Memory for places: A navigational model in support of Marr's theory of hippocampal function

Michael Recce and Kenneth D. Harris
Department of Anatomy and Developmental Biology
University College London
London WC1E 6BT
United Kingdom

Short title: Memory for places

32 text pages, 8 figures, no tables

Correspondence author: Michael Recce, Department of Anatomy and Developmental Biology, University College London, WC1E 6BT, UK. (tel) +44 171 380 7839; (fax) +44 171 380 7839; e-mail m.recce@ucl.ac.uk

Key words: associative memory, spatial

Abstract

In this paper we describe a model that applies Marr's theory of hippocampal function to the problem of map based navigation. Like many others we attribute a spatial memory function to the hippocampus, but we suggest that the additional functional components required for map based navigation are located elsewhere in the brain. One of the key functional components in this model is an egocentric map of space, located in the neocortex, that is continuously updated using ideothetic (self motion) information. The hippocampus stores snapshots of this egocentric map. The modelled activity pattern of head direction cells is used to set the best egocentric map rotation to match the snapshots stored in the hippocampus, resulting in place cells with a non-directional firing pattern. We describe an evaluation of this model using a mobile robot, and demonstrate that with this model the robot can recognise an environment and find a hidden goal. This model is discussed in the context of prior experiments that were designed to discover the map based spatial processing of animals. We also predict the results of further experiments.

Introduction

Evidence has accumulated in support of several somewhat different views of hippocampal function. While there is almost general agreement that spatial memories are formed in the hippocampus, there is less agreement on hippocampal involvement in non-spatial memory, or in specialised spatial function other than memory (e.g. the ability to take short-cuts). This debate has been complicated by the difficulties in classifying forms of spatial or navigational abilities, and in precisely defining a verifiable type of non-spatial memory that requires the hippocampus.

The most influential theory of hippocampal involvement in spatial function is the *cognitive map* theory of O'Keefe and Nadel (1978). O'Keefe and Nadel (1978) distinguished between a method for finding spatial locations based on maps (the *locale* system) and one based on routes (the *taxon* system). Further they proposed that the locale system or cognitive map is anatomically located in the hippocampus. The proposed cognitive map is allocentric, in that the location of the animal and the external cues are represented in an fixed, world-centered frame. During the past twenty five years the spatial role of hippocampal function has received wide-ranging support, including data from lesion experiments in rats (see Jarrard(1993) for a review), and data on the firing properties of hippocampal *place cells* (O'Keefe & Dostrovsky, 1971).

An alternative theory holds that the hippocampus is a temporary memory store. 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. The patterns are stored after a single presentation, and the recall process reconstructs the initially stored patterns. Marr derived a neuronal architecture suitable for performing this function (now called *auto-associative 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. This theory has received renewed attention recently (McNaughton & Morris, 1987; Treves & Rolls, 1994; Willshaw & Buckingham, 1990). Marr's view is consistent with the amnesic effects of medial temporal lesion in humans (Scoville and Milner (1957); see Squire (1992) for a recent review). 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 (now called *hetero-associative memory*).

What is the goal of the model?

The goal of the model described in this paper is to show that the auto-associative memory theory of Marr is also consistent with the involvement of the hippocamus in spatial behaviour. In describing this model, we ascribe some of the locale spatial functions to brain structures other than the hippocampus. The model is also intended as a step toward the development of a robust navigational system. Most other navigation models either rely on error free path integration, or require highly accurate identification of landmarks in an environment. We propose a model that uses real, and therefore imprecise, sensory inputs, and we evaluate it using a real mobile robot.

We present the model at two levels, beginning with the conceptual level. At this level, the model ascribes functions to various brain structures, and describes how they might interact to enable the animal to perform spatial behaviour. We then describe the model at an implementational level, giving a description of the workings of the robot and the computer program used to evaluate the conceptual model. For ease of programming, the required functions have not all been implemented in the same way believed to occur in the brain. However, the implementation can still provide initial verification of the functional decomposition of the hypothesised navigational system. In the Discussion section we present prior work that postulates mechanisms by which the required functions could be carried out by systems of neurons.

What are the data to be addressed?

The spatial abilities of animals. Animals are clearly capable of heading directly towards a visible goal. In addition to this rather simple ability, experiments have shown that in a familiar environment, they are able to head towards goals that are not directly detectable by vision (or other immediate sensory input). This ability exists for goals located at a fixed direction and distance from the starting point of a search (egocentric location) (Save & Moghaddam, 1996), or at fixed positions with respect to the distal cues in an environment (allocentric location) (Morris, 1981). Furthermore, there is evidence to suggest that animals are capable of taking short-cuts through unexplored areas of familiar environments (Tolman, 1932; Chapius & Scardigli, 1993).

Spatial abilities have also been demonstrated in environments that have not been previously explored. For example, animals can use *path integration* to return to a starting position. Path integration is the ability to deduce a homing vector on the basis of ideothetic (self motion) information, and has been demonstrated to occur in the absence of specific sensory markers of the goal location. Evidence for path integration has been found both in mammals (Mittelstaedt & Mittelstaedt, 1980) and invertebrates (e.g. Wehner *et.al*, (1996)).

In some cases, animals appear to deduce homing vectors on the basis of the constellation of visual cues (Morris, 1981), even if they are disoriented before being placed in the location from which they must deduce the homing vector (Alyan, 1994). We refer to the ability to deduce location information solely from distal cues as *cue-based localisation.* Initial experiments suggested that cue-based localisation was possible in parts of the environment that the animal had not previously visited (Morris, 1981; Sutherland & Linggard, 1982), and this ability has been called *instantaneous transfer* (Morris, 1981). Recently however, the suggestion that animals can perform instantaneous transfer has become more controversial (Sutherland *et al.*, 1987; Keith & McVety, 1988; Chew *et al.*, 1989; Keith, 1989; Whishaw, 1991; Alyan, 1994).

The effect of brain lesions on spatial ability. Lesions to the cerebral cortex in rats have been shown to disrupt their performance in a wide variety of spatial tasks (e.g. (Kolb *et al.*, 1994; Save & Moghaddam, 1996). Data suggest that the cortical areas involved in spatial function include the posterior parietal cortex and medial prefrontal cortex (but see de Bruin *et. al* (1994)). In humans, lesions to the right parietal cortex have been linked to *spatial hemineglect* (see Bisiach (1993) for a review). In this condition, the patient seems unable 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 & Luzzatti, 1978).

Hippocampal lesions in rats have been shown to disrupt navigation to invisible goals at fixed allocentric locations, while not affecting navigation to visible goals ((Morris *et al.*, 1982); see (Jarrard, 1993) for a review). In humans, recent analysis of the affects of medial temporal lesions and data from brain imaging experiments has implicated the hippocampal formation in visual-spatial memory (Pigott & Milner, 1993; Smith *et al.*, 1995; Feigenbaum *et al.*, 1996).

Properties of place cells. Recordings in rat hippocampus have revealed the existence of *place cells*, which fire when the rat is in a specific part of an environment (O'Keefe & 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, and this stability has been demonstrated to persist for many days (Thompson & Best, 1990). When rats randomly search for food in an open field environment, the firing of these cells is largely independent of head direction (Muller *et al.*, 1994), but when rats are constrained to move along a fixed one-dimensional trajectory, the cells have been shown to fire preferentially in one direction of movement (McNaughton *et al.*, 1983; O'Keefe & Recce, 1993; Markus *et al.*, 1995).

Several studies have demonstrated that the firing of place cells can be driven by the location of distal cues. For example, if an entire environment is rotated, the firing areas of place cells rotate with it (O'Keefe & Conway, 1978; Muller & Kubie, 1987). However, the firing of place cells is not exclusively controlled by distal cues. In one experiment, rats were placed in a familiar environment with the room lights switched off (Quirk *et al.*, 1990). In this case, a new set of place cells begins to fire, and the pattern of activity persists after the lights have been switched back on. New place cells are also formed when a rat is allowed access to a previously unreachable part of a familiar environment (Wilson & McNaughton, 1993).

Finally, neurons in other areas near the hippocampus proper, such as the dentate gyrus (Jung & McNaughton, 1993), the subiculum (Sharp & 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 of medial entorhinal cells are topologically transformed when the shape of the environment changes (Quirk *et al.*, 1992).

Head direction cells Cells in the postsubiculum of rats have been shown to have a firing pattern that is sharply tuned to the animals allocentric head direction (Ranck, 1985; Taube *et al.*, 1990a). The firing of these *head direction* cells does not appear to be modulated by the animals location in an environment. 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 & Williams, 1993), the retrosplenial and parietal cortical areas (Chen *et al.*, 1994; McNaughton *et al.*, 1991) and the striatum (Wiener, 1993).

Head direction cells have been shown to sustain their activity after salient visual cues have been removed (Taube *et al.*, 1990b; McNaughton *et al.*, 1991; Mizumori & Williams, 1993). In addition, head direction cells have been shown to be coupled to the place cells, in that both of the cell types rotate together, when the salient cues are rotated in a familiar environment (McNaughton *et al.*, 1994).

How should the reader evaluate the model?

A good model or theory should explain a number of experimental observations and make predictions for feasible experiments. These positive features need to be balanced by the number of unproven assumptions made by the theory and by the complexity of the theory.

The model described here seeks to unify the two hippocampal theories of map-based navigation and of intermediate memory, and can be judged successful if it reduces the apparent conflict between these theories. We propose two experiments which could verify or refute the ideas and assumptions underlying this model. Furthermore, we present the model as a potential solution to the general robot localisation problem, which has as yet not been solved by the engineering community.

The Conceptual Model

In this section we describe the model at a conceptual level, and its relation to current experimental findings. The computational details are presented in the next section, where we describe how the model was implemented with a robot.

A block diagram of the model we evaluate is shown in figure 1. The model consists of two parts. The hippocampus is modelled by an auto-associative memory, and the modelled neocortex performs four separate spatial functions which together allow it to represent an egocentric map of the environment.

———Figure 1 here———

The neocortex

Processing of sensory data. Marr (1970) proposed that one the functions 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. In general, sensory input is processed through several successive cortical stages prior to reaching the hippocampus (Swanson, 1983), each of which may lead to redundancy compression.

In a spatial context, redundancy compression could result in the classification of cues and determination of their location. In our model, the neocortical activity pattern represents the locations of external cues and a possible goal in egocentric coordinates. We will refer to this set of cue locations as the *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).

In order to explain the continued firing of place cells in the absence of direct sensory input (O'Keefe & Speakman, 1987; Quirk *et al.*, 1990), our model requires the neocortex to store the nature and location of cues in working memory, so they remain in the egocentric map even if the animal is not directly attending to them.

Path integration. We suggest that the neocortex uses motor efference copy and ideothetic input to update the positions of remembered cues in an egocentric coordinate frame. This is in contrast to an approach in which the animal's position is updated in an allocentric coordinate system. We propose that the positions of all the cues are updated simultaneously and in parallel.

This mechanism is similar to one suggested for the monkey visual 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 & Berthoz, 1991).

Coordinate transformations. Recordings from neurons in the posterior parietal cortex of monkeys have identified neurons with retinotopic visual receptive fields, whose activity level is modulated by eye and head position (see (Andersen, 1995) for a review). It has been proposed that these neurons provide a distributed representation for the spatial location of stimuli, from which location information can be extracted in non-retinotopic coordinate systems (Zipser & 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) suggested that the head direction cells may be part of a system for maintaining a fixed allocentric bearing to landmarks. This is accomplished by rotating the direction to a landmark by the angle coded for by the instantaneous activity of the population of head direction cells. We have incorporated this role for the head-direction system in our model. All of the features in the cortical egocentric map are rotated by the by the angle coded for by the head direction system upstream from the hippocampus. In this way the location of features in the egocentric map is represented by an egocentric bearing and an egocentric distance, and the features stored in the modelled hippocampus are represented by an allocentric bearing and an egocentric distance. As a result the input to the hippocampus is independent of head direction. This produces non-directional place fields, that are consistent with the place cell activity patterns recorded in open-field environments (Muller *et al.*, 1994). We propose that the brain uses similar mechanisms to perform this rotation as it uses to perform other coordinate transformations, such as the conversion from retinal to head centered coordinates.

———Figure 2 here———

Planning and execution of movements. In order to evaluate the model, it is necessary to include a rudimentary system to drive the modelled rat's behaviour in an environment. This was achieved by including a drive to reach a goal and a drive to explore.

In our model, the animal is in one of three behavioural modes (see figure 2). When the animal is first placed in an environment it is in *orientation mode*. In this mode, the animal directs is behaviour towards finding the locations of salient features in order to construct a useful egocentric map of the environment. There are two conditions that can lead to a change from this mode: high map quality and the recall of an eogcentric map from memory. The map quality is is a measure of the fraction of the sensory inputs that are consistent with features that are currently part of the egocentric map. The assumption here is that when the animal is initially placed in an environment it has no knowledge of the location of features. As features are detected the memory is searched for a match to the egocentric map. In this search, the rotational angle coded by the head direction cells is systematically varied to maximise the fit between the current egocentric map and hippocampally stored maps.

If the egocentric map is of sufficient quality and no hippocampal map is recalled for any rotational angle, the animal goes into *exploration mode*. In this mode, the animal explores the environment, and the hippocampus stores the cortical maps at each location, as described below. When the animal first enters exploration mode the allocentric angle, coded for by head direction cells, is set to an arbitrary value. As exploration continues, the head direction angle is updated using ideothetic input, so that it measures the animals head direction relative to this initial, arbitrary bearing. As a result, the maps stored in the modelled hippocampus all have the same orientation.

Alternatively, if an egocentric is found in the auto-associative memory and recalled, then the animal goes into *recall mode.* In this mode, the animal is not driven to explore, and so may engage in other behaviour such as returning to the site of a goal. In recall mode, the head direction cell coding is updated by ideothetic input.

The hippocampus

Consistent with the theory of Marr (1971), the modelled hippocampus stores snapshots of cortical input. At a later time, when presented with a fragment of a stored snapshot, the hippocampus reactivates the full cortical activity pattern that was present when the snapshot was stored. New snapshots are stored when no stored pattern is found that completes the current cortical activity pattern. This occurs in all behavioural modes except orientation mode, in which snapshot storage is inhibited.

In a spatial context, the cortical pattern represents a set of cue descriptions which make up an egocentric map. Therefore, the modelled hippocampus stores the egocentric maps that were active at the locations the animal has visited. If the animal is later at one of these locations with a partial map, the full map is reactivated.

In our model, the neocortex distinguishes between cues which have been directly observed, and cues whose existence is suggested by recall of hippocampally stored maps. Cues contained in the maps recalled by the hippocampus are treated as *provisional* until their existence is later confirmed by observation. Provisional cues are not included in the pattern sent to the hippocampus. However, they may be used to direct the animals behaviour. For example, the animal may direct its attention towards the recalled location of a salient but as yet unobserved cue, in order to more quickly determine its location.

The details of the interaction between the egocentric map in the neocortex, the behavioural modes, the place units, and the head direction cells are summarised in figure 3.

The population activity of the head direction cells codes for an angle α which is the difference between the allocentric bearing of the environment and the current head direction. When the animal is initially placed in the environment, in orientation mode, the allocentric bearing is set to an arbitrary value. All of the egocentric maps that are stored in a single environment have the same allocentric bearing.

———Figure 3 here———

Implementation of the model

In this section we describe the robot that was used to test the conceptual model described in the previous section, and details of the computer programs used to implement the model. Each of the functional components is described separately. The implementational details are not intented to match a specific method used by brain regions. Instead, the implementation is matched to the requirements of the sonar sensor, which is unlike any of the sensory systems of a rat. Also, the data structures and techniques are selected for easy implementation on a digital computer, and not to match a particular system of neurons.

———Figure 4 here————

The Robot

We evaluated the model using a real robot, rather than with a computer simulation. As discussed below, our view is that this is a more stringent test of the model. The purpose built robot, ARNE, has a circular base with a 25 cm diameter (see figure 4). There are two movements that the robot can make: moving a specified distance forward, and rotating by a specified angle without translation. ARNE is equipped with shaft encoders on the two drive wheels, that provide a measure of the distance moved and the rotation angle of a turn. A single sonar sensor, mounted on a motorised pivot, is used to measure the positions of objects in the environment. After each movement of the robot, the sonar transducer is rotated through a complete 360° scan in 18° steps, and the distances to objects are calculated from the time-of-flight of the sonar signal (range 2.5 meters, maximum resolution 0.01 meters).

For several reasons single sonar distance measurements are not reliable. The sonar system can only accurately measure the distance to a wall that is perpendicular to the wavefront of the sound signal. An acute angle between the beam and a wall often results in a specular reflection, resulting in an exaggerated distance measurement. Also, spreading of the sonar beam results in imprecise angular localisation of the sonar reflection point. The properties of the sonar system are illustrated in the typical sonar scan, shown in figure 5. There are 19 sonar distance measurements in a single scan (there is no measurement in the direction directly behind the robot), and each of these is represented by a radial dashed line in the figure. The solid lines, in the figure, show a superimposed plan view of the true wall locations in an environment. The erroneous sonar readings that extend beyond the wall are due to specular reflections of the beam. The readings fall short of the wall when the nearest point is not at the center of the spreading beam. Previous work has shown that scans from a sequence of movements can be combined to produce an accurate map of an environment (Leonard & Durrant-Whyte, 1992; Lee & Recce, 1996). In comparison with a rat, ARNE has very limited sensory input.

In order to simulate a hidden goal, an overhead camera was placed in one part of the environment. A lamp was mounted on the central axis of the circular base of the robot, so that when ARNE moved under the camera its presence at the goal was signaled. The camera is above the environment and therefore cannot be observed by the robot's sonar sensor.

———Figure 5 here———

The neocortex

Processing of sensory data. The simulated neocortex processes sonar data to produce a map consisting of three types of feature: walls, free space, and goals. A wall feature is a line segment corresponding to part of a straight wall in the environment, and is represented by the cartesian coordinates of its endpoints. In order to represent free space a 72x72 grid is superimposed on the plan view of the space surrounding the robot. When the robot is initially placed in an environment each of the grid locations is marked as unknown, and after a location is known not to contain obstacles it is marked as free. The location of a goal in an environment is represented by cartesian coordinates. The algorithms which are used to extract stimuli from raw sonar data are quite complex, and are only described in outline here. They have been described fully elsewhere (Lee & Recce, 1996; Harris *et al.*, 1996).

After a sonar scan is taken, the entire scan is analysed, and the new data are incorporated into the egocentric map. For each distance reading, the egocentric map, held in working memory, is searched for a wall feature which might correspond to the reading. If a corresponding wall feature is found, then its position is updated to take the new reading into account. If no corresponding wall is found in the working memory, then the reading is compared with readings taken from the previous two viewpoints. If three consecutive readings have been observed, which are consistent with the existence of a new wall feature, then this feature is added to the map. The decision on whether to combine new readings is made using mathematical heuristics that are based on the egocentric position of the feature and the length and angle of the sonar reading (see (Harris *et al.*, 1996) for full details). When a reading is ascribed to a new or existing wall, all free space map grid segments along the path of the sonar beam are marked as free.

Working memory. The working memory contains the egocentric map, which is represented as a list of observed features. Once found, an observed feature remains in the map until the robot is removed from the environment. Other proposed features of working memory, like volatility or limited capacity, were not simulated.

Path integration. Path integration information is obtained from measurement of the amount of rotation of the two wheels (using shaft encoders). After each movement ARNE predicts the expected distance and direction to previously observed features in the environment, and updates the map accordingly.

Coordinate transformations. The population activity of modelled head direction cells was represented by a single floating point number (the angle α). In orientation mode the angle is varied to maximise hippocampal recall, as measured by the number of active place units (see below). In the other two behavioural modes the rotational angle is updated using ideothetic inputs. The feature and free space maps are rotated by this angle, before they are passed to the simulated hippocampus.

Planning and execution of movement. In orientation and exploration modes, the robot's aim is to construct a good map of the environment. Previous work with ARNE evaluated the success of different exploration strategies (Lee & Recce, 1996). For the simulations described here we used a simple wall following strategy, in which the robot takes steps of one ARNE length (25 cm), keeping a wall one ARNE length to its left.

In recall mode, the robot's behaviour depends on the status of the hidden goal. If the goal location is unknown, ARNE continues to follow the wall following strategy described above. Otherwise, ARNE heads directly towards the goal, following the shortest route that passes only through parts of the environment marked as free on the free space map. The shortest path is calculated using a recursive search algorithm (Lee, 1996).

The hippocampus

The auto-associative memory was implemented by a single layer of *place units* operating independently and in parallel. Each place unit is only trained once, and stores a complete copy of the neocortical egocentric map that was active at the time it was trained. A place unit in this implementation corresponds to a distributed firing pattern of all of the place cells active in one place in a neural implementation. Each of the stored egocentric maps in one environment have the same allocentric bearing, but the maps are displaced from each other corresponding the the location of the robot in the environment.

Once a place unit has been trained, its firing is calculated at each time step on the basis of a similarity measure between its stored map and the current rotated egocentric map, that is presented, as input, to the hippocampus. The similarity measure is based on geometrical heuristics (Harris *et al.*, 1996) and is calculated as a floating point number between zero and one. If the similarity measure exceeds a fixed threshold, the place unit fires. The size of place fields and overlap of place cells both depend on the value of the threshold applied to this similarity measure. In the current implementation there is no restriction on the number of place units active at the same time. The place units do not have a graded response, so they are either active or inactive.

When a place unit fires, the map it stores is returned to the cortex. If more than one place unit fires, the maps stored in each active place unit are merged into the cortical map. For each wall feature in the map of a firing place unit, the cortical map is examined for a corresponding wall. If there is a corresponding wall, its position is updated according to the position of the recalled feature. Otherwise, the recalled feature is added to the cortical map. Free space segments are treated similarly: if a free space segment is recalled from the hippocampus, and the corresponding segment in the cortical map is marked as unknown, then the free space mark is added. Features added as a result of hippocampal recall are marked as provisional in the neocortical map, to distinguish them from features directly observed by sonar, as described above.

Experiments

The model has been evaluated using two experiments. For these tests two irregular shaped environments (A and B) were constructed (see figure 6). The shapes were chosen to reduce the similarity of particular feature subsets that might be observed with the sonar sensor. Using only sonar in a rectangular room, the robot would not be able to distinguish between diagonally opposite locations. The environments were sufficiently large so that it was never possible for the sonar system to detect all of the walls from one location.

Experiment 1

The goal of the first experiment was to evaluate place cell firing patterns in the two environments. ARNE was placed in environment A and was allowed to explore for 120 time steps. ARNE was then removed and placed in environment B for 120 time steps. ARNE was then reintroduced to environment A for 80 time steps, and then to environment B for 80 time steps.

During the first visit to environment A, ARNE spent 56 steps in orientation mode before switching to exploration mode. In the first visit to environment B, ARNE spent 55 steps in orientation mode before switching to exploration mode. The switch to exploration mode indicates that the egocentric map has a sufficient quality to be stored in the auto-associative memory.

In the second exposure to environment A, ARNE spent 15 steps in orientation mode before switching to recall mode, and in the second exposure to environment B, ARNE spent 12 steps in orientation mode before switching to recall mode. The switch to recall mode indicates that an egocentric map has been recalled by the modelled hippocampus.

———Figure 6 here———

Figures 6 shows maps of the two environments that ARNE has learned and example firing rate maps of four typical place units. Each row contains four firing maps of a particular place unit. The first two columns show the activity of each of the units during the first experience in each of the two environments, followed by two columns showing place unit activity from the second exposure to the two environments. There were a total of 186 place units created during the four runs in the environment, of which 94 fired only in the first environment, 90 fired only in the second environment, and two fired in both environments. Seventy nine of the place units were active only in one of the two visits to an environment. Nearly all of the place units shown coded for a single location in one environment. There were two exceptions, one of which is shown as cell 4 in figure 6.

———Figure 7 here———

Figure 7 shows the cortical map active at the eleventh time step of the second exposure to environment A. Figure 7A shows only features that have been directly observed by sonar, while figure 7B also shows features that have been added to the cortical map by firing place units. At this stage of exploration, two place units had fired, both of which had been trained in the first exposure to environment A. The map shown in figure 7B demonstrates the pattern completion performed by the auto-associative memory. It is the result of merging the maps stored by these two cells into the sonar-derived map of figure 7A.

Experiment 2

This experiment tested the robot's ability to return to a hidden goal in a previously explored environment. ARNE was introduced to an environment containing a goal, and was allowed to explore for 200 time steps. ARNE was then removed from the environment, and successively replaced at 6 test locations. In each of these test locations the robot started with a different orientation relative to the long axis of the room. On each trial, ARNE was immediately removed on reaching the goal.

ARNE's trajectories in the testing and recall phases are shown in figure 8. In the initial exploration of the environment, ARNE spent 22 steps in orientation mode before switching to exploration mode. During the testing phase, ARNE switched from orientation mode to recall mode after, on average, nine time steps. Note that when ARNE was placed in the test locations, the initial segment of its trajectory was along the wall. Once a sufficient map had been constructed to remember the remainder of egocentric features, ARNE moved directly to the goal.

———Figure 8 here———

Discussion

What do we know now that is new?

We have presented a conceptual framework for a navigational system consisting of an egocentric map and an auto-associative memory, and suggested that a similar system may exist in the brain. In this framework the egocentric map is located in the neocortex, and the hippocampus serves as the auto-associative memory.

We implemented the model on a real mobile robot, and performed two experiments. In the first experiment, we allowed the robot to explore two environments, while monitoring the activity of place units in the auto-associative memory. We found that most place units showed localised place fields in one of the two environments. Furthermore, the place fields were stable, and did not change when the robot was reintroduced into the environment. Examination of the modelled neocortical map active a short time after reintroduction into a familiar environment showed that the modelled hippocampus is capable of performing the desired function of pattern completion.

Nearly all simulated place cells fired only in one of the two environments. However, the wall and free space features that constitute the representation of space in the simulated neocortex occur in both environments. A parallel may be drawn to the results of Quirk *et. al.* (1992), which found that, while hippocampal place cells usually fire only in one of two similar environments, entorhinal cells usually fire in both. A direct comparison is impossible, however, as our symbolic implementation of the neocortex does not allow for an analog of entorhinal unit recording.

In the second experiment, the robot was allowed to explore an environment containing a hidden goal. The robot was then reintroduced into the environment at various test locations, and with varied orientations. After a short period of exploration, the robot headed along a straight path to the goal location. During the exploration it determined its position and orientation using egocentric maps previously stored in auto-associative memory. This experiment demonstrated that the robot is able to use this memory to perform cue-based localisation, and to recall the egocentric bearing to the goal. In returning to the goal, the robot was able to take short-cuts across the center of the environment, through an area of space that it had not previously explored.

Without the modelled hippocampus, the robot would have explored the environment in the same manner as occurred in the initial visit. In this case the robot would have found the goal eventually. The snapshot memory simply shortens the search process.

These experiments also imply several properties of the model, which were not explicitly tested. For example, since the environment is recognised only on the basis of sonar information, if the entire environment had been rotated, in absolute space, both the simulated place fields and the head direction angle (representing simulated head direction cell activity) would have rotated a corresponding amount. This is consistent with the experimental data on the relationship between place cells and head direction cells (McNaughton *et al.*, 1994). Furthermore, in a symmetric environment place units will be active at all points with indistinguishable sensory inputs. This is consistent with some of the place cells described by Sharp and co workers (1990), from an experiment rats search for food in a cylindrical environment with mirror symmetry. However, due to the paucity of sensory input, ARNE is more likely than a rat to find two places indistinguishable.

What couldn't have been accomplished by qualitative reasoning?

In making a model of how a system of neurons performs a particular task, two questions need to be addressed. Firstly, is the model capable of performing the task, and secondly, does the model perform the task in the same way as the brain. To answer the latter question, the model must explain experimental results and predict the results of new experiments. In this respect, qualitative reasoning and simulation are both helpful.

However, to answer the first question, qualitative reasoning alone is of less use. It is easy to underestimate the difficulty of tasks which animals perform apparently effortlessly. Computer simulation helps avoid this problem by forcing the modeller to clearly define the task and model, and by allowing for the performance of models to be quantified. However, many computer simulations still make unreasonable assumptions about the tasks the brain is required to perform. In the example of models of spatial function, it is often assumed that cues can be identified and localised with complete accuracy. By contrast, the sonar data available to ARNE is very unreliable: features are often badly localised or missed altogether (see figure 5). Testing a model with a real robot provides a more satisfactory answer to this question.

On the basis only of qualitative reasoning alone, or simulation without the use of a robot, we could not be confident that the model can guide navigation in the real world. If the the model had not been able to successfully guide the robot, we would know to reject it.

How does this model relate to others?

The present model is proposed as a means of unifying the spatial function and intermediate memory function attributed to the hippocampus. In this section we compare our model with other associative memory models and with a cross-section of other navigational models.

Purely spatial models. The first and most influential model of hippocampal involvement in spatial function is the cognitive map theory of O'Keefe and Nadel (1978). This theory proposed that the hippocampus is specialised for performing geometrical computations in cartesian space, and allows for hippocampal involvement in processing of other information that may be encoded in a spatial manner, such as language. By contrast, our model does not require an allocentric map for navigation and the modelled hippocampus does not perform specialised spatial functions other than the memory of an egocentric map. Therefore in our model the hippocampus could act as a more general memory store.

Neuron-level models. In this paper, we have described a system-level model. We have proposed functions for various brain structures, and implemented these functions using methods suitable for implementation on a digital computer. Other authors have already described neural networks that perform some of these functions. Future work will aim to incorporate these neuron-level models into our model.

The idea that the hippocampus acts as an auto-associative memory was originally proposed by Marr (1971). Marr's model described how the hippocampus could act as a temporary memory store for cortical activation patterns, and described putative roles for each region, cell type and connection. Gardner-Medwin (1976) developed a theory to show that a gradual recall process increases the memory capacity of the auto-associative memory. More recently, Willshaw and Buckingham (1990) implemented and re-evaluated Marr's model. Treves and Rolls (1994) developed a formal model of each anatomical hippocampal region, which builds on the initial model developed by Marr. Gibson and Robinson (1992) proposed a mathematical theory of auto-associative memory recall performance, and Hirase and Recce (1996) applied this theory to investigate the thresholding strategy which results in the highest storage capacity.

In addition to the hippocampus, our model ascribes many spatial processing functions to the neocortex. Droulez and Berthoz (1991) presented a neuron-level model that can perform some of these functions. In their model, the egocentric location of stimuli is represented by neuronal activity in a topographically arranged three layered neural net. The network has two input systems, one carrying egocentric visual information, and the other carrying efference copies of motor commands. The latter input is used by the network to accurately update the location of stimuli, even in the absence of visual input. The network performs the functions of working memory and path integration, which our model requires of the neocortex.

Hetero-associator models. In two proposed models (Sutherland & Rudy, 1989; McNaughton & Nadel, 1990) the hippocampus learns appropriate behavioural responses to configurations of stimuli or, equivalently, it acts as a hetero-associative memory. In configural association theory inputs to the associator are all sensory (Sutherland & Rudy, 1989), and McNaughton and Nadel's (1990) model the inputs are the current location of the rat and the direction of movement. In our model, however, the hippocampus is an auto-associative memory which stores sets of cue configurations, but does not associate responses. When the animal later encounters a subset of the stored cues, the modelled hippocampus performs pattern completion to reproduce the original full set. The recalled cues may also include a goal. Neither of the hetero-associative models includes a memory of the direction and distance to a distant feature or goal.

The model of Touretzky and Redish. Touretzky and Redish (1995) have presented a system-level model hippocampal involvement in navigation, which has some features in common with our model. In particular, they produce non-directional place cells by rotating an egocentric map through the head direction cell population angle. However, there are many differences. For example, in our model, the egocentric position of cues is updated by the neocortex using ideothetic information, while in the model of Touretzky and Redish a separate *path integrator* module is used to update the animals coordinates in allocentric space. In their model the hippocampus is a hetero-associator which learns associations between the egocentric local view and the allocentric path integrated coordinates. In contrast, the hippocampus in our model stores snapshots of cortical activity, including the current egocentric map. We suggest that this auto-associator function is need to unify the proposed spatial and memory functions of the hippocampus.

Path integration models. Various authors have proposed neuron-level models of ideothetic control of head direction cell firing (McNaughton *et al.*, 1991; Zhang, 1996). Our model requires a similar function, and does not conflict with these models.

More recently models have been proposed (McNaughton *et al.*, 1996) to extend this scheme to path integration of the animal's position in allocentric space, and have proposed that the hippocampus is the substrate of this function. In contrast, in our model the hippocampus is an auto-associative memory and the neocortex performs the path integration function.

Competitive learning models. Zipser (1985) proposed that a mechanism of competitive learning (Rumelhart & Zipser, 1986) could be responsible for the formation of hippocampal place cells. The idea that competitive learning, involving feedback inhibition and Hebbian synaptic plasticity, occurs in hippocampal processing can be found in the work of Marr (1971). Treves and Rolls (1994) more specifically proposed that this competitive learning mechanism occurs in the dentate gyrus. Sharp (1991) implemented a computer simulation of a competitive net, and found that the simulated neurons do indeed show realistic-looking place fields. Place cells formed in this way have been used in later work as input to a navigational system in two distinct models (Brown & Sharp, 1995; Burgess *et al.*, 1994). While our model is consistent with the existence of a competitive net in the early stages of hippocampal processing, it differs from these models in the use to which the place cell representation is put. In the model of Brown and Sharp, the place cell representation is used as input to a reinforcement learning system located in the nucleus accumbens, and in the model of Burgess *et. al.* it is used to compute a population vector indicating the direction to a goal. In our model, the place cell representation is fed back to the cortex where it reactivates representations of features present when the place cell was trained. These features could include, but are not restricted to, goals. However, the above models do not necessarily conflict with ours. It is possible that the brain uses all three mechanisms.

To the best of our knowledge ours is the only model that has been tested in several environments, and the only one that has been implemented on a mobile robot. In most other models, all of the cues in an environment must be distinguishable, which makes implementation on a robot more difficult.

What new experimental directions are suggested?

The simulations we have described demonstrate that the robot can perform some of the spatial tasks that are thought to be dependent on the hippocampus. The model also makes some predictions for experiments that have not yet been conducted.

The effect of lesions on path integration ability. In our model, 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. Evidence for the effects of cortical (Save & Moghaddam, 1996) and hippocampal (Matthews *et al.*, 1988) lesions on path integration has been claimed. However to our knowledge no study has quantitatively differentiated the effects of these lesions on path integration performance.

Instantaneous transfer. Our model makes predictions for the behaviour of animals in unexplored parts of familiar environments. In our view, 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 may be able to follow an accurate trajectory to a goal. However, the animal will not have place cells that code for the new location, and so might be expected to engage in exploratory behaviour. By contrast, if the animal is disoriented before being introduced into 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. This suggestion is consistent with data published by Alyan (1994).

Summary

We have shown that Marr's model of hippocampal function is consistent with the involvement of the hippocampus in spatial tasks. In doing this we have attributed well defined spatial roles to regions outside the hippocampus. The model was evaluated with a mobile robot, and was able to construct and use maps with real, rather than simulated, sensory input.

References

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

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

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

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

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

Bisiach, E., & Luzzatti, C. 1978. Unilateral neglect of representational space. *Cortex*, **14**, 129–133.

Brown, M. A., & Sharp, P. E. 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.

Burgess, N., Recce, M., & O'Keefe, J. 1994. A model of hippocampal function. *Neural Networks*, **l7**, 1065–1081.

Chapius, N., & Scardigli, P. 1993. Shortcut ability in hamsters *(Mesocricetus auratus)*: The role of environmental and kinesthetic information. *Animal Learning and Behaviour*, **21**, 255–265.

Chen, L. L., Lin, L.-H., Green, E. J., Barnes, C. A., & McNaughton, B. L. 1994. Head direction cells in the rat posterior cortex. I. Anatomical distribution and behavioral modulation. *Exp. Brain Res.*, **101**, 8–23.

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

de Bruin, J. P. C., Sanchez-Santed, F., Heinsbroek, R. P. W., Donker, A., & Postmes, P. 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., & Berthoz, A. 1991. A neural network model of sensoritopic maps with predictive and short-term memory properties. *Proc. Nat. Acad. Sci.*, **88**, 9653–9657.

Feigenbaum, J. D., Polkey, C. E., & Morris, R. G. 1996. Deficits in spatial working memory after unilateral temporal lobectomy in man. *Neuropsych.*, **34**, 163–176.

Gibson, W. G., & Robinson, J. 1992. Statistical analysis of the dynamics of a sparse associative memory. *Neural Networks*, **5**, 645–661.

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

Harris, K. D., Lee, D. C., & Recce, M. 1996. Environment memory for a mobile robot using place cells. *In: Proceedings of Fourth International Conference on Simulation of Adaptive Behavior (SAB96).*

Hirase, H., & Recce, M. 1996. A biologically simple thresholding strategy in an associative memory performs best. *Network: Computation in Neural Systems*, submitted.

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

Jung, M. W., & McNaughton, B. L. 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., & McVety, K. M. 1988. Latent place learning in a novel environment and the influences of prior training in rats. *Psychobiology*, **156**, 146–151.

Kolb, B., Buhrmann, K., McDonald, R., & Sutherland, R. J. 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.

Lee, D., & Recce, M. 1996. Quantitative evaluation of the exploration strategies of a mobile robot. *Int. J. Robotics Research*, in press.

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., & Durrant-Whyte, H. F. 1992. *Directed sonar sensing for mobile robot navigation.* Cambrage, MA: Kluwer Academic Publishers.

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

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

Marr, D. 1970. A theory for cerebral neocortex. *Proc. R. Soc. Lond. B*, **176**, 161–234.

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

Matthews, B. L., Campbell, K. A., & Deadwyler, S. A. 1988. Roatational stimulation disrupts spatial learning in fornix-lesioned rats. *Behavioural Neuroscience*, **102**, 35–42.

McNaughton, B. L., & Morris, R. G. M. 1987. Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends in Neuroscience*, **10**, 408–415.

McNaughton, B. L., & Nadel, L. 1990. Hebb-Marr networks and the neurobiological representation of action in space. *Pages 1–64 of:* Gluck, M. A., & Rumelhart, D. E. (eds), *Neuroscience and Connectionist Theory.* Lawrence Erlbaum Associates.

McNaughton, B. L., Barnes, C. A., & O'Keefe, J. 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., Chen, L. L., & Markus, E. J. 1991. Dead reckoning, landmark learning, and the sense of direction - a neurophysiological and computational hypothesis. *Journal of Cognitive Neuroscience*, **3**, 192–202.

McNaughton, B. L., Knierim, J. J., & Wilson, M. A. 1994. Vector encoding and the vestibular foundations of spatial cognition. *Pages 585–595 of:* Gazzaniga, M. (ed), *The Cognitive Neurosciences.* Cambridge, Mass.: MIT Press.

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

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

Mizumori, S.J.I., & Williams, J.D. 1993. Directionally selective mnemonic properties of neurons in the lateral dorsal nucleus of the thalamus of rats. *J. Neurosci.*, **13**, 4015–4028.

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

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

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

Muller, R. U., Bostock, E., Taube, J. S., & Kubie, J. L. 1994. On the direction firing properties of hippocampal place cells. *Journal of Neuroscience*, **14**, 7235–7251.

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

O'Keefe, J., & Conway, D. H. 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., & Dostrovsky, J. 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., & Nadel, L. 1978. *The Hippocampus as a Cognitive Map.* Oxford, UK: Clarendon Press.

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

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

Pigott, S., & Milner, B. 1993. Memory for different aspects of comples visual scenes after unilateral temporal- or frontal-lobe resection. *Neuropysch.*, **31**, 1–15.

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

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

Rumelhart, D. E., & Zipser, D. 1986. Feature discovery by competitive learning. *Pages 151–193 of:* Rumelhart, D. E., & McClelland, J. L. (eds), *Parallel Distributed Processing 1.* MIT Press.

Save, E., & Moghaddam, M. 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.

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

Sharp, P. E. 1991. Computer simulation of hippocampal place cells. *Psychobiology*, **19**, 103–115.

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

Sharp, P. E., Muller, R. U., & Kubie, J. L. 1990. Firing properties of hippocampal neurons in a visually-symmetrical stimulus environment: contributions of multiple sensory cues and mnemonic processes. *J. Neurosci.*, **10**, 3093–3105.

Smith, M.L., Leonard, G., Crane, J., & Milner, B. 1995. The effects of frontal-lobe or temporal-lobe lesions on susceptibility to interference in spatial memory. *Neuropysch.*, **33**, 275–285.

Squire, L. R. 1992. Memory and the hippocampus: a synthesis from findings with rats, monkeys and humans. *Psych. Rev.*, **99**, 195–231.

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

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

Sutherland, R. J., Chew, G. L., Baker, J. C., & Linggard, R. C. 1987. Some limitations on the use of distal cues in place navigation by rats. *Psychobiology*, **15**, 48–57.

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

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

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

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

Thompson, L. T., & Best, P.J. 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.

Tolman, E. C. 1932. Cognitive maps in rats and men. *Psychol. Rev.*, **55**, 189–208.

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

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

Wehner, R., Michael, B., & Antonsen, P. 1996. Visual navigation in insects - coupling of egocentric and geocentric information. *J. Exper. Biology*, **199**, 129–140.

Whishaw, I.Q. 1991. Latent learning in a swimming pool lace task by rats: evidence for the use of associative and not cognitive mapping processes. *Quaterly Journal of Experimentimental Psychology*, **43B**, 83–103.

Wiener, S. I. 1993. Spatial and behavioral correlates of striatal neurons in rats performing a self-initiated navigational task. *J. Neurosci.*, **13**, 3802–3817.

Willshaw, D. J., & Buckingham, J. T. 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., & McNaughton, B. L. 1993. Dynamics of the hippocampal ensemble code for space. *Science*, **265**, 1055–1058.

Zhang, K. 1996. Representation of spatial orientation by the intrinsic dynamics of the head-directon cell ensemble: A theory. *J. Neurosci.*, **16**, in press.

Zipser, A. 1985. A computational model of hippocampal place fields. *Behav. Neuroscience*, **99**, 1006–1018.

Zipser, D., & Andersen, R. A. 1988. A back-propagation programmed neural network that simulates response properties of a subser of posterior parietal neurons. *Nature*, **331**, 679–684.

Figure Captions

Figure 1 A block diagram of the navigational model. The neocortex processes sensory input to deduce the egocentric positions of objects in the world. These are held in short-term memory in the form of an egocentric map, and updated using ideothetic input. The modelled hippocampus is an auto-associative memory which stores snapshots of the egocentric map and performs pattern completion on fragmentary egocentric maps.

Figure 2 Behavioural modes. When the animal is first placed in an environment, it is in orientation mode. After a good egocentric map has been constructed in the neocortex, the animal switches into either exploration mode or recall mode, depending on the quality of hippocampal recall. If the animal is in exploration mode and recall quality becomes high, it switches to recall mode.

Figure 3 This figure illustrates conversion from the egocentric cortical map to place coding, and back. The angle α represents the animal's allocentric head direction as coded for by the population activity of head direction cells. Egocentric map input from the modelled neocortex is rotated through $-\alpha$ before reaching the hippocampus and by α before returning to the neocortex. Each of the place units operates independently, and the egocentric maps stored in active place units are combined before they recalled memory is returned to the neocortex. The behavioural mode input determines the plasticity of place units. In orientation mode the head direction code changes in order to maximise the recall quality in the hippocampus.

Figure 4 The robot, ARNE, used to test the proposed navigational model. The light above ARNE is detected using an overhead camera. Directly below the light, and mounted at the centre of ARNE's circular base there is a single ultrasonic sensor on a pivot mount, that is used to find the distance to features in the environment. The wheels are centred to allow ARNE to turn on the spot.

Figure 5 A typical sonar scan, superimposed on a map of the environment. The 19 radial dashed lines are individual sonar readings, taken with an angular spacing of 18^o. No reading is taken directly behind the robot. The sonar readings are superimposed on a plan view map of the environment, in which each of the solid lines is a wall. To guide the eye, a dashed line has been drawn connecting the sonar reflection points. The sonar readings that appear to pass through the wall are specular reflections, that result from a small angle between the direction of travel of the beam and the wall. The readings fall short of the wall when the nearest point in the spreading sonar beam is not at the centre of the beam. The longest wall is 4.7 meters, while the maximum range of the sonar is 2.5 meters.

Figure 6 Firing rate maps for simulated place cells. 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 ARNE steps and the second run lasted 80 ARNE steps. The firing rate maps are computed as the ratio of the number of spikes fired over the time spent in each location.

Figure 7 An example of pattern completion, taken from the eleventh time step of the second exposure to environment A. A) Only features that have been directly detected by sonar. B) The complete map, including features returned to the neocortex by active place cells. The filled circle represents ARNE. The lines represent walls. The speckled area represents free space.

Figure 8 The trajectories followed by ARNE in the second experiment. 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. When the robot was tested with the auto-associative memory disabled, it displayed a similar trajectory to that shown in A.

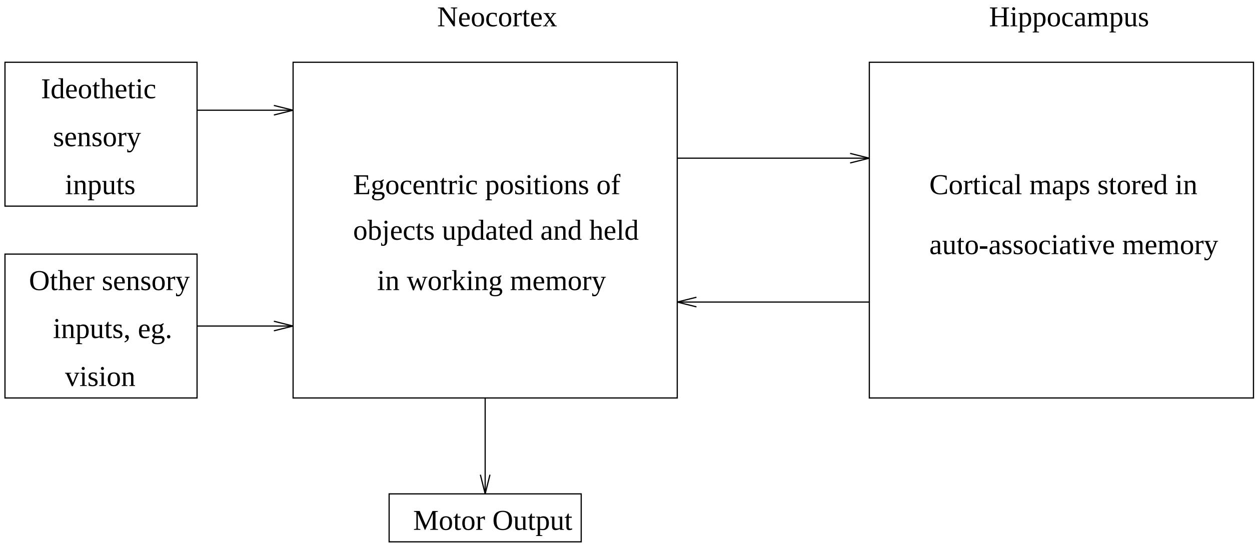
Neocortex
Hippocampus
Ideothetic sensory inputs
Other sensory inputs, eg. vision
Egocentric positions of objects updated and held in working memory
Cortical maps stored in auto-associative memory
Motor Output

Figure 1:

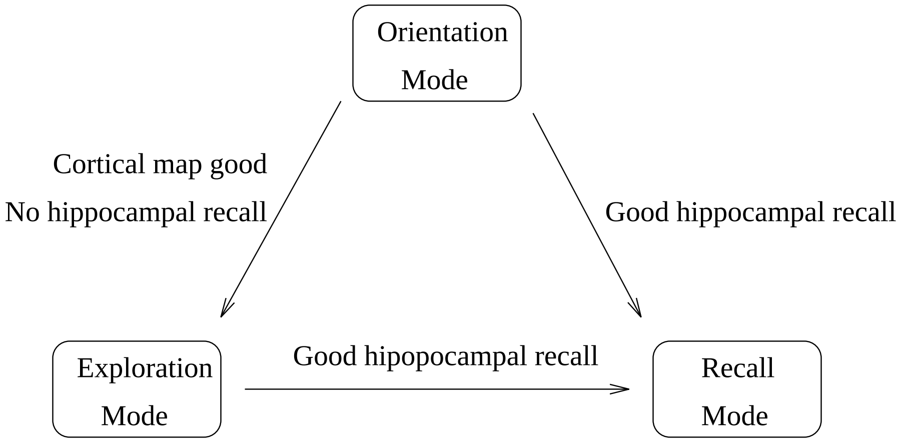
Orientation Mode
Cortical map good
No hippocampal recall
Good hippocampal recall
Exploration Mode
Good hipopocampal recall
Recall Mode

Figure 2:

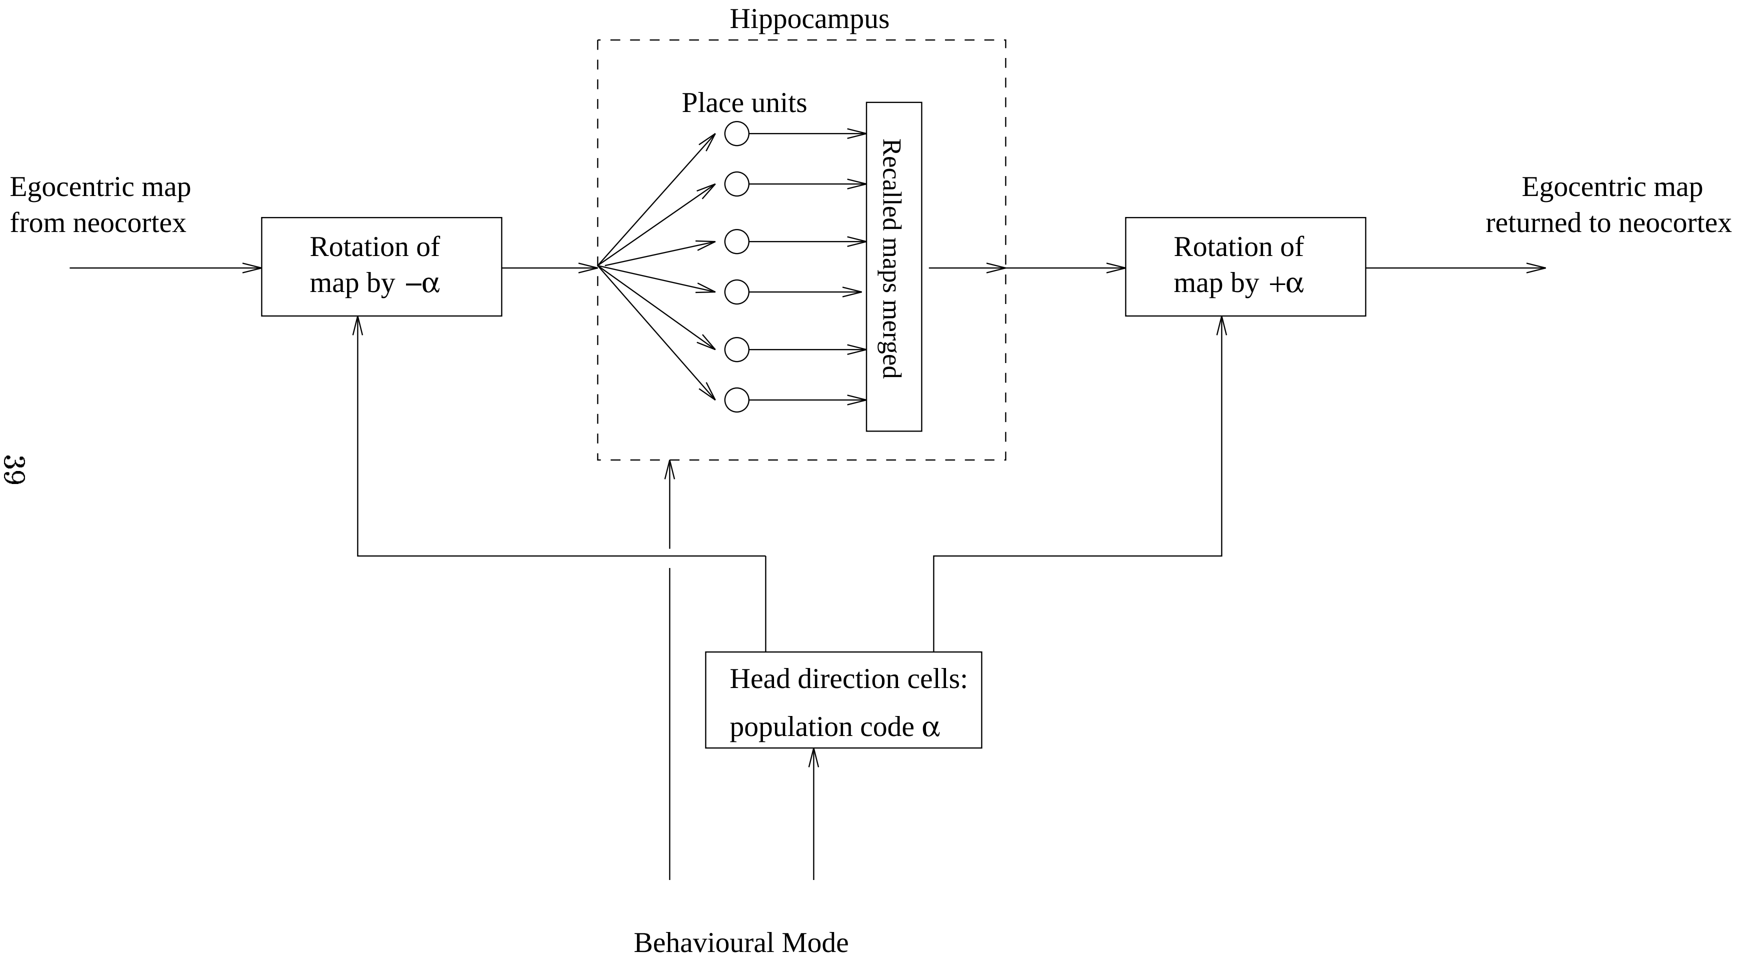
Hippocampus
Place units
Recalled maps merged
Egocentric map
from neocortex
Rotation of
map by −α
Rotation of
map by +α
Egocentric map
returned to neocortex
Head direction cells:
population code α
Behavioural Mode

Figure 3:

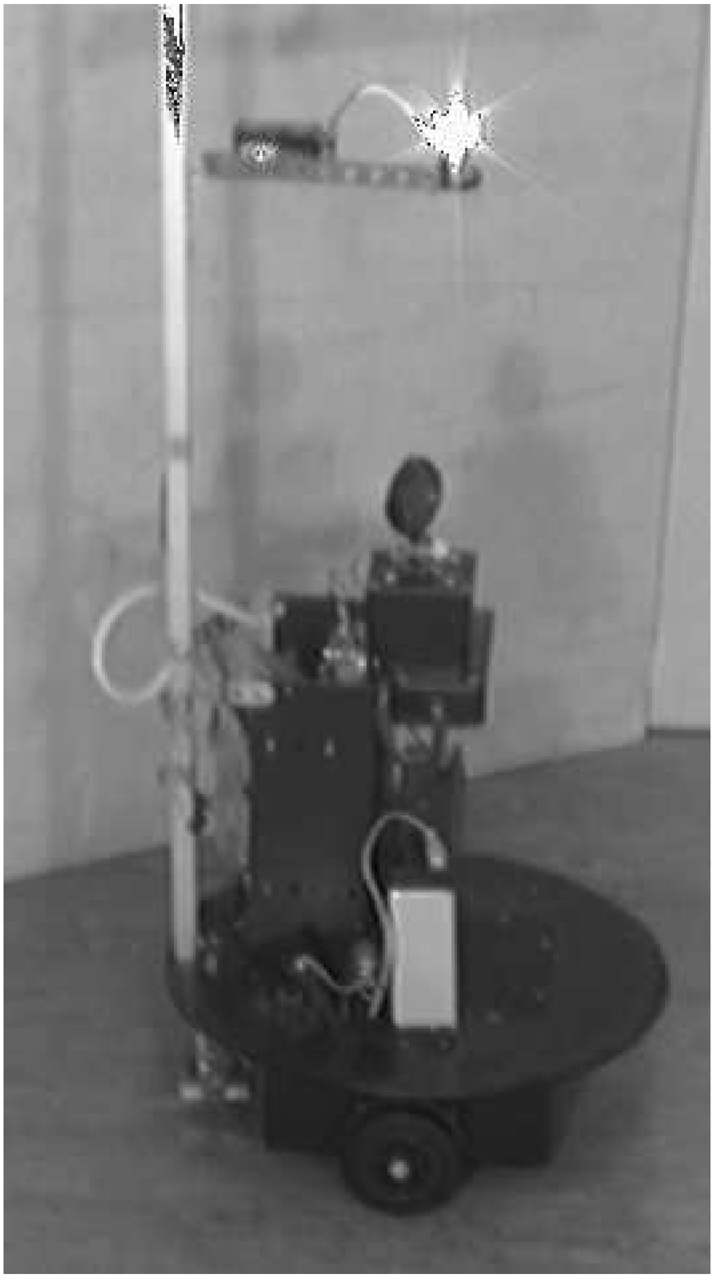

Figure 4:

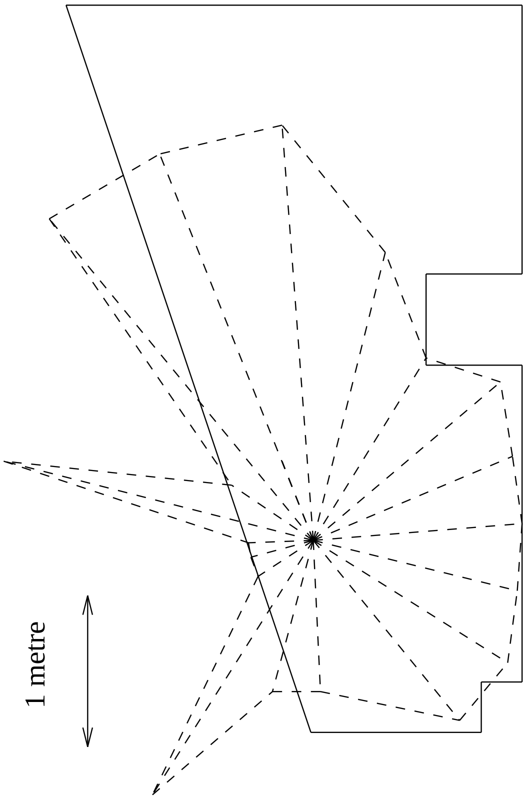

Figure 5:

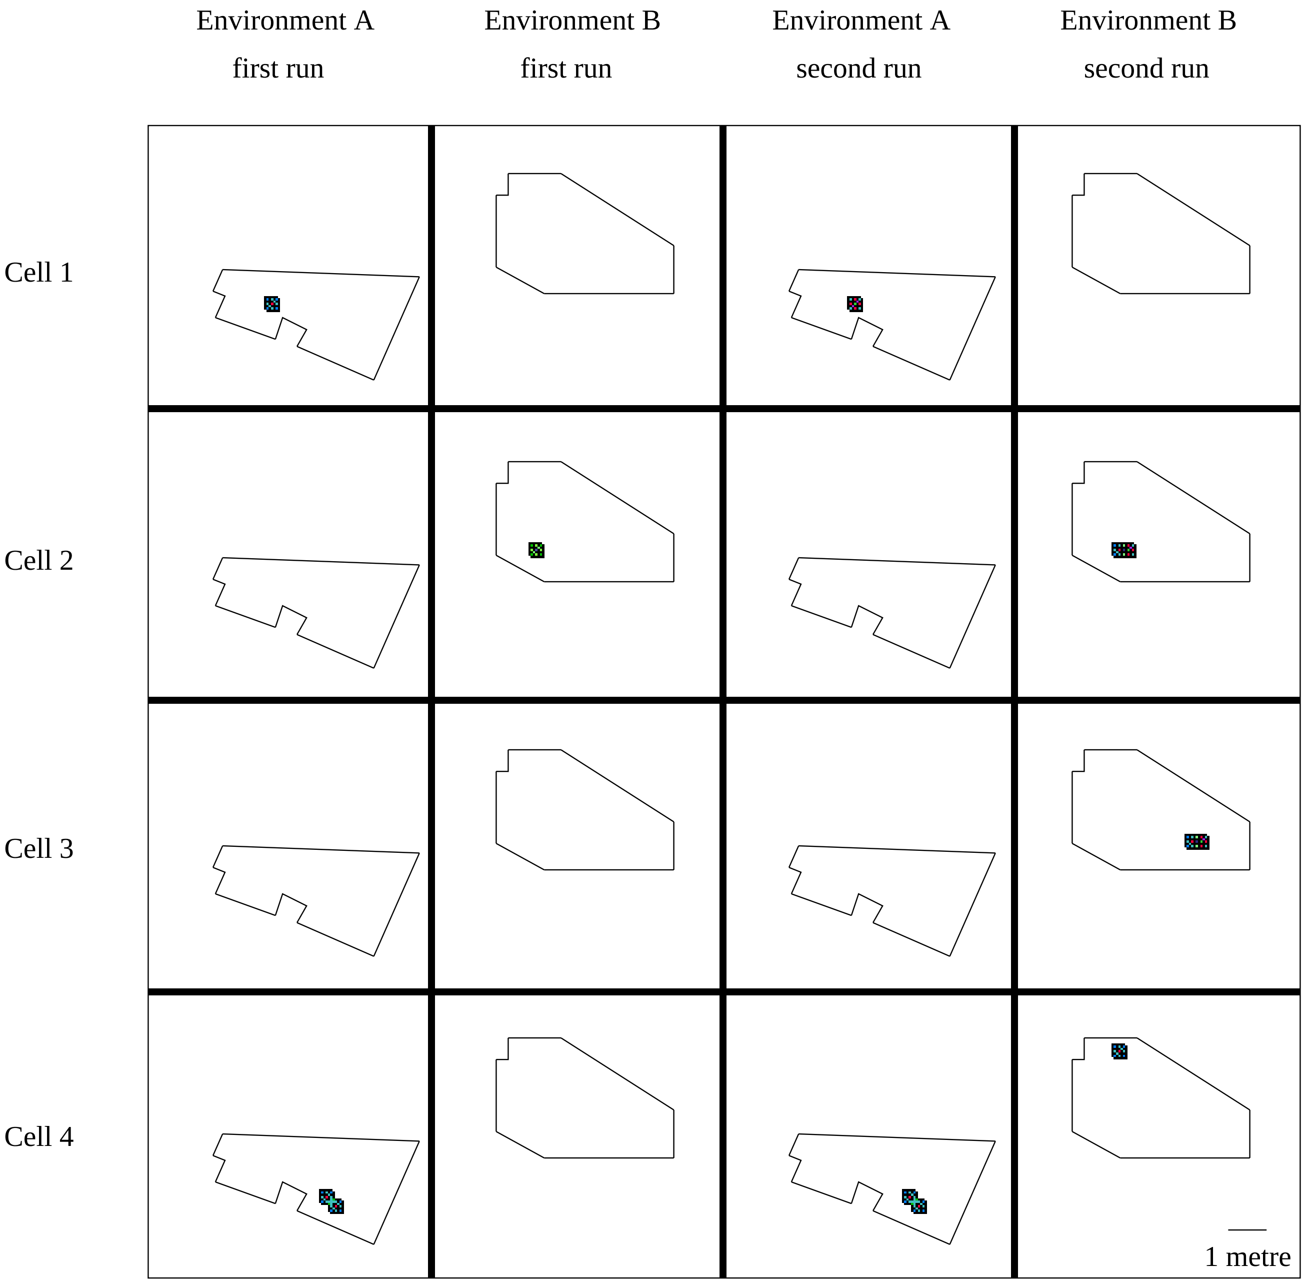
Environment A first run
Environment B first run
Environment A second run
Environment B second run
Cell 1
Cell 2
Cell 3
Cell 4
1 metre

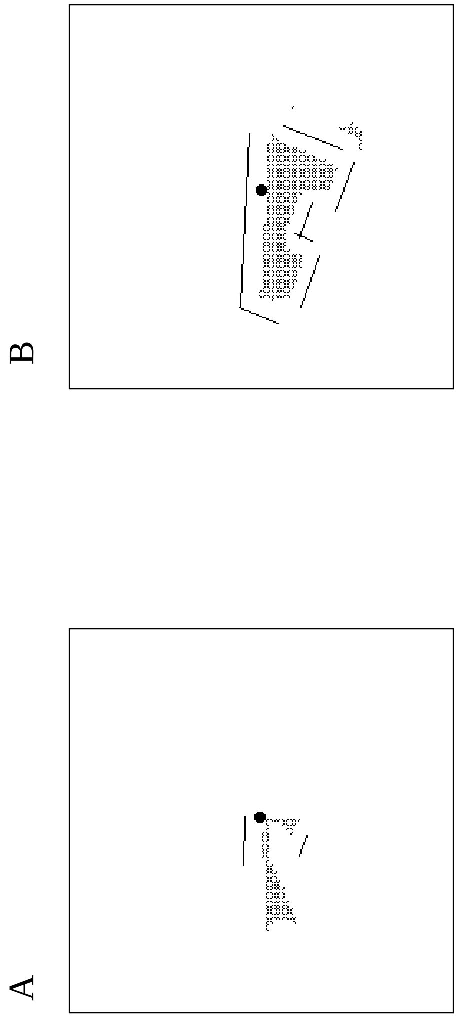

Figure 7:

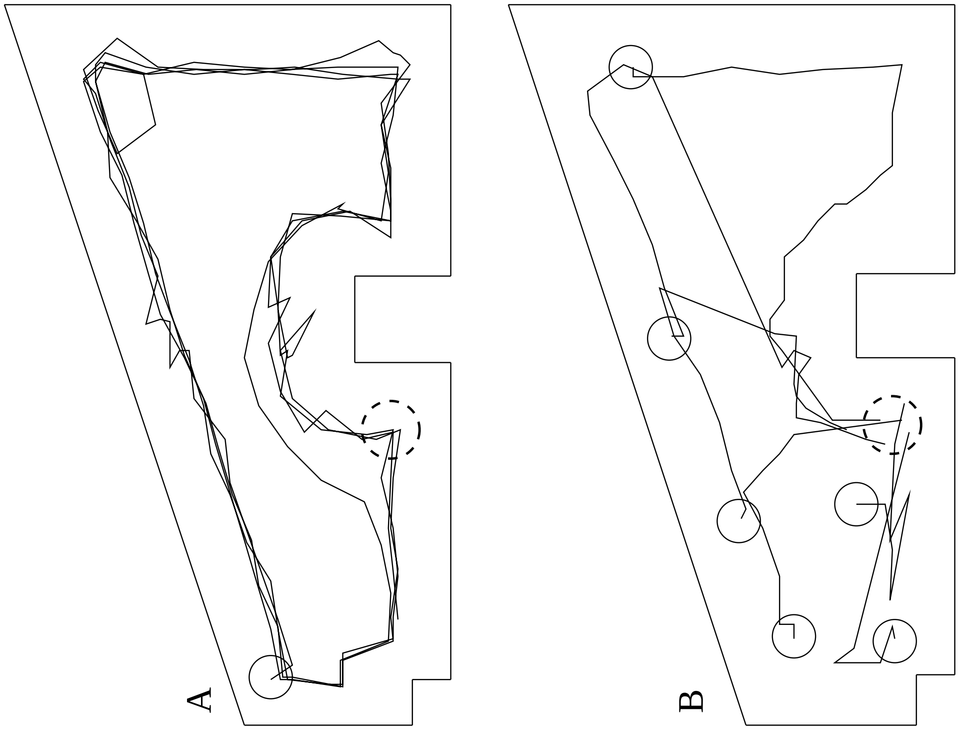

Figure 8: