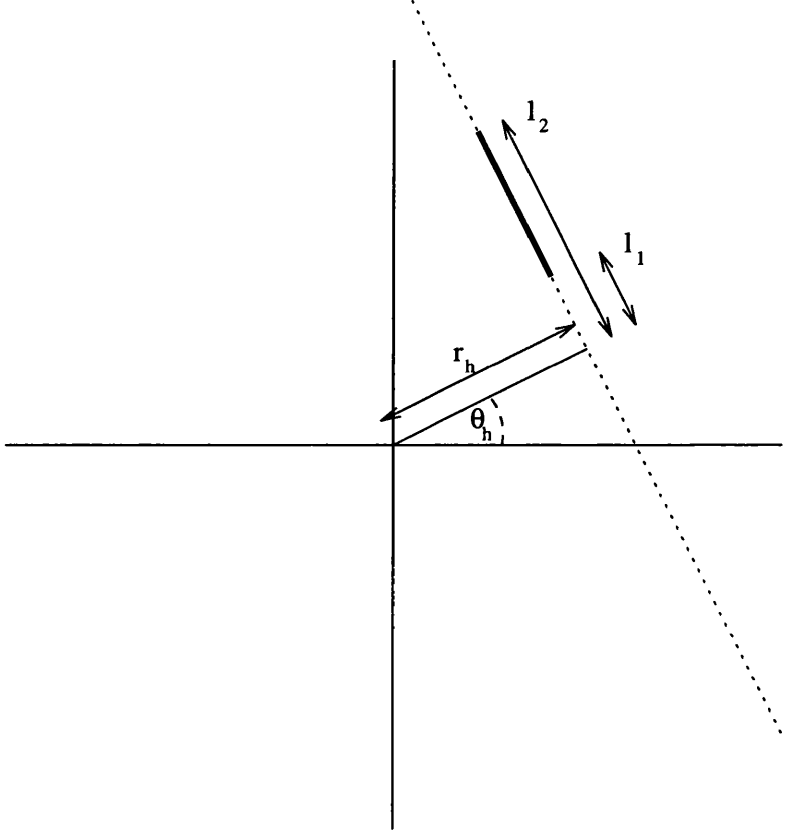

Figure 9.7: The coordinate system used to describe the position of wall features. The coordinates r_h and θ_h describe the infinite line of which the wall is a subsegment in Hough coordinates. l_1 and l_2 give the positions of the two endpoints, relative to the point of closest approach to the origin.

over all lines, we don't need to recalculate p, μ, and σ for every line – we need only calculate them for each value of r_h and θ_h. Calculating these parameters is what takes most of the search time, so this gives us a vast increase in speed.

Because of the simple model of section 9.3.1, the effect of the wall endpoint positions l_1 and l_2 is simply to determine which readings α are included in the sum of equation 9.6. To find the score for any value of l_1 and l_2, we use the method illustrated in figure 9.8. We define two functions $f(l)$ and $g(l)$ of the lateral distance l along the wall. The first function, $f(l)$, illustrated in figure 9.8a, is the sum of the scores of all readings for which the left intersection point of the sonar beam lies to the left of l. The second function, $g(l)$, illustrated in figure 9.8b, is the total of the scores of all readings for which the right intersection point of the sonar beam lies to the left of l. The score is then

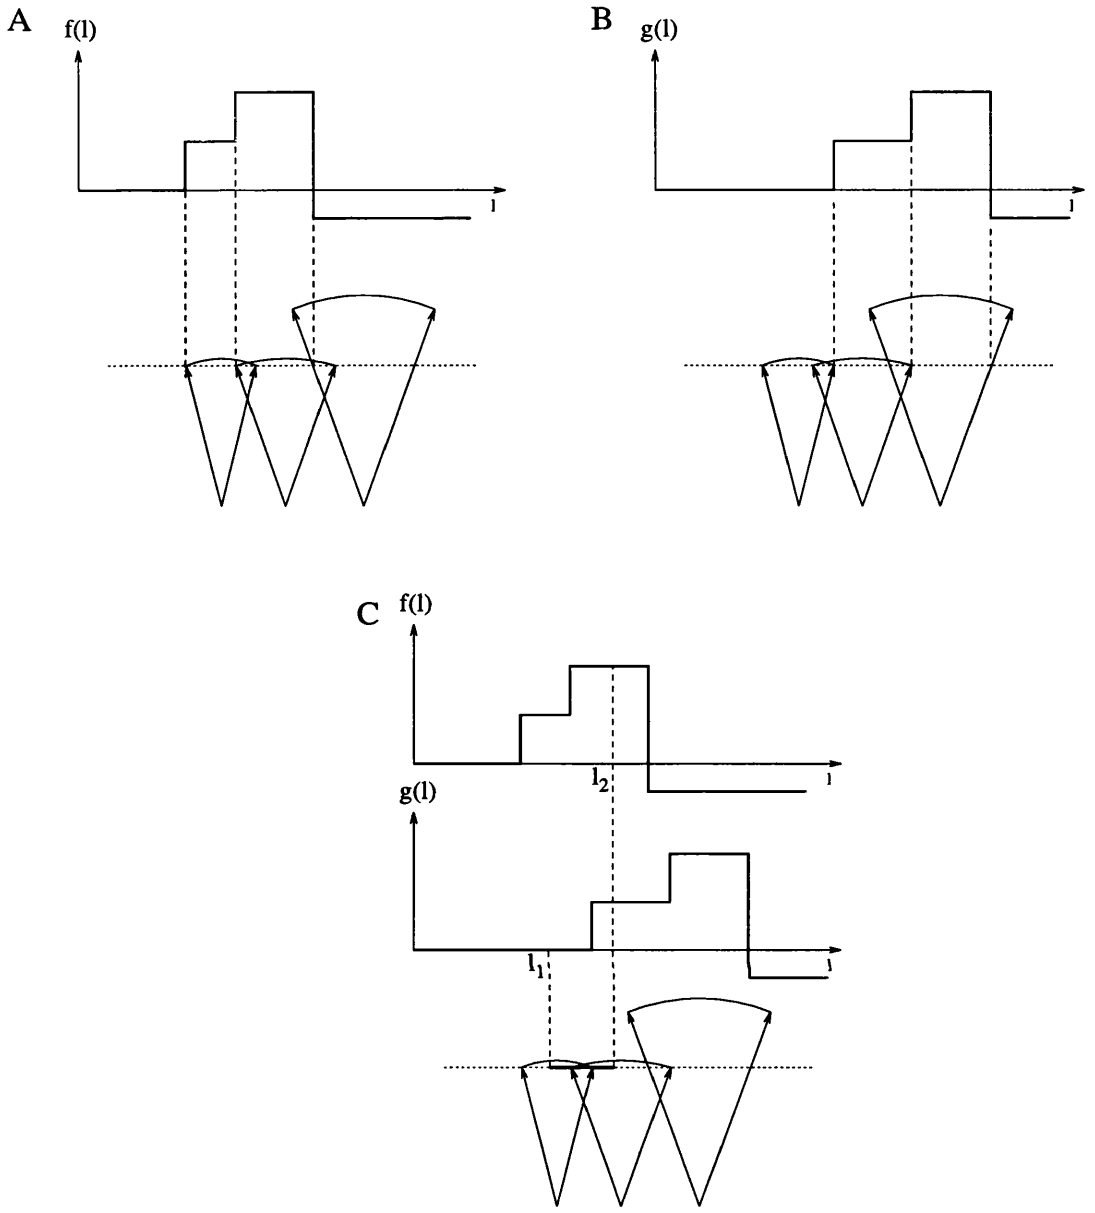

Figure 9.8: The method used to speed up the computation of scores for wall features. The score $log\,[A_\alpha P(X_{0\alpha} \geq r_\alpha) - B_\alpha P(X_{0\alpha} > r_\alpha)]$ is computed and stored for each reading α. For the (fictitious) readings shown here, the left two have a positive score, and the right one has a large negative score. Two functions of l are defined: A) the function $f(l)$ is defined to be the sum of scores for all readings whose left beam intersection point lies to the left of l. B) the function $g(l)$ is defined to be the sum of scores for all readings whose right intersection point lies to the left of l. C) The score for a given line segment is given by $f(l_2) - g(l_1)$.

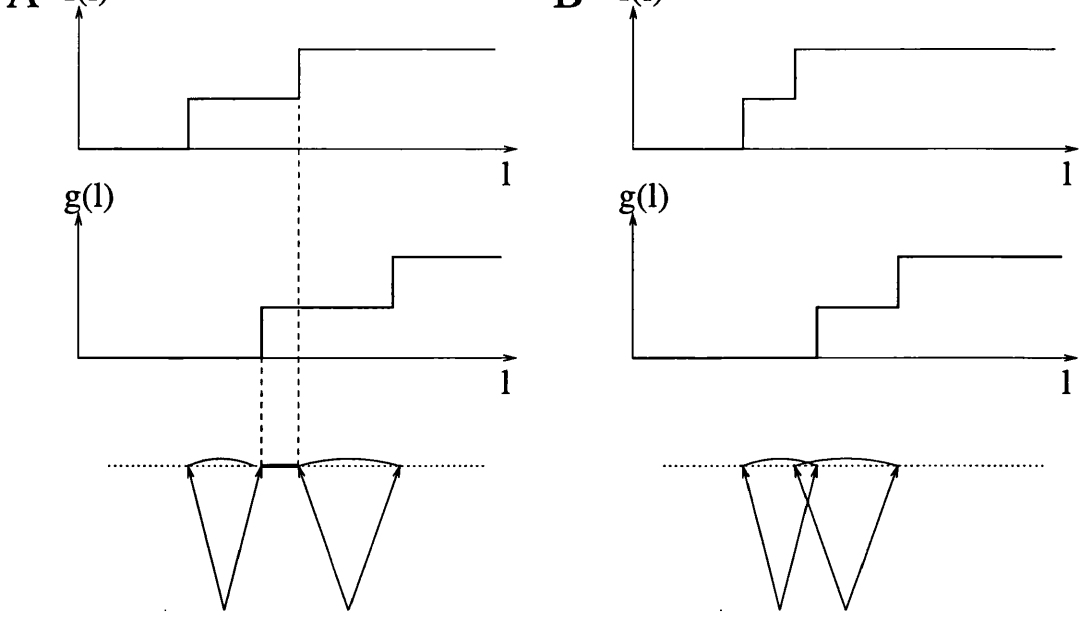

Figure 9.9: Detection of local maxima in l_1 and l_2. Local maxima occur when $g(l_1)$ and $f(l_2)$ both take a step up, provided that $l_2 > l_1$. A) An example of a local maximum. B) An example of a case where there is no local maximum.

calculated as $f(l_2) - g(l_1)$.

Now let us consider how, for a given infinite line, we can find the optimal values of l_1 and l_2. As $f(l)$ and $g(l)$ are piecewise constant functions, we have some freedom of choice for endpoint positions. We choose to make the detected walls as short as possible, and so take the left wall endpoint at the right of a flat section of $g(l)$, and the right wall endpoint at the left of a flat section of $f(l)$.

We find all local maxima by looking for those pairs of endpoints where the left endpoint l_1 corresponds to a step up for $g(l_1)$, and the right endpoint l_2 corresponds to a step up for $f(l_2)$, subject to the constraint that $l_1 < l_2$. A couple of examples of this are shown in figure 9.9. In figure 9.9a, there is one case where the criteria are fulfilled, and a single wall feature is detected. But in figure 9.9b, there are no endpoint pairs that fulfil the criteria, because all the step up points for $g(l_1)$ lie to the right of those for $f(l_2)$. In this case, a wall feature will not be detected.

One might worry that a wall feature should be detected for the case of

edge feature. In the case of figure 9.9a, however, the two readings are too far apart to be explained by a single edge feature.

We must make one more approximation in order to reduce the exhaustive search to 2 dimensions. For every infinite line, we are going to only consider one choice of endpoints, those that give the maximum score possible. We find these endpoints by exhaustive search over the local maxima of l_1 and l_2. Because this is a small, discrete space, the search is very fast.

The final algorithm for finding local maxima in the space of wall features is as follows: we exhaustively search through the space of all infinite lines (with a grid spacing of 10cm for r_h and 30^o for θ_h). For each each infinite line, we store the best choice of finite sub-line, and its mean field score. We then find all grid points that are local maxima of the stored best scores. For every such point, we start a local search in the space of infinite lines, using the simplex method, to find the exact location of the local maximum. Again we use a fractional tolerance of 0.05, and searches taking more than 50 iterations are abandoned. For each local maximum produced by the simplex method, we produce a single finite wall, choosing the endpoints by the method described above.

9.3.4 Combining features

The method described in the previous subsection gives us a set of local maxima of $SC_{mf}(f_0; \{r_\alpha\})$, which correspond to plausible features in the environment.

What we want, though, is a *set* of features that, together, give a good score for $SC_{map}(\{f_i\}; \{r_\alpha\})$. We will find one by searching the space of all maps consisting combinations of the plausible features found by the mean-field approximation. This is a discrete space, and so we now have to solve a discrete optimisation problem.

One might wonder why we cannot just take a map consisting every single local maximum of $SC_{mf}(f_0; \{r_\alpha\})$. But such a map would not necessarily be a maximum of $SC_{map}(\{f_i\}; \{r_\alpha\})$. In practical terms the problem is that the

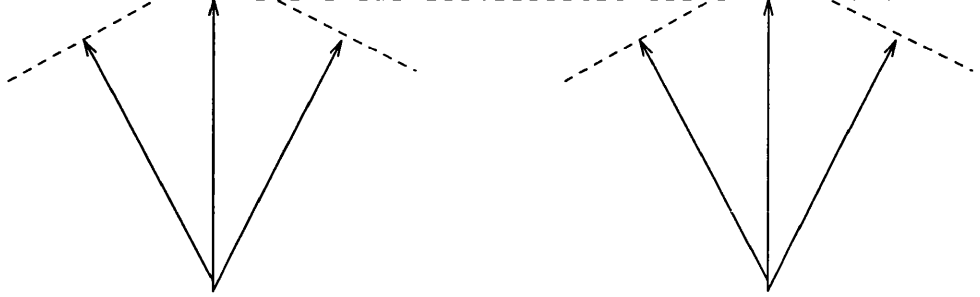

Figure 9.10: For a given set of sonar readings, there will be many local maxima of the mean-field approximated score. Not all of these will correspond to actual features in the environment. In the example shown here, all the dotted lines correspond to local maxima of the mean-field feature score. However, the long horizontal wall is enough to explain all of the readings. The map score will be higher if only this one feature is included, than if they all are.

mean field approximation finds too many features. In figure 9.10, for example, are shown 5 wall features that are local maxima given a set of 6 (fictitious) sonar readings. The total map score will be lower if all of these features are included than if just the horizontal line is. We need a way of selecting the subset of the plausible features that gives us the highest score.

Many advanced methods for discrete optimisation problems exist in the literature (see e.g. Reeves 1993). But we will use a very simple hill-climbing method. We start with an empty map, and add features to it iteratively. On each iteration, we loop through the plausible features, and calculate the change in map score that would be gained by adding that feature. If the best feature increases the map score, we add it and repeat the iteration; otherwise we stop.

One might worry that this hill-climbing search method will get stuck in local maxima. But the score landscape is so simple in this case that it does not seem to happen.

The calculation of scores can by dramatically sped up by caching the values of $P(X_{i\alpha} \geq r_\alpha)$ and $P(X_{i\alpha} > r_\alpha)$ for every feature i and reading α. The computational complexity of our search method is quadratic in the number of plausible features, but with this caching method is very quick to execute.

The algorithm as described above uses *all* observations made by the robot to build a map. In theory this is will make for very accurate maps, but in practice there are two problems. The first problem is computational complexity. The amount of time taken to build a map is linear in the number of observations used to build it, so if we use all observations then execution time will scale without limit. The second problem is odometry errors. Incremental localisation will reduce these errors, but if the robot revisits a previously-explored part of the environment it is better to use the more recent readings and disregard the old ones, which are more likely to be afflicted with odometry errors.

To overcome these problems, we used an iterative technique to filter the readings used to make the map. A "filtered list" of readings is initialised to be empty. The observed readings are examined in reverse order, starting with the most recent. For each reading considered, if there is no reading in the filtered set taken from a viewpoint within 20cm of the current one, the current reading is added to the filtered list, otherwise it is ignored.

In this way we produce a filtered list in which no two readings are taken from viewpoints within 20cm of each other, and in which more recent readings will be included over older readings.

9.4 Quantitative assessment

In order to compare the results of our new mapping system to previous ones, we need a quantitative map quality metric. We will use the metric defined by Lee and Recce (1997), as described in section 4.5.6.

9.4.1 Generation of a free-space map

The map quality metric requires a grid-based free-space map. But our mapping algorithm only produces a list of features. We therefore need a system for making a free-space map out of a feature map.

We construct a free-space map on a 10cm grid using a system similar to

cell on the grid is marked as unknown. For every feature in the map we mark the cells in which the feature lies as occupied. We then go through all the observations in the filtered list. For each observation, we find the nearest feature for which the probability of producing the reading according to our probabilistic model is greater than p_{min}, the probability floor limit used in the algorithm (10^{-4}). We then mark all points in the arc defined by the sonar beam for this reading, going up to the length of the nearest reflecting feature, as free.

Finally, all grid points that are within 1 robot width of an occupied cell and would otherwise be marked as free, are marked as dangerous, and all points that the robot has actually passed through are marked as free.

9.4.2 Localisation

The aim of this chapter is to describe and test a new mapping algorithm. But for this mapping system to work well, it also needs an (incremental) localisation system. Incremental localisation is a well-researched topic, and lots of techniques exist for its solution. For the experiments described below, we are using tracefiles of data collected previously to evaluate different exploration strategies using the mapping system of Lee and Recce. Lee and Recce's system contains an incremental localisation function. To save re-implementing it, we simply saved the output of their incremental localisation system in another tracefile, and used it as an input to our system.

9.4.3 Data sets used

Lee and Recce collected a large amount of sonar data, using several different exploration strategies, in several different environments. Here we are not concerned with the efficiency of different exploration strategies – instead we would like to compare the efficiency of our mapping system to that of Lee and Recce. So the data we are going to use was all collected using the same exploration strategy – simple wall following. We will test the system on data collected in 3

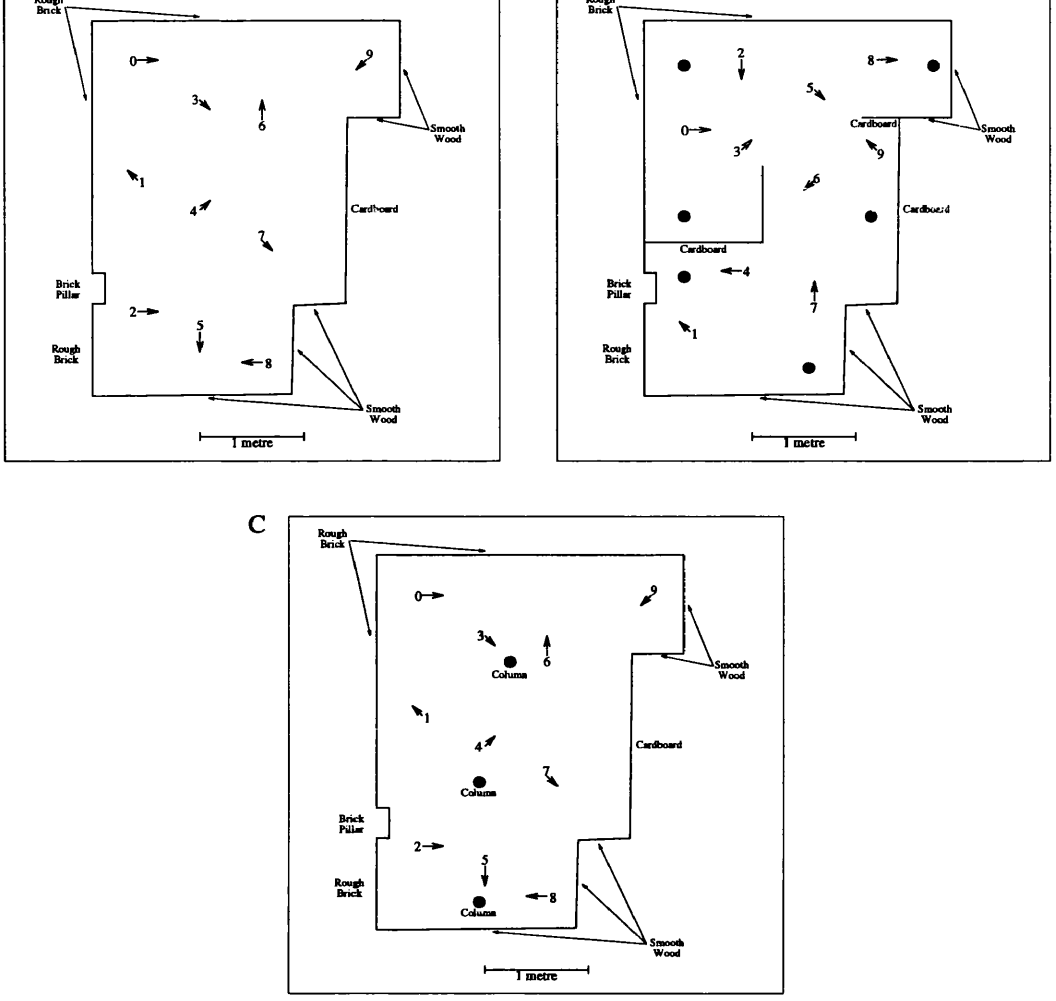

C

Figure 9.11: The test environments where sonar readings were collected to test the mapping system. 10 runs were executed in each environment, from the starting points shown, using a simple wall-following exploration strategy. A) The "empty" environment. B) The "walls" environment C) The "columns" environment.

different environments – the "empty" environment, the "walls" environment, and the "columns" environment, shown in figure 9.11.

As you might expect, the empty environment is an empty room. But it is not of a regular shape, and the different walls are made of different materials of differing degrees of specularity. The columns environment is the same room but with three cylindrical cans placed in it. The walls environment contains six cylindrical cans, and also cardboard dividing walls.

In the empty and columns environments, the whole room is visible from

ment is visible from any one viewpoint, so the robot must explore thoroughly before being able to map the room.

For each of the rooms, we evaluated the performance of the mapping system using trace files consisting of the results of 10 different wall-following exploration sessions from different starting points.

9.5 Results

9.5.1 Example maps produced by the algorithm

Figure 9.12 shows the features detected by the new mapping algorithm, and the old algorithm of Lee and Recce, after one, five, and forty sonar scans the walls environment.

The new algorithm has detected the north wall after just one sonar scan. After five scans, the new algorithm has detected the north and west walls, two of the columns, and two point reflectors corresponding to wall endpoints. The old algorithm has not detected the west wall, the northwest column, or the central endpoint, but has detected another column that the new algorithm has missed. The new algorithm has also produced a spurious feature at this time. However, this feature is on the edge of the area that has been visible to the robot until this time, and therefore does not result in a significant lowering of map quality.

After 40 timesteps, the maps produced by both algorithms look substantially similar. However, the map produced by the new system has a much higher quality metric. This is because of the way Lee and Recce's map quality metric works. The metric is the fraction of all test journeys that can be completed successfully on the basis of the current map. In the walls environment, a large fraction of the test journeys either go through a narrow passageway, or must go past a collision "hot spot" (Lee 1996, section 13.3.1). As a result, there are certain critical map features that, if even slightly mislocated, will result in much lower map quality. Careful examination of figure 9.12 shows

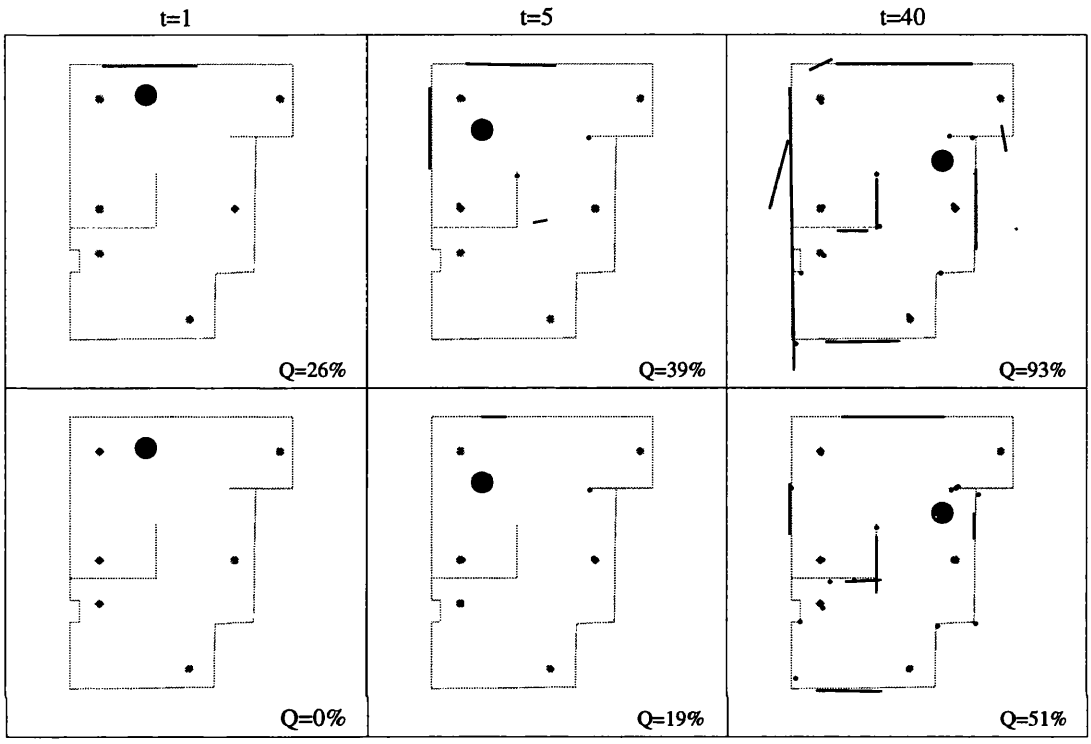

Figure 9.12: Some example maps built from data collected in the walls environment, after 1, 5, and 40 timesteps. The top row shows the output of the new mapping system, and the bottom row shows the output of the old system of Lee and Recce on the same data. The light lines denote the actual feature locations, and the dark lines and dots represent features detected. The large dark filled circle represents the robot. At the bottom right of each map is map quality, using the metric of Lee and Recce.

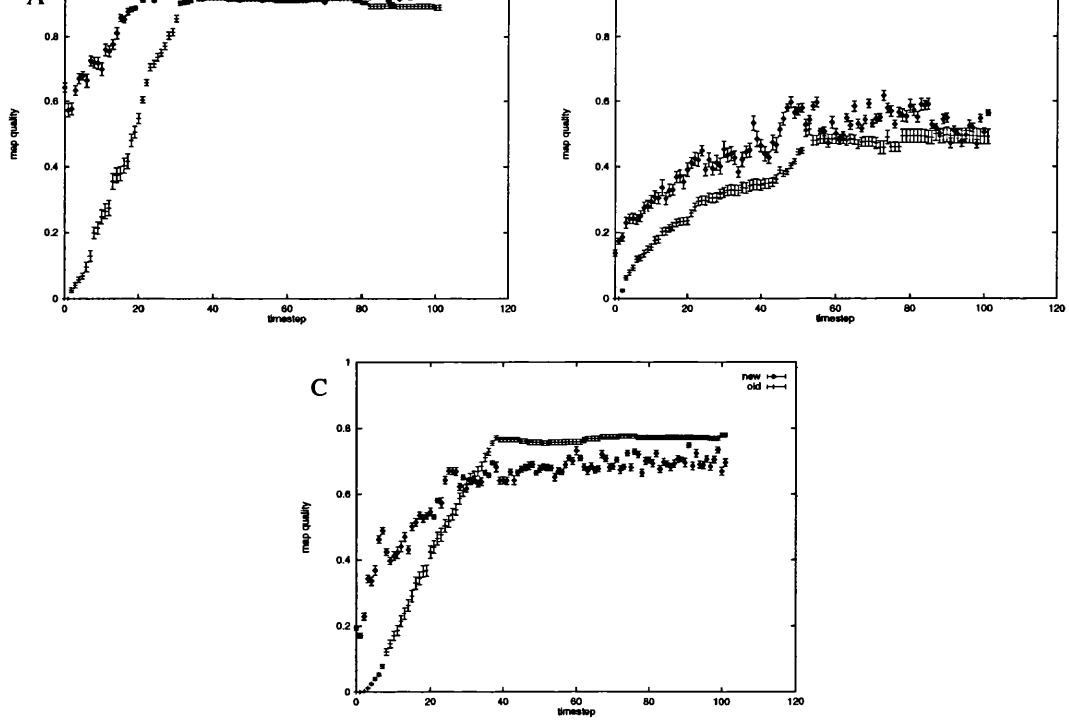

Figure 9.13: Average map quality against number of scans for the new method and the old method of Lee and Recce, in A) the empty environment, B) the walls environment, and C) the columns environment. The average is taken over 10 exploration runs from different starting points. The two algorithms were fed identical trace-file data. Error-bars denote standard error of the mean.

that, after 40 timesteps, the old system has slightly mislocated both ends of the central cardboard wall. This slight mislocalisation results in much lower map quality.

9.5.2 Quantitative comparison to previous method

Figure 9.13 shows a graph of average map quality against number of scans in the 3 environments, averaged over all 10 runs, for the new system and the system of Lee and Recce. In all cases, the new system gives a higher map quality in the early stages of exploration, confirming the trend seen in the specific examples of figure 9.12.

In the case of the empty and walls environments, the new algorithm produces higher asymptotic map quality, again confirming the specific examples of figure 9.12. But in the columns environment, the new algorithm produces

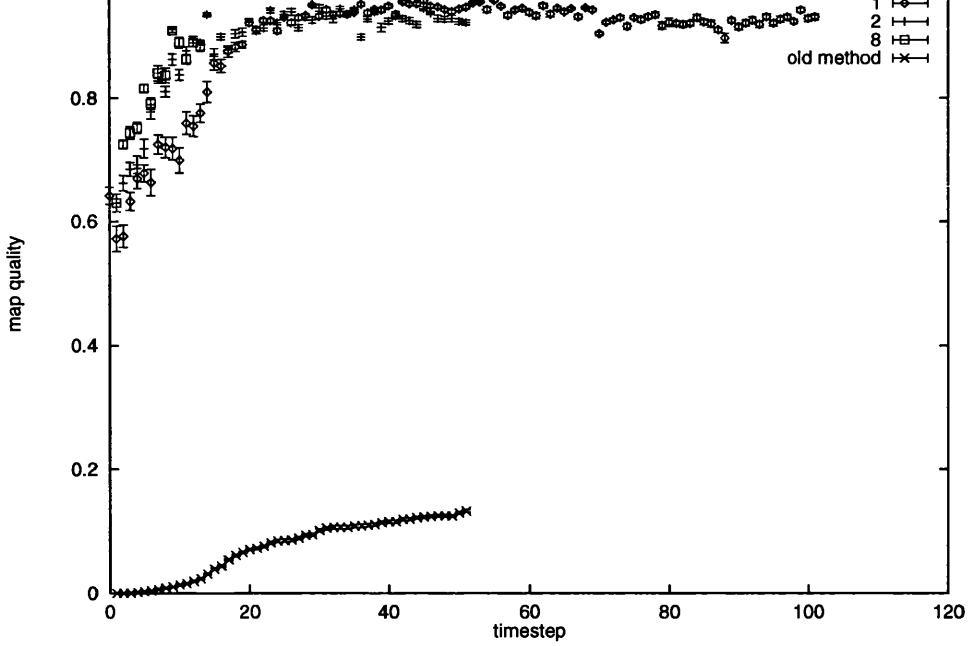

Figure 9.14: Subjunctive evaluation of effect of step size on map quality in the empty environment. The graphs show map quality against viewpoint number where only 1 out of n sonar readings has been used to construct the map, simulating step sizes n times larger than usual. Graphs are shown for n =1, 2, and 8. With the new algorithm, performance actually improves with increased step size. The bottom trace shows performance of the Lee and Recce system with step size of 2. Error-bars denote standard error of the mean.

a lower asymptotic map quality. Likely reasons for this will be discussed in section 9.6.3

9.5.3 Influence of step size on map quality

The Lee and Recce mapping system, as currently implemented, requires scans from three closely-space viewpoints in order to detect a single feature. As a result, the robot must take short steps in order to map an environment. It would be much quicker to map an environment if the robot could take large step sizes.

We investigated the effect of large step sizes on map quality subjunctively, by *decimating* an existing tracefile of data, that is, deleting all but 1 out of n scans recorded in it. Figure 9.14 shows the effect of n on map quality, for this system and the system of Lee and Recce. For the new system, performance

is much lower for step sizes greater than 1. The lower trace in figure 9.14 shows the map quality resulting from the Lee and Recce system, with $n = 2$, corresponding to a virtual step length of 60cm. At this step size, the algorithm does not detect a single feature. Consequently, the only areas marked as free are those that the robot has passed through, and the map has a very low quality.

9.6 Discussion

9.6.1 Improved performance in early stages of exploration

For all environments, the new mapping system gives a significantly higher map quality in the early stages of exploration than the old system of Lee and Recce.

Why might this be? Let us start by looking at the maps built after just one sonar scan, as shown in the left column of figure 9.12. The new algorithm has already detected a feature corresponding to the north wall but Lee and Recce's method has not detected any features. The Lee and Recce method, as currently implemented, will not add any features to the map unless they have been observed from three consecutive sensor positions. This vastly reduces the number of sensor readings that are actually used in map construction. The new method uses all readings, and will therefore detect features from less sonar data. The difference between our map quality and Lee and Recce's goes down with time because, given time three consecutive readings will be seen for each feature.

As seen in figure 9.13, the increase in map quality for both mapping algorithms is slower in the walls environment than in the other two environments. The reason is that, in the empty and columns environments, wherever the robot is, it has an unobstructed view of the entire room. But in the walls environment, only half the room is visible from any one viewpoint. Non-averaged quality plots in the walls environment, show step-like increases in map quality,

they occur at different times on different exploration runs, the averaged plot shows a slow gradual increase.

9.6.2 Improved asymptotic quality in the walls and empty environments

In the walls and empty environments, asymptotic map quality is significantly higher with the new system than with Lee and Recce's system.

The most likely reason for this is the way feature positions are determined by the two systems. As reviewed in chapter 6, the way Lee and Recce's system calculates feature positions is to take the positions of all sonar returns that were ascribed to a particular feature, and calculate the feature position by producing, in the case of edge features a point corresponding to the average reflection point, and in the case of wall features a line by linear regression.

This has two drawbacks compared with our method. The first is that they have to ascribe every observation to a single feature (the correspondence problem). If a reading is wrongly ascribed, map quality will be lowered permanently. But in the new system, any observation can contribute to the positioning of any feature, and there is no correspondence problem. The second drawback is the *ad-hoc* method of finding feature positions. The new method, being based on an optimality principle derived from an empirical probabilistic model, localises features much more accurately, and so produces higher map quality.

For both mapping systems, the asymptotic map quality in the walls environment is much lower than in the empty and columns environments. The reason for this is the extreme sensitivity of the quality metric to the position of certain critical features, described in section 9.5.1. While the asymptotic quality of the new method in the walls environment is in general higher than the old method in the walls environment, it is worth noting that it is still lower than the asymptotic quality of either method in the empty environment. This indicates that the new method is more accurate than the old one when

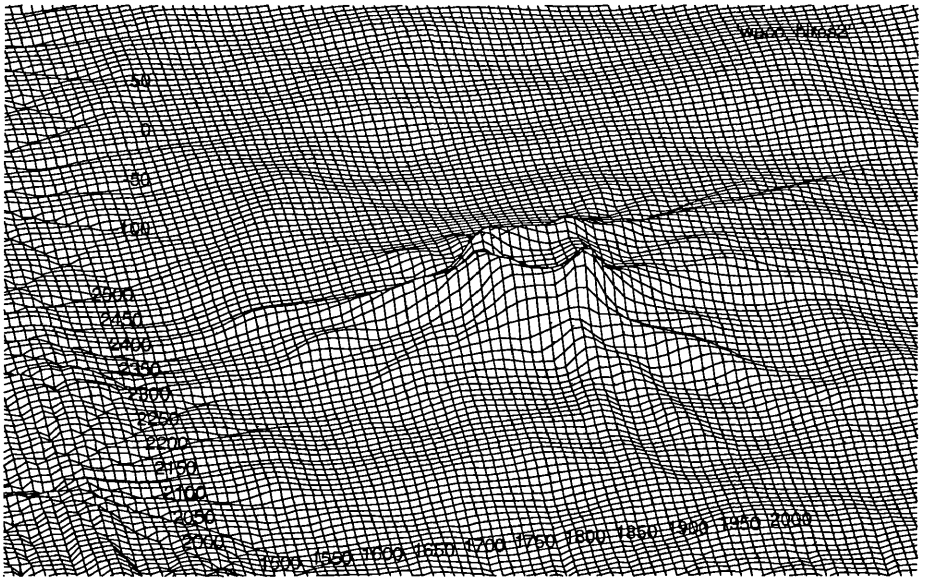

Figure 9.15: A close-up of the mean-field score for edge features in the neighbourhood of one of the columns in the columns environment. The score is highest on a circular rim around the edge of the feature, leading to mislocalisation of the feature.

localising these critical features, but is not always accurate.

9.6.3 Worse asymptotic quality in the columns environment

Why does the new method produce worse asymptotic quality in the columns environment? The answer is to do with the actual columns themselves. These columns are cylindrical cans, about 10cm in diameter. But our model of edge reflectors was derived from two objects – the corner of a metal filing cabinet, and the edge of a metal meter rule, which both act like acoustic knife-edges. This is what Kuc and Siegel originally had in mind for edge reflectors. However, the 10cm columns do not act like acoustic knife-edges and so the new algorithm does not do as well.

Figure 9.6.3 shows a close-up of the mean-field score for edge features in the vicinity of a column. There is no local maximum where the column is; instead there is a "volcano crater", corresponding to the edge of the cylindrical column. The new mapping system sometimes gets the location of the columns

is. This mislocalisation leads to a lower score.

One could imagine of two ways to get around this problem. One would be to add another type of feature to the existing mapping system, corresponding to 10cm cylindrical columns. However, this is not an ideal solution, because if we take this approach, we must manually construct a probabilistic model for every type of environmental feature the robot must deal with.

A better solution is for the robot to learn its own models of the reflectivity of environmental objects. Possible methods for implementing this are discussed in section 9.7.2.

9.6.4 Suitability of Bayesian methods

The mapping system works by finding the map which, given the observed readings, has highest *a posteriori* probability. However, a map that is optimal in this sense may not be optimally useful to the robot. The map quality measure of Lee and Recce, which we are using to evaluate performance, is particularly sensitive to the positions of obstacles which the robot must navigate around, like the cardboard wall in the walls environment. If such an obstacle is mislocalised by a small amount, but enough to put it in the wrong grid segment, the map quality can go down by as much as 50%. So, a system optimised to give the highest quality, according to Lee and Recce's measure, might care more about the location of this feature than others, which the Bayesian score does not.

Another problem with all statistical methods is the possibility of overfitting, i.e. modelling observed distributions in too much detail, rather than capturing the underlying causes of those observations. This can be seen in the sample maps of figure 9.12. The map at t=5 contains a small wall feature in a region of empty space in the centre of the room, and the map at t=40 contains three extraneous wall features, one of which is actually outside the environment! This overfitting does not lead to a decrease in map quality, however, because the spurious features are always located at the edges of

features in the centre of a well-explored area because these features would contradict many sonar readings and so drastically lower the overall map score.

9.6.5 Execution speed

The space of maps is very large, of variable dimension of up to 400. Searches of spaces this large are very hard. We have found an approximate way of searching this space, which gives an approximate solution in a tractable amount of time. But execution speed is still a primary concern. The complexity of the search method is linear in the number of observations used to construct the map, and we use a filtering method to cut it down further to linear in the total area explored.

We used several other tricks to decrease computation time. In the first stage, we performed exhaustive searches on a low resolution, and began a local optimisation from each local maximum located at the low resolution. We also used a particular model of the effect of wall endpoints on reflection because it allowed for faster searches of the space of wall features. In the second stage, we saved time by caching calculated probabilities. This used up to 10 megabytes of memory, but for modern computers this is not a problem.

To map an environment of the size we considering here took up to one minute of processor time on our machine (SPARC Ultra, 296 MHz). This is not fast enough to do as the robot explores in real time, but as machines become faster and the algorithm becomes more fine-tuned, execution speed will not be a problem.

9.6.6 Ramifications for choice of exploration strategy

Lee and Recce's (1997) primary purpose was quantitative evaluation of exploration strategies. It would be interesting to see if repeating their experiments using the new mapping algorithm would lead to different results.

One reason to suspect that it might is that the new system can build maps with any step length, whereas the original Lee and Recce system suffers a

$30cm$. Large step sizes, however, can lead to localisation errors, which may also prove to be an important factor in determining the optimal exploration strategy.

9.7 Directions for future work

The work of this chapter has demonstrated that mapping by Bayesian search is practical in heterogeneous environments. But there are still many avenues of research which could lead to an improved method. We will list a few here.

9.7.1 Localisation

The system we described in this chapter is a mapping system, with no associated method of incremental localisation. We have relied on a previous system to give us the localisation information that we need to build accurate maps. It would be nice if we could integrate the localisation functions into this mapping system. As the system here is derived in a Bayesian framework, a sensible approach would be to consider the positions of the robot to be random variables. The mapping system would then combine the search for optimum feature positions with a search over the space of robot trajectories.

Thrun et al. (1998) have described a technique for concurrent mapping and localisation which performs a search over the joint space of trajectories and maps using an "E-M" algorithm (Rabiner and Juang 1986). Such an approach may well work in the current context.

9.7.2 Learning

In theory, it would be quite straight forward for the robot to build up its own probabilistic models of the reflectivity properties of different environmental objects. Once the robot has a map, it would then have to solve the correspondence problem – i.e. decide which sonar readings were caused by which features – and then derive a probabilistic model. There are plenty of learning

neural networks.

For this to work in practice, two problems would have to be overcome. The first is a "chicken and egg" problem. The robot would not be able to build a map without a sensor model, but the model cannot be learned without a valid map. This could be overcome by starting with an approximately correct model, which is then corrected as time goes by.

The second problem is how to perform the correspondence. One possibility would be to use the current sensor model to perform the correspondence, and ascribe a observation to a feature if the probability that the feature produced the observation under the current model exceeds a fixed threshold. However, if the current model is non-optimal, readings may be allocated incorrectly, and the model may deteriorate. A better possibility would be to use some rough spatial criterion, and say the observation belongs to the feature if the sonar beam was pointed roughly towards the feature and it gave roughly the distance to the feature. In effect, this is what we did in chapter 8: we classed a return as direct if the measured distance was within 10cm of the actual distance, and built the model from the direct returns.

If the mapping system was capable of learning, it would be able to deal with novel objects without needing an explicit reflectivity model to be empirically determined. This would not only apply to the cylindrical pillars that caused lower asymptotic map quality in the columns environment, but also to more complex objects like chairs. Ideally, the robot would detect when a new feature type is found, and make a new category automatically.

9.7.3 Relation to neural information processing

The problem we are dealing with here, that of object recognition and localisation from sensory input, is a central one to the brain. In this section we will discuss some previous work in the field of neural information processing which could be relevant to this case. The exact form of the problem is different here, because, except for bats, animals do not have sonar sensors. But one

might expect that some general principles about how the brain solves this kind of problem might be relevant to this instance. Most of the previous research described in this section has been to do with processing of visual information, but the information processing strategies may well generalise to the current case.

The sonar mapping task we are concerned with here involves extracting a variable number of features from sensory data. This is different to the usual pattern classification task that most neural networks are set. We will now describe some work in the field of neural processing that is concerned with extraction of multiple features.

The temporal binding hypothesis

The temporal binding hypothesis is a theory for how the brain might represent the nature and location of several objects. One way this might be done is with a "grandmother cell" approach, in which a single neuron represents the presence of a certain object in a certain location. However, this would lead to a problem of combinatorial explosion, because the total number of cells required would be the product of the number of feature types and the number of locations.

This has led to the *temporal binding hypothesis* (see e.g. Singer and Gray 1995). In this scheme, there are two sets of neurons, one set coding for the location of a stimulus, and one set for its nature. The neurons from both populations that are relevant to a given feature fire synchronously to describe its position and location, and two different objects are represented by synchronous firing patterns at different times. In this scheme, the total number of neurons needed is equal to the sum of the number feature types and the number of locations. It has been proposed that the pacemaker for temporal binding is the 40Hz gamma rhythm observed in cortex, so that neurons that code for different features are active on successive gamma cycles.

This idea has some experimental support. The main support comes from evidence that separate brain areas are involved in representing the nature and location of stimuli (Ungerleider and Mishkin 1982), and also from experiments

levels of synchronisation in two neurons involved in the representation of a single feature, above two neurons representing separate features (Gray et al. 1989).

To my knowledge the temporal binding hypothesis, in its original form, has not yet been translated into a specific computational network model.

The model of Wiskott and Von der Malsburg

C. Von der Malsburg was the first to suggest that coherent firing of neurons may serve a representational purpose in the brain. Recently, he has given some specific computational network models that instantiate this idea. But, in their detail, they are quite different to the original temporal binding theory described above.

Wiskott and von der Malsburg (1993) describe a model that is relevant to the problem here, of detecting and locating multiple features from (visual) sensory data. The model is not instantiated as a neural network per se, because of implementational constraints on a serial computer, but is inspired by a model of visual cortical function. It works as follows: all input images are preprocessed by wavelets. For every feature that is to be recognised, a template is stored of the wavelet transformed image. Then, when a new image is presented for analysis, the stored templates are matched to it. Every template is matched, at every possible position, and a best fit position is obtained for each feature.

An iterative process is then employed to determine which of the stored features are actually present in the input image. On every cycle, a list of candidate features is stored, corresponding to those features scores exceeding a threshold (which depends on the number of features already accepted). Then, an analysis is performed to see which of the candidate features is nearest the front of the image, i.e. is not occluded by other candidate images. This feature is then accepted. The iteration then proceeds again, terminating when no features qualify as candidates.

There are some differences between this model and the model presented in this chapter, the main one being that theirs is a model of visual feature detection, and ours uses sonar. But structurally, the algorithms are very similar. In particular, their model uses a search procedure to find the locations of plausible features, and then uses an iterative process to decide which of them are present in the input image.

The model of Ambros-Ingerson *et al.*

Another relevant model is that of Ambros-Ingerson et al. (1990). This is a model of olfactory pattern recognition, based on a simulation of the olfactory bulb and olfactory cortex.

In this model, sensory input arrives at the olfactory bulb, where it activates a set of mitral cells. These cells then excite pyramidal cells in the olfactory cortex. The pyramidal cells in cortex are arranged in groups, with each group having a winner-take all mechanism of reciprocal inhibition. The winning cortical cells fire, and by connections to inhibitory interneurons in the olfactory bulb, inhibit the set of mitral cells that would optimally excite the winning cortical cells. The mitral cells which continue to fire after this will then cause a different of cortical cells to fire.

In this way, the network performs a hierarchical clustering operation. The set of cortical cells firing on the first timestep corresponds to the first level of the hierarchy on which the input odours are classified. The mean vector for this cluster is then subtracted out of the input, and the remaining pattern is activates a cortical activity pattern on the second timestep that corresponds to the subcluster that best describes the input.

Although this network is presented as a method of hierarchical clustering of odour stimuli, it is actually performing a similar function to the multiple feature mapping we described above. In our system, the feature with strongest support is detected, and then the vectors A_α and B_α are recalculated. Those readings which are "explained" by the detected feature have B_α set to be small, so they are less likely to be used as evidence for future features. This is similar

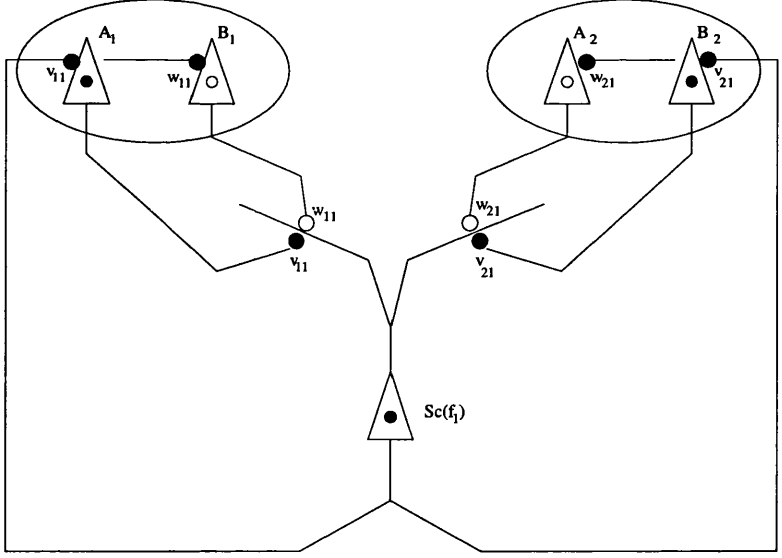

Figure 9.16: A neural network capable of implementing the mapping system described here, that functions on a similar basis to that of Ambros-Ingerson *et al.* . For each input reading there is a column which contains one excitatory and one inhibitory cell. These project to a population of feature cells (only one is shown), via excitatory and inhibitory weights whose values correspond to the probabilities $P(X_{i\alpha} \geq r_\alpha)$ and $P(X_{i\alpha} > r_\alpha)$. The feature cells are connected together by a mutually inhibitory winner-take-all mechanism (not shown). An essential feature of the network is the non-linear interactions between the dendrites of the feature cells (see text for further details).

to the subtractive inhibition of Ambros-Ingerson et al.

Possible neural implementation of the model of this chapter

The preceding paragraphs have described neural models that operate along similar principles to the mapping algorithm described in this chapter. However, there are some important differences. Firstly, the problems that these other models set out to solve are different – namely visual object recognition, and hierarchical clustering of odours. Secondly, these other models have come from a set of ad-hoc heuristics, but ours has come from a mathematical framework. It would be interesting to find a neural network that could implement our model, and to see in what ways the mathematical framework and specific problem domain force it to be different to those described above.

A possible neural implementation of the algorithm is shown in figure 9.16.

arranged in columns (shown enclosed by an ellipse) containing 1 excitatory and 1 inhibitory cell; and a set of map cells, which contains one cell for every possible feature (only one shown). Both of the cells in one column synapse onto the same dendrite of the feature cell.

The membrane potential on the excitatory cell in column α corresponds to the value of A_α, and that of the inhibitory cell corresponds to B_α. The excitatory cells have a real-valued output, which is multiplied by the value of the synaptic weight to produce a post-synaptic potential in a target map cell. The weights are fixed, with values $w_{\alpha i} = P(X_{i\alpha} \geq r_\alpha)$ and $v_{\alpha i} = P(X_{i\alpha} > r_\alpha)$.

The membrane potential on a feature cell i corresponds to a score of that feature, which is equal to the total map score with that feature added:

$$SC(f_i; \{A_\alpha, B_\alpha\}) = \prod_\alpha [A_\alpha P(X_{i\alpha} \geq r_\alpha) - B_\alpha P(X_{i\alpha} > r_\alpha)]$$

In order for the feature cells to compute this score, there must be non-linear interactions between the cell's dendrites. The potential on each dendrite is equal to the excitatory input to that dendrite minus the inhibitory input:

$$A_\alpha P(X_{i\alpha} \geq r_\alpha) - B_\alpha P(X_{i\alpha} > r_\alpha)$$

For the cell to compute $SC(f_i; \{A_\alpha, B_\alpha\})$, the required membrane potential for the soma must be the *product* of the potentials of its dendrites.

When a set of sensory data is collected, the columns for which the appropriate readings were detected are activated, with the potentials of both cells in the column equal to 1. These then cause the feature cells to compute the scores for each feature on the basis of all observed readings. The cells are connected by an inhibitory winner-take-all mechanism (not shown), so that, at any time, only the cell with the highest potential will fire. This cell fires one spike only, which signals that that feature has been detected, and the firing cell then causes the values of A_α and B_α to be updated, by its backprojecting inhibitory synapses. Then, a new cycle begins, where the new potential on the sensory cells causes a second feature to be detected in the map layer.

mation of section 9.3.2. The mean field approximation was employed to avoid performing a search over all feature positions at each timestep. However, in a neural implementation, this search would be performed rapidly in parallel, so the approximation would not be required.

Part IV

Chapter 10

Conclusions and future work

The work of this thesis has been aimed at finding out about possible mechanisms of spatial function in animals, and in artificial systems. The main contributions have been:

- A framework for hippocampal involvement in spatial function in the brain, and the implementation of this theory on a mobile robot.
- A grid-based neural network sonar mapping system, inspired by models of visual cortex, which uses lateral connections to overcome the assumption of independence of different grid segment occupancies.
- A probabilistic model of time-of-flight sonar reflection from sharp edges, rough walls, and specular walls.
- A Bayesian feature-based sonar mapping system based on a search for the map with highest *a posteriori* probability, given a set of sonar readings.

This chapter will summarise the results of these investigations, and propose some possible avenues for future research

10.1 Framework for spatial function in the mammalian brain

In chapter 5 we described a framework for spatial function in the mammalian brain, in which the neocortex was responsible for building and representing

responsible for absolute localisation.

The key idea was that absolute localisation does not necessarily require the production of a pair allocentric coordinates to describe the animal's position, but instead could simply involve recalling a complete egocentric map from an observed partial map fragment. In this way, absolute localisation could be performed by an autoassociative memory, which is the architecture proposed by Marr (1971) for the hippocampus.

We considered the results of several previous experiments in relation to the new framework. The Morris (1981) water maze is often cited as the archetypal spatial task dependent on the hippocampus. Our framework would explain the ability of rats at this task as follows: when a rat is introduced to an environment it observes a partial map consisting of the cues visible from its start location, and this is completed to a full map which also contains the location of the hidden platform. We also examined various other experimental results under the new framework, including: short-cut ability of animals; latent learning; the differences between the spatial firing patterns of hippocampal and entorhinal place cells; the effect of environmental manipulations on place cell firing patterns; and non-spatial factors influencing place cell firing.

Based on the theory, we made two experimental predictions. The first is that instantaneous transfer does not occur in mammals. The second is that the hippocampus is *not* the substrate for path-integration ability. Since this proposal was first made (Recce and Harris 1996), some preliminary evidence has been found that this is indeed the case (Alyan et al. 1997).

Chapter 6 described how the framework was instantiated in a mobile robot control system. To make this control system, we needed software modules to perform the functions that, in chapter 5, were ascribed to the neocortex and hippocampus. For the neocortical module, we adapted a previous system of Lee and Recce (1997). This system was designed to evaluate exploration strategies, and was not intended as a model of any neural structure. However, with some small modifications, it is capable of performing the mapping

sents maps in a symbolic rather than neural code, the hippocampus module could not be instantiated by a standard recurrent autoassociative memory. Instead, we used a custom-designed autoassociative memory that performs pattern completion in the same symbolic representation used by the neocortex module.

The control system was tested in various ways, including: getting the robot to perform an analog of the Morris water maze task; demonstrating increased map quality in the early stages of exploration, which indicates recall of the previous environment; and examination of the firing maps of the simulated place cells that compose the simulated hippocampus.

10.2 Neural network grid-based mapping system

In chapter 7, we pointed out that the grid-based sonar mapping methods introduced by Elfes and Moravec are in fact equivalent to systems of integrate-and-fire neurons. Using this analogy, we proposed an extension to these mapping systems, inspired by models of visual cortex, in which the neurons have lateral connections whose aim is to overcome the assumption that the occupancy probabilities of different grid segments are independent.

The network contains two set of cells: a set of "line neurons" that code for the presence of walls; and a set of "free-space neurons" which code for unoccupied areas of the environment. The line neurons have elongated receptive fields, and so are activated by sonar reflections corresponding to straight lines, in a similar way to the "edge-detector" cells of primary visual cortex. The free-space neurons have receptive fields corresponding to the parts of space immediately beyond the neuron's grid location, from which sonar reflections would not be expected if the neuron's grid location was occupied.

Lateral connections in the network enhance firing patterns corresponding to maps of high *a priori* probability, such as those containing collinear line neuron

lower *a priori* probability, such as random scatter.

A critical feature of this system is that there are a range of line neurons for every grid segment, corresponding to lines of different angles passing through that point. Because the reflection properties of specular walls are highly angle-dependent, such a mechanism is essential for a grid-based system to detect specular walls.

The new system was tested with real sonar data collected in a heterogeneous environment. Like the mapping system of Lim and Cho (1992), but unlike the original mapping systems of Elfes and Moravec, this new system was capable of detecting specular walls. The more accurate prior probabilities enforced by the lateral connections also led to the absence of fictitious obstacles, which are seen in the results of the Lim and Cho system.

However, this mapping system was not derived from any mathematical principles, and contained many free parameters, the values of which needed to be tuned by hand, and were not guaranteed to be optimal.

10.3 Probabilistic models of time-of-flight sonar

Chapter 8 describes the production of a set of probabilistic models for sonar reflection from various types of environmental feature.

We collected a large quantity of experimental data for time-of-flight sonar reflections from rough walls, specular walls, and sharp edges, at various distances and incidence angles. Using standard statistical techniques including linear and logistic regression, we constructed probabilistic models for the sonar sensor output, as a function of the distance from the object and incidence angle of the sonar beam.

The modelled probability distributions are characterised by three parameters: the probability of a direct return, and the mean and standard deviation of range readings in the case of a direct return. In the case that a direct return is not observed, the model makes no prediction for the returned range. In this case, there are three possibilities: no return will be received at all; a reflection

the case of edge reflectors); or a secondary specular reflection will occur (in the case of specular walls). In the latter two cases, the return produced depends not only on the position of the primary target but also on an interaction with the position of the secondary target. A suggestion for how to model this interaction was given in section 8.5.3, but was not implemented.

10.4 Bayesian feature-based mapping system

In chapter 9, we described a novel feature-based mapping system, which was derived from a mathematical principle, Bayes' theorem. The probabilistic models of chapter 8 were used to construct a probability function on the space of all maps, and the problem of map construction was rephrased as a search for the map of highest probability. Due to the very large size of the space to be searched, "off-the-shelf" search techniques such as simulated annealing could not be used, and instead a special search technique was developed, based on a mean field approximation, followed by an iterative process.

The mapping system was tested using tracefiles of sonar data previously collected by Lee and Recce (1997), recorded in multiple exploration sessions in three different environments, using a simple wall-following exploration strategy. Map quality was assessed using a previously described map quality metric, and the new mapping system was quantitatively compared to the Lee and Recce mapping system.

We found that, in all cases, map quality was significantly higher in the early stages of exploration under the new mapping system. We concluded that this was because the new system, being based on Bayesian logic, uses all data to construct a map, but the old system throws away many sonar readings.

We also found that asymptotic map quality was significantly higher with the new system in two of the three environments. We concluded that this was due to better accuracy in localisation of features due to the more accurate sonar model used by the new system. However, in the third environment, asymptotic map quality was significantly worse. We concluded that this was

model was explicitly built.

In any algorithm based on a large optimisation problem, computational complexity will be a serious issue. However, by using a series of approximations and tricks, the CPU time required to map a room was kept to approximately 1 minute.

10.5 Directions for future work

There are several avenues for future work suggested by the results of this research. We will now discuss some of them.

10.5.1 Neurobiology

On the neurobiological side, the main avenues of research suggested are experimental. In chapter 5 we made two strong experimental predictions: that instantaneous transfer does not occur, and that path integration ability is not primarily dependent on the hippocampus. On the first issue, there has been a lot of controversy (Sutherland et al. 1987; Keith and McVety 1988; Chew et al. 1989; Keith 1989; Whishaw 1991; Alyan 1994). Recent research seems to suggest that instantaneous transfer does not occur (Alyan 1994), but a conclusive experiment is needed to settle this issue. On the second issue, less research has been done. Alyan et al. (1997) have given some preliminary evidence in support of our prediction, but no paper has yet been published.

10.5.2 Sensor simulators

One piece of work which would be of considerable practical importance to robotics in general, would be to use the probabilistic models of chapter 8 to construct a sonar simulator. Simulators are clearly useful for development of robot control systems, especially if learning techniques or genetic algorithms are employed. Current sonar simulators, however, are based on very simple models, and the development of more realistic simulators would increase the

real world.

Before such a simulator could be constructed, it would be necessary to extend the current models to account for multiple specular reflections. A way this might be achieved was described in section 8.5.3.

10.5.3 Grid-based mapping

The grid-based system of chapter 7 suffers from the problem that it has many free parameters which are not systematically optimised or derived from mathematical principles. One solution to this problem would be to perform some kind of search, perhaps using genetic algorithms to optimise the network. The accurate simulator described in the previous paragraph would be useful for this. Another possibility, originally suggested by Elfes, would be to use a Markov random field approximation to derive the update rules for a grid-based system. If the system is to detect specular walls, it will not be sufficient to have just one variable stored per grid segment, but, like we did here, and like Lim and Cho do, to have several for each grid segment, representing the probability distribution of wall orientations passing through that segment.

To use a Markov random field approximation, it is necessary to have an estimate of the *a priori* probability function on the space of maps. This function would contain parameters, so it could be argued that following this approach would merely replace one set of undetermined parameters with another. Furthermore, the parameters of the *a priori* model would depend on the type of environment the robot is likely to be found in. However, given floor-plans of some rooms in which the robot will be typically found, it may be possible to estimate the parameters of the *a priori* distribution.

10.5.4 Feature-based mapping

The feature-based mapping system of chapter 9 also suggests some lines for future research.

As discussed previously, this system would benefit greatly from a mechanism by which it could learn the probability distributions of readings expected from various types of objects, and by which it could identify and classify new types of objects. Some of techniques which could implement this were described in section 9.7.2.

Neural Implementation

Also suggested in chapter 9 was a way of implementing the mapping system with a neural network. However, in my opinion this is not a good direction for research, at least in the near future. The reason is that to simulate a system with the required number of neurons would be very slow on a serial computer.

Localisation

The system we described in chapter 9 is a mapping system with no associated incremental localisation system. It would be useful to integrate a localisation function with the mapping system. Methods by which this could be achieved were suggested in section 9.7.2. In particular, it is possible that the search for the best feature positions could be combined with a search for the most likely robot position along the robot's trajectory.

Implications for exploration strategies

Finally it would be interesting to repeat the experiments of Lee and Recce (1997) on the merits of different exploration strategies to see if the new mapping system leads to different results. Our intuition is that it would – but only experiments can show for sure.

Appendix A

Glossary of biological terms

When describing biological research, it is necessary to use biological terms. This appendix provides a glossary of some of the more obscure biological and psychological terms used in the thesis, for readers who may not be familiar with them.

Allocentric. A coordinate system in which the position of an animal, and of any cues, are represented with respect to a fixed, world-centred frame.

Cerebellum. A structure at the back of the vertebrate brain, usually thought to be involved in motor function.

Cerebral cortex. The outer surface of the mammalian brain. The cerebral cortex can be divided into **neocortex**, a six-layered structure which accounts for the majority of cortex in humans; **archicortex**, which includes the hippocampus; and **paleocortex**, which includes the olfactory cortex.

Classical conditioning. A simple type of learning involving the pairing of a previously neutral **conditioned stimulus** (e.g. the sound of a bell) with an **unconditioned stimulus** (e.g. the presentation of food) that elicits a response (e.g. salivation). As a result of repeated pairings, the conditioned stimulus comes to elicit a **conditioned response**.

Contralateral. On or relating to the opposite side of the body.

EEG or electroencephalogram. The large scale electrical activity produced by a part of the brain. EEG activity is thought to reflect the coherent firing of large number of synapses, rather than the activity of any individual synapse or cell.

Egocentric. A coordinate system in which the positions of cues are represented with respect to a frame centred on the animal.

Eukaryotic. A cell which contains a membrane-bound nucleus, rather than having the nucleic acids mixed in with the rest of the cytoplasm. Plants and animals cells are eukaryotic, but bacteria are not.

Fimbria-fornix. A bidirectional nervous pathway that connects the hippocampus to subcortical brain structures.

Hippocampus. A structure located in the medial temporal area of the mammalian brain, thought to be involved in either spatial or memory function. The hippocampus consists of three major cell layers, called the **dentate gyrus**, **CA3**, and **CA1**.

Ideothetic. Ideothetic sensory information provides an animal with an estimate of its own motion. Ideothetic information can be derived from tactile, vestibular, and proprioceptive senses.

Ipsilateral. On or relating to the same side of the body.

Lesion. A discrete area of tissue damage.

LTP or **long-term potentiation**. A long-lasting increase in synaptic efficacy.

Nucleus accumbens. A brain structure which receives projections from the hippocampal area and is thought to be involved in reinforcement learning.

Proprioceptive. Proprioceptive senses provide information about the physical configuration of the body, for example muscle tensions or joint angles.

response to a spatially localised stimulus. For example a cell in the visual system might fire when light is incident on a certain discrete area of retina. The spatial area that excites the cell is called the cell's receptive field.

REM sleep. A phase of sleep characterised by high levels of brain activity and rapid eye movements. In humans, dreams occur during REM sleep.

Slow wave sleep. A phase of sleep characterised by low frequency EEG oscillations.

Subiculum and **postsubiculum**. Brain structures located close to the hippocampus.

Thalamus. A structure in the centre of the mammalian brain, shaped like a small pair of eggs. Almost all sensory information arriving at the cerebral cortex passes through the thalamus, and the thalamus also receives extensive projections from the cerebral cortex and other brain areas.

Appendix B

Notation for chapter 9

In general, capital letters denote random variables, small letters denote numerical values, and script characters denote sets.

α	Index for sonar scan number
ξ_α	Transducer position from which scan α was taken
R_α	The range measured on scan α (as a random variable)
r_α	An actual value for R_α
$\mathcal{R}$	The set $\{R_\alpha\}$
i	Index for feature number
N_f	Number of features
F_i	Description of feature i (as a random variable)
f_i	An actual value for F_i
$\mathcal{F}$	The set $\{F_i\}$
$P_f(f)$	Prior probability distribution for a single feature
λ	*A priori* expected number of features
$X_{i\alpha}(F_i, \xi_\alpha)$	The reflection produced by feature i to scan α
$SC_{map}(\{f_i\}; \{r_\alpha\})$	Map score of $\{f_i\}$ given a set of measurements $\{r_\alpha\}$
$SC_{reading}(r_\alpha; \{f_i\})$	Reading score of r_α given a map $\{f_i\}$
$SC_{mf}(f_0; \{r_\alpha\})$	Mean field score of a feature f_0 given a set of measurements $\{r_\alpha\}$
A_α, B_α	Intermediate values used in calculation of mean field score
r_h, θ_h	Hough coordinates used to describe an infinite line
l_1, l_2	Endpoint positions of a finite line
f(l), g(l)	Functions used in computing optimal endpoint positions

Appendix C

Mean-field approximation

The *mean field* score of a feature f_0 is the log likelihood of the observed readings given the presence of the *single* feature f_0 in the map $\mathcal{F}$:

$$SC_{mf}(f_0; \{r_\alpha\}) = logP(\mathcal{R} = \{r_\alpha\} | f_0 \in \mathcal{F})$$

How do we calculate the mean field score? Let us first look at the likelihood of a single observation r_α, given that f_0 is contained in the map. We can calculate it by summing over all possible maps that contain f_0:

$$P(R_\alpha = r_\alpha | f_0 \in \mathcal{F}) = \sum_{N=0}^{\infty} \sum_{f_1 \ldots f_N} P(R_\alpha = r_\alpha | \mathcal{F} = \{f_0, f_1, \ldots, f_N\}) P(\mathcal{F} = \{f_0, f_1, \ldots, f_N\} | f_0 \in \mathcal{F})$$

By equation 9.2,

$$P(R_\alpha = r_\alpha | f_0 \in \mathcal{F}) = \sum_{N=0}^{\infty} \sum_{f_1 \ldots f_N} P(R_\alpha = r_\alpha | \mathcal{F} = \{f_0, f_1, \ldots, f_N\}) \frac{e^{-\lambda} \lambda^{N+1}}{(N+1)!} \prod_{i=1}^{N} P_f(f_i)$$

Substituting in equation 9.5,

$$\sum_{N=0}^{\infty} \sum_{f_1 \ldots f_N} \left[\prod_{i=0}^{N} P(X_{i\alpha} \geq r_\alpha) - \prod_{i=0}^{N} P(X_{i\alpha} > r_\alpha) \right] \frac{e^{-\lambda} \lambda^{N+1}}{(N+1)!} \prod_{i=1}^{N} P_f(f_i)$$

$$= \sum_{N=0}^{\infty} \frac{e^{-\lambda} \lambda^{N+1}}{(N+1)!} \left[P(X_{0\alpha} \geq r_\alpha) \prod_{i=1}^{N} \sum_{f_i} P_f(f_i) P(X_{i\alpha} \geq r_\alpha) \right.$$

$$\left. - P(X_{0\alpha} > r_\alpha) \prod_{i=1}^{N} \sum_{f_i} P_f(f_i) P(X_{i\alpha} > r_\alpha) \right]$$

We can do the sum over N using this identity:

$$\sum_{N=0}^{\infty} \frac{e^{-\lambda} \lambda^{N+1}}{(N+1)!} x^N = \frac{e^{-\lambda}(e^{\lambda x} - 1)}{x}$$

Substituting this in,

$$P(R_\alpha = r_\alpha | f_0 \in \mathcal{F}) = A_\alpha P(X_{0\alpha} \geq r_\alpha) - B_\alpha P(X_{0\alpha} > r_\alpha)$$

where

$$\begin{aligned} A_\alpha &= e^{-\lambda}(e^{\lambda x_\alpha} - 1)/x_\alpha \\ B_\alpha &= e^{-\lambda}(e^{\lambda y_\alpha} - 1)/y_\alpha \end{aligned}$$

are given in terms of the expected probabilities of a feature producing a reflection of range greater than or equal to r_α, or greater than r_α:

$$\begin{aligned} x_\alpha &= \sum_f P_f(f) P(X_{f\alpha} \geq r_\alpha) \\ y_\alpha &= \sum_f P_f(f) P(X_{f\alpha} > r_\alpha) \end{aligned}$$

To use the mean field approximation, all we need to do is get some values for x_α and y_α. To do this from first principles would be very complicated. Instead we will make the approximation that, averaged over feature positions, a return is produced with fixed probability p_{expret} (equal to 0.4 in the implementation described below), and that if a return is produced, all range readings up to the maximum possible are equally likely. Notice that we did not use an explicit model for $P_f(f_0)$, but gave an explicit model for the expectation values x_α and y_α.

The mean-field score for a feature f_0 is therefore given by:

$$SC_{mf}(f_0; \{r_\alpha\}) = \sum_\alpha log\left[A_\alpha P(X_{0\alpha} \geq r_\alpha) - B_\alpha P(X_{0\alpha} > r_\alpha)\right] \qquad \text{(C.1)}$$

Bibliography

Agresti, A. (1990). *Categorical Data Analysis.* Wiley Series in Probability and Mathematical Statistics. Wiley.

Alberts, B., D. Bray, J. Lewis, M. Raff, K. Roberts, and J. D. Watson (1994). *Molecular biology of the cell, third edition.* Garland Publishing, Inc.

Alyan, S., B. Paul, E. Ellsworth, R. White, and B. McNaughton (1997). Is the hippocampus needed for path integration? In *Society for neuroscience abstracts.*

Alyan, S. H. (1994). Evidence against instantaneous transfer of spatial knowledge in the house mouse *(mus musculus). Psychobiology 22*, 328–337.

Ambros-Ingerson, J., R. Granger, and G. Lynch (1990). Simulation of paleocortex performs hierarchical clustering. *Science 247*, 1344–1348.

Andersen, R. A. (1995). Encoding of intention and spatial location in the posterior parietal cortex. *Cerebral Cortex 5*, 457–469.

Antelman, G. (1997). *Elementary Bayesian Statistics.* Edward Elgar.

Baddeley, A. D. (1986). *Working memory.* Oxford University Press.

Barnes, C. A., B. L. McNaughton, S. Y. J. Mizumori, B. W. Leonard, and L.-H. Lin (1990). Comparison of spatial and temporal characterisics of neuronal activity in sequential stages of hippocampal processing. *Progresss in Brain Reseach 83*, 287–300.

Barto, A. and R. Sutton (1981). Landmark learing: an illustration of associative search. *Biological Cybernetics 42*, 1–8.

Bisiach, E. (1993). Mental representation in unilateral neglect and related disorders: The twentieth bartlett memorial lecture. *Quaterly Journal of Experimental Psychology 46A*, 435–461.

space. *Cortex 14*, 129–133.

Bozma, O. and R. Kuc (1994). A physical model-based analysis of heterogenous environments using sonar - ENDURA method. *IEEE Transactions on Pattern Analysis and Machine Intelligence 16*, 497–506.

Braitenberg, V. (1987). *Vehicles: Experiments in synthetic psychology.* MIT Press.

Brooks, R. A. (1986). A robust layered control system for a mobile robot. *IEEE J. of Robotics and Autommation 2*, 14–23.

Brown, M. A. and P. E. Sharp (1995). Simulation of spatial-learning in the Morris water maze by a neural network model of the hippocampal-formation and nucleus accumbens. *Hippocampus 5*, 171–188.

Burgard, W., A. Cremers, D. Fox, D. Hahnel, G. Lakemeyer, D. Schulz, W. Steiner, and S. Thrun (1998). The interactive museum tour-guide robot. In *AAAI-98, To appear.*

Burgess, N., M. Recce, and J. O'Keefe (1994). A model of hippocampal function. *Neural Networks 7*, 1065–1081.

Buzsaki, G. (1989). Two-stage model of memory trace formation: A role for "noisy" brain states. *Neuroscience 31*, 551–570.

Cartwright, B. A. and T. S. Collett (1982). How honey bees use landmarks to guide their return to a food source. *Nature 295*, 560–564.

Chen, L. L., B. L. McNaughton, C. A. Barnes, and E. R. Oritz (1990). Head direction and behavioral correlates of posterior cingulate and medial prestriate cortex neurons in freely moving rats. *Soc. Neurosci. Abstr. 16*, 441.

Chew, G. L., R. J. Sutherland, and I. Q. Whishaw (1989). Latent learning does not produce instantaneous transfer of place navigation: a rejoinder to Keith and McVety. *Psychobiology 17*, 207–209.

Cho, D. W. (1990). Certainty grid representation for robot navigation by a Bayesian method. *Robotica 8*, 159–165.

Cox, I. and J. Leonard (1994). Modeling a dynamic environment using a Bayesian multiple hypothesis approach. *Artificial Intelligence 66*, 311–344.

and P. Postmes (1994). A behavioural analysis of rats with damage to the medial prefrontal cortex using the morris water maze: evidence for behavioural flexibility, but not for impaired navigation. *Brain Research 652*, 323–333.

Droulez, J. and A. Berthoz (1991). A neural network model of sensoritopic maps with predictive and short-term memory properties. *Proc. Nat. Acad. Sci. 88*, 9653–9657.

Drumheller, M. (1987). Mobile robot localization using sonar. *IEEE Transactions on Pattern Analysis and Machine Intelligence 9*(2), 325–332.

Duhamel, J. R., C. L. Colby, and M. E. Goldberg (1992). The updating of the representation of visual space in parietal cortex by intended eye movements. *Science 255*, 90–92.

Eichenbaum, H., M. Kuperstein, A. Fagan, and J. Nagode (1986). Cue-sampling and goal-approach correlates of hippocaqmpal unit activity in rats performing an odor discrimination task. *Journal of Neuroscience 7*, 716–732.

Elfes, A. (1987). Sonar-based real-world mapping and navigation. *IEEE J. Robotics and Automation 3*, 249–265.

Flash, T. and N. Hogan (1985). The coordination of arm movements: an experimentally confirmed mathematical model. *Journal of Neuroscience 5*, 1688–1703.

Gardner-Medwin, A. R. (1976). The recall of events through the learning of associations between their parts. *Proc. R. Soc. Lond. B 194*, 375–402.

Gluck, M. and C. Myers (1993). Hippocampal mediation of stimulus representation: A computational theory. *Hippocampus 3*, 491–516.

Goldman-Rakic, P. S. (1990). Cortical localization of working memory. In J. McGauch, N. Weinberger, and G. Lynch (Eds.), *Brain organization and memory: cells, systems, and circuits.* Oxford University Press.

Gothard, K., W. Skaggs, and B. McNaughton (1996). Dynamics of mismatch correction in the hippocampal ensemble code for space: interaction between path integration and environmental cues. *Journal of Neuroscience 16*, 8027–8040.

Gray, C. M., P. König, A. Engel, and W. Singer (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar sychronization which reflects global stimulus properties. *Nature 338*, 334–337.

Hallam, J. (1986). Analysing specular echoes in active acoustic range data. In A. Cohn and J. Thomas (Eds.), *Artificial intelligence and its applications*, pp. 165–177. John Wiley and Sons.

Hawkins, R., E. Kandel, and S. Siegelbaum (1993). Learning to modulate transmitter release - themes and variations in synaptic plasticity. *Annual Review of Neuroscience 16*, 625–665.

Hill, A. and P. Best (1981). Effects of deafness and blindness on the spatial correlates of hippocampal unit activity in the rat. *Experimental neurology 74*, 204–217.

Jarrard, L. E. (1993). On the role of the hippocampus in learning and memory in the rat. *Behav. Neural Biology 60*, 9–26.

Judd, S. and T. S. Collett (1998). Multiple stored views and landmark guidance in ants. *Nature 392*, 710–714.

Jung, M. W. and B. L. McNaughton (1993). Spatial selectivity of unit activity in the hippocampal granular layer. *Hippocampus 3*, 165–182.

Keith, J. R. (1989). Does latent learning produce instantaneous transfer of place navigation. *Psychobiology 17*, 210–211.

Keith, J. R. and K. M. McVety (1988). Latent place learning in a novel environment and the influences of prior training in rats. *Psychobiology 156*, 146–151.

Knierim, J., H. Kudrimoti, and B. McNaughton (1995). Place cells, head direction cells and the learning of landmark stability. *J. Neurosci. 15*, 1648–1659.

Kolb, B., K. Buhrmann, R. McDonald, and R. J. Sutherland (1994). Dissociation of the medial prefrontal, posterior parietal, and posterior temporal cortex fo r spatial navigation and recognition memory in the rat. *Cerebral Cortex 6*, 664–680.

Kolb, B. and I. Whishaw (1995). *Fundamentals of human neuropsychology.* W.H. Freeman.

Current Opinion in Neurobiology 7, 826–831.

Kuc, R. and M. W. Siegel (1987). Physically based simulation model for acoustic sensor robot navigation. *IEEE Trans. Pattern Analysis and Machine Intelligence 9*, 766–778.

Lee, D. and M. Recce (1997). Quantitative evaluation of the exploration strategies of a mobile robot. *Int. J. Robotics Research 16*, 413–447.

Lee, D. C. (1996). *The Map-Building and Exploration Strategies of a Simple, Sonar-Equipped Mobile Robot; An Experimental, Quantitative Evaluation.* Distinguished Dissertations in Computer Science. Cambridge University Press.

Leonard, J. J. and H. F. Durrant-Whyte (1992). *Directed sonar sensing for mobile robot navigation.* Cambrage, MA: Kluwer Academic Publishers.

Levy, W. (1996). A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. *Hippocampus 6*, 579–590.

Lim, J. H. and D. W. Cho (1992). Physically based sensor modeling for a sonar map in a specular environment. In *Proceedings of the 1992 IEEE international conference on robotics and automation*, pp. 1714–1719.

Madhavan, R., M. Dissanayake, and H. Durrant-Whyte (1998). Autonomous underground navigation of an lhd using a combined icp-ekf approach. In *Proceedings of the 1998 IEEE international conference on robotics and automation*, Volume 4, pp. 3703–3708.

Maguire, E., M. Burgess, J. Donnett, R. Frackowiak, C. Frith, and J. O'Keefe (1998). Knowing where and getting there: a human navigation network. *Science 280*, 921–924.

Markus, E. J., Y. Qin, B. J. Leonard, W. E. Skaggs, B. L. McNaughton, and C. A. Barnes (1995). Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *J. Neurosci. 15*, 7079–7094.

Marr, D. (1969). A theory of cerebellar cortex. *J. Physiol. Lond. 202*, 437–370.

161–234.

Marr, D. (1971). Simple memory: a theory for archicortex. *Phil. Trans. R. Soc. Lond. B 176*, 23–81.

Maskarov, D. and H. Durrant-Whyte (1995). Mobile vehicle navigation in unknown environments: a multiple hypothesis approach. *IEE Proc. Control Theory Appl. 142*, 385–400.

McKerrow, P. J. (1991). *Introduction to Robotics.* Addison-Wesley.

McNaughton, B. L., C. A. Barnes, J. L. Gerrard, K. Gothard, M. W. Jung, J. J. Knierim, H. Kudrimoti, Y. Qin, W. E. Skaggs, M. Suster, and K. L. Weaver (1996). Deciphering the hippocampal polyglot: the hippocampus as a path integration system. *J. Exper. Biology 199*, 173–185.

McNaughton, B. L., C. A. Barnes, and J. O'Keefe (1983). The contributions of position, direction and velocity to single unit activity in the hippocampus of freely-moving rats. *Exp. Brain Res. 52*, 41–49.

McNaughton, B. L., L. L. Chen, and E. J. Markus (1991). Dead reckoning, landmark learning, and the sense of direction - a neurophysiological and computational hypothesis. *Journal of Cognitive Neuroscience 3*, 192–202.

Mittelstaedt, M. L. and H. Mittelstaedt (1980). Homing by path integration in a mammal. *Naturwissenschaften 67*, 566–567.

Miyashita, Y. (1988). Neuronal correlate of visual associative long-term memory in the primate temporal cortex. *Nature 335*, 817–820.

Mizumori, S. and J. Williams (1993). Directionally selective mnemonic properties of neurons in the lateral dorsal nucleus of the thalamus of rats. *Journal of Neuroscience 13*, 4015–4028.

Moravec, H. P. (1988). Sensor fusion in certainty grids for mobile robots. *AI Magazine 9*, 61–74.

Morris, R. G. M. (1981). Spatial localization does not require the presence of local cues. *Learn. Motiv. 12*, 239–260.

Morris, R. G. M., P. Garrud, J. N. P. Rawlins, and J. O'Keefe (1982). Place navigation impared in rats with hippocampal lesions. *Nature (Lond) 297*, 681–683.

direction firing properties of hippocampal place cells. *Journal of Neuroscience 14*, 7235–7251.

Muller, R. U. and J. L. Kubie (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *J. Neurosci. 7*, 1951–1968.

Muller, R. U., J. L. Kubie, E. M. Bostock, J. S. Taube, and Q. G. (1987). Spatial firing correlates of neurons in the hippocampal formation of freely moving rats. *J. Neurosci. 7*, 1951–1968.

Muller, R. U., J. L. Kubie, and R. J. B. Jr. (1991). Spatial firing patterns of hippocampal complex spike cells in a fixed environment. In J. Paillard (Ed.), *Brain and Space*, pp. 296–333. Oxford UP.

Myers, C. and M. Gluck (1994). Context, conditioning, and hippocampal re-representation. *Behavioural Neuroscience 108*, 835–847.

Myers, C., M. Gluck, and R. Granger (1995). Dissociation of hippocampal and entorhinal function in associative learning: a computational approach. *Psychobiology 23*, 116–138.

Myers, C. E., B. R. Ermita, K. D. Harris, M. Hasselmo, P. Solomon, and M. A. Gluck (1996). A computational model of cholinergic disruption of septo-hippocampal activity in classical eyeblink conditioning. *Neurobiology of Learning and Memory 66*, 51–66.

Nilsson, N. J. (1984). Shakey the robot. Technical Note 323, SRI International, Menlo Park, California.

O'Keefe, J. (1976). Place units in the hippocampus of the freely moving rat. *Exp. Neurol. 51*, 78–109.

O'Keefe, J. and D. H. Conway (1978). Hippocampal place units in the freely moving rat: why they fire and where they fire. *Exp. Brain Res. 31*, 573–590.

O'Keefe, J. and J. Dostrovsky (1971). The hippocampus as a spatial map. preliminary evidence from unit activity in the freely-moving rat. *Brain Research 34*, 171–175.

O'Keefe, J. and L. Nadel (1978). *The Hippocampus as a Cognitive Map*. Oxford, UK: Clarendon Press.

O'Keefe, J. and M. L. Recce (1993). Phase relationship between hippocampal place units and the eeg theta rhythm. *Hippocampus 3*, 317–330.

O'Keefe, J. and A. Speakman (1987). Single unit activity in the rat hippocampus during a spatial memory task. *Exp. Brain Res. 68*, 1–27.

Owen, A. (1997). Cognitive planning in humans: neuropsychological, neuroanatomical, and neuropharmacological perspectives. *Progress in Neurobiology 53*, 431.

Pavlov, I. P. (1927). *Conditioned reflexes; an investigation into the physiological activity of the cerebral cortex.* Oxford University Press.

Polaroid (1991). *Ultrasonic Ranging System - User Guide.* 119 Windsor Street - 2B, Cambridge, MA 02139, USA: Polaroid Corporation, Ultrasonic Componenets Group.

Press, W., B. Flannery, S. Teukolsky, and W. Vetterling (1993). *Numerical recipes in C.* Cambridge University Press.

Quirk, G. J., R. U. Muller, and J. L. Kubie (1990). The firing of hippocampal place cells in the dark depends on the rat's recent experience. *J. Neurosci. 10*, 2008–2017.

Quirk, G. J., R. U. Muller, J. L. Kubie, and J. B. J. Ranck (1992). The positional firing properties of medial entorhinal neurons: description and comparison with hippocampal place cells. *J. Neurosci. 12*, 1945–1963.

Rabiner, L. and B. Juang (1986). An introduction to hidden markov models. *IEEE ASSP Magazine.*

Ranck, J. B., J. (1985). Head direction cells in the deep cell layer of dorsal presubiculum in freely moving rats. In G. Buzsaki and C. H. Vanderwolf (Eds.), *Electrical activity of the archicortex*, pp. 217–220. Budapest: Publishing House of the Hungarian Academy of Sciences.

Recce, M. (1994). *The representation of space in the rat hippocampus as revealed using new computer-based methods.* Ph. D. thesis, University College London.

Recce, M. and K. D. Harris (1996). Memory for places: a navigational model in support of Marr's theory of hippocampal function. *Hippocampus 6*, 735–748.

tional neuroscience theory of rodent navigation. Ph. D. thesis, Carnegie Mellon University.

Redish, A. and D. Touretzky (1997). Cognitive maps beyond the hippocampus. *Hippocampus 7*, 15–35.

Reeves, C. R. (1993). *Modern heuristic techniques for combinatorial problems.* Blackwell Scientific.

Richardson, S. and P. Green (1997). On bayesian analysis of mixtures with an unknown number of components. *Journal of the Royal Statistical Society B 59*, 731–792.

Save, E., A. Cressant, C. Thinus-Blanc, and B. Poucet (1996). Early visual deprivation does not prevent hippocampal place cell firing in the rat. *Society for neuroscience abstracts 22*, 912.

Save, E. and M. Moghaddam (1996). Effects of lesions of the associative parietal cortex on the acquisition and use of spatial memory in egocentric and allocentric navigation tasks in the rat. *Behavioural Neuroscience 110*, 74–85.

Schmaltz, L. and J. Theios (1972). Acquisition and extinction of a classically conditioned response in hippocampectomized rabbits. *J Comp Physiol Psychol 79*, 328–333.

Schnitzer, M. (1993). Theory of continuum random walks and application to chemotaxis. *Physical Review E 48*, 2553–2568.

Scoville, W. B. and B. Milner (1957). Loss of recent memory after bilateral hippocampal lesions. *J. Neurol. Neurosurg. Psychiatry 20*, 11–21.

Sharp, P. E. and C. Green (1994). Spatial correlates of firing patterns of single cells in the subiculum of the freely moving rat. *Journal of Neuroscience 14*, 2339–2356.

Singer, C. and W. Gray (1995). Visual feature integration and the temporal correlation hypothesis. *Annual review of neuroscience 18*, 555–586.

Strong, S., B. Freedman, W. Bialek, and R. Koberle (1998). Adaptation and optimal chemotactic strategy for *E. coli. Physical Review E 57*, 4604–4617.

Some limitations on the use of distal cues in place navigation by rats. *Psychobiology 15*, 48–57.

Sutherland, R. J. and R. Linggard (1982). Being there: A novel demonstration of latent spatial learing. *Behavioural and Neual Biology 36*, 103–107.

Sutherland, R. J. and J. W. Rudy (1989). Configural association theory: The role of the hioppocampal formation in learning, memory, and amnesia. *Psychobiology 17*, 129–144.

Swanson, L. W. (1983). The hippocampus and the concept of the limbic system. In W. Seifert (Ed.), *Neurobiology of the hippocampus*, pp. 3–19. Academic Press.

Taube, J. and H. Burton (1995). Head direction cell-activity monitored in a novel environment and during a cue conflict situation. *Journal of Neurophysiology 74*, 1953–1971.

Taube, J. S. (1995). Head direction cells recorded in the anterior thalamic nuclei of freely moving rats. *Journal of Neuroscience 15*, 70–86.

Taube, J. S., R. U. Muller, and J. Ranck, J. B. (1990a). Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *Journal of Neuroscience 10*, 420–435.

Taube, J. S., R. U. Muller, and J. Ranck, J. B. (1990b). Head-direction cells recorded from the postsubiculum in freely moving rats. II. Effects of environmental manipulations. *Journal of Neuroscience 10*, 436–447.

Thompson, L. T. and P. Best (1990). Long-term stability of the place-field activity of single units recorded from the dorsal hippocampus of freely behaving rats. *Brain research 509*, 299–308.

Thrun, S., D. Fox, and W. Burgard (1998). A probabilistic approach to concurrent mapping and localisation for mobile robots. *Machine Learning 31*, 29–53.

Tolman, E., B. Ritchie, and D. Kalish (1946). Studies in spatial learning. I. Orientation and the short-cut. *J. Exp. Psychol. 36*, 13–24.

Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychol. Rev. 55*, 189–208.

on interacting representations of space. *Hippocampus 6*, 247–270.

Touretzky, D. S. and A. D. Redish (1995). Landmark arrays and the hippocampal cognitive map. In L. Niklasson and M. Boden (Eds.), *Current trends in connectionism - Proceedings of the 1995 Swedish Conference on Connectionism*, pp. 1–13. Lawrence Erlbaum.

Treves, A. and E. T. Rolls (1994). A computational analysis of the role of the hippocampus in memory. *Hippocampus 4*, 373–391.

Ungerleider, L. and M. Mishkin (1982). Two cortical visual systems. In *Analysis of visual behaviour*, pp. 549–586. MIT press.

Vanderwolf, C. H. (1969). Hippocampal electrical activity and voluntary movement in the rat. *Electroencephalogr. Clin. Neurophysiol. 26*, 407–418.

Wallenstein, G. and M. Hasselmo (1997). GABAergic modulation of hippocampal population activity: sequence learning, place field development, and the phase precession effect. *J. Neurophysiol. 78*, 393–408.

Walter, W. G. (1950). An imitation of life. *Scientific American 182*, 42–45.

Watson, J. (1907). Kinaesthetic and organic sensations: their role in the reactions of the white rat to the maze. *Psychological Review 8*, 43–100.

Wehner, R., B. Michael, and P. Antonsen (1996). Visual navigation in insects - coupling of egocentric and geocentric information. *J. Exper. Biology 199*, 129–140.

Whishaw, I. (1991). Latent learning in a swimming pool place task by rats: evidence for the use of associative and not cognitive mapping processes. *Quaterly Journal of Experimentimental Psychology 43B*, 83–103.

Whishaw, I. Q., J.-C. Cassel, and L. E. Jarrard (1995). Rats with fimbria-fornix lesions display a place response in a swimming pool: a dissociation between getting there and knowing where. *J. Neurosci. 15*, 5779–5788.

Wiener, S., C. Paul, and H. Eichenbaum (1989). Spatial and behavioural correlates of hippocampal neuronal activity. *Journal of Neuroscience 9*, 2737–2763.

Wiener, S. I. (1993). Spatial and behavioral correlates of striatal neurons

science 13, 3802–3817.

Willshaw, D. J. and J. T. Buckingham (1990). An assessment of Marr's theory of the hippocampus as a temporary memory store. *Phil. Trans. R. Soc. Lond. B 329*, 205–215.

Wilson, M. A. and B. L. McNaughton (1993). Dynamics of the hippocampal ensemble code for space. *Science 265*, 1055–1058.

Wiskott, L. and C. von der Malsburg (1993). A neural system for the recognition of partially occluded objects in cluttered scenes - a pilot study. *International Journal of Pattern Recognition and Artificial Intelligence 7*, 935–948.

Young, B. J., G. D. Fox, and H. Eichenbaum (1994). Correlates of hippocampal complex-spike cell-activity in rats performing a non-spatial radial maze task. *J. Neurosci. 14*, 6553–6563.

Zhang, K. (1996). Representation of spatial orientation by the intrinsic dynamics of the head-directon cell ensemble: A theory. *J. Neurosci. 16*, 2112–2126.

Zipser, D. and R. A. Andersen (1988). A back-propagation programmed neural network that simulates response properties of a subset of posterior parietal neurons. *Nature 331*, 679–684.