

Cognitive Models of Spatial Navigation from a Robot Builder's Perspective

Gordon Wyeth and Brett Browning

Department of Computer Science and Electrical Engineering

The University of Queensland

Brisbane Qld 4072, Australia

wyeth@elec.uq.edu.au

browning@elec.uq.edu.au

Phone: +61 7 3365 3770

Fax: +61 7 3365 4999

Adaptive Behaviour Journal

Special Issue on Biologically Inspired Models of Spatial Navigation

November 12, 1997

Abstract

Complete physically embodied agents present a powerful medium for the investigation of cognitive models for spatial navigation. This paper presents a maze solving robot, called a micromouse, that parallels many of the behaviours found in its biological counterpart, the rat. A cognitive model of the robot is presented and its limits investigated. Limits are found to exist with respect to biological plausibility and robot applicability. It is proposed that the fundamental representations used to store and process information are the limiting factor. A review of the literature of current cognitive models finds a lack of models suitable for implementation in real agents, and proposes that these models fail as they have not been developed with real agents in mind. A solution to this conundrum is proposed in a list of guidelines for the development of future spatial models.

1 Introduction

This paper presents a complete physically embodied agent for a complex spatial navigation task; namely exploring and solving a large maze. This agent was designed from an engineering perspective, rather than the biologically inspired designs addressed in this special issue. However, the process of engineering design reveals key aspects that are relevant to the design of biologically plausible cognitive models of spatial navigation tasks. These aspects range from issues in sensor and motor interfacing to issues in the representation of space. This paper reviews current biologically inspired models of maze navigation and contrasts their design with practices adopted in our engineered solution. We argue that in order to develop accurate and useful models of spatial navigation, such models should be evaluated in a framework that embodies the agent in space and time (whether real or realistically simulated) and interfaces to sensors and actuators (either real or realistically virtual).

Our robot, CUQEE III, is a micromouse (see Figure 1). The robot was developed to compete in national and international micromouse competitions: competitions where maze solving robots compete against the clock to find the fastest path to the centre of an unknown maze, and then run the remembered path. CUQEE III has been highly successful at these contests, winning the Australian contests in 1995 and 1996, and winning the prestigious international APEC contest in 1997. The robot is self-contained, performing all necessary computation for maze navigation, exploration and solving on its own microprocessor in real time. Real time for this robot means generating robust behaviour as it moves at speeds of up to 5ms^{-1} , making decisions as it passes through junctions at around 1ms^{-1} . The cognitive processes developed for this robot produce rapid and reliable spatial navigation behaviour, and, in contrast to biological agents, the processes are well understood.




Figure 1: Photograph of CUQEE III. The robot measures 250 mm across the front sensor wing, but the body fits in the palm of the hand. The robot weighs a mere 260 g.

The behaviour observed in the robot bears similarities to the maze navigation behaviour observed in rats. As robot designers, we have had to address issues in generating these behaviours and their integration in a complete functioning system. This process of *downhill invention*¹ has produced a well understood complete model of spatial navigation, which we will present in this paper. The robot has an *explore* behaviour that is used to investigate the maze to build a spatial representation, and a *navigating* behaviour that is used when negotiating sections of the maze that appear in the spatial map. The robot also has low level reactive behaviours for corridor following and corner navigation, and higher level motivational behaviour for generating the best contest winning strategy. As we will later show, parts of the overall behaviour bear similarities to specific behaviours of a rat in a maze, as well as the obvious similarity in the overall behaviour. We will also highlight the key limitations of the robot, and illustrate the reasons for those limitations.

We now seek to extend our model beyond its current limitations, and have turned to the literature for inspiration from biology. Our review of literature contrasts our experience with the downhill invention of a complete agent with other less complete models for the generation of individual behaviours or cognitive processes derived from the uphill analysis of the behaviour and corresponding neuroanatomy of the rat.

The following section (Section 2) describes the robot and its task. A cognitive model of the robot is presented, based on our knowledge as the robot's creators. The behaviour of the robot is then compared and contrasted with rat behaviour, and the extent of the biological similarity established. Based on these findings, we propose a number of benefits that might be gained from a neural implementation of the robot's cognitive model. Section 3 reviews current cognitive models of spatial navigation and investigates them with respect to the cognitive model of the robot. The review shows that these models could not form the basis of a complete navigation system, and highlights the difficulties with making a complete biologically plausible agent using these models. The final section (Section 4) proposes a framework for the development of complete biologically plausible models of spatial navigation.

2 The Robot

This section presents a description of our maze solving robot, CUQEE, that illustrates the internal workings of a real spatial navigating agent. The task of the robot is presented as an overview of the rules for the micromouse contest. A description of the physical construction

¹A term used by Valentino Braitenberg [Braitenberg, 1984] to highlight the pleasures of invention, in contrast to the hard work of analysis (*uphill analysis*).

of the robot follows, sufficient to describe the operation of the robot from a behavioural perspective. We present a cognitive architecture of the robot that shows the operation of the various processes and algorithms required to control a contest winning robot.

The description is then compared with CUQEE's biological counterpart: the rat. We present the key similarities and differences between CUQEE's behaviour and the behaviour of a rat in a maze. This comparison sets the scene for a discussion of current spatial navigation models from our perspective as robot builders, which follows in Section 3.

2.1 The Micromouse Contest

The overall behaviour of the robot is defined, in many respects, by the rules of the micromouse contest. The contest takes place in a large maze (Figure 2) with reconfigurable walls. The robot tries to find the centre of the maze (the goal), starting from one of the corners facing in a clockwise direction. After finding the goal, the robot will continue to explore until it has found the fastest path from the start to the goal. With that solution in mind, the robot then runs as quickly as possible from start to goal. The winner of the contest is the robot that achieves the lowest score, where the score is calculated based on the time for final run, plus a component of the exploration time. Typically the score formula is something like:

$$\text{Score} = \text{Run time} + 1/30 \text{ Explore time}$$

The maze itself has 16 x 16 cells, with each cell being 180 x 180 mm in dimensions. The floor of the maze is painted matt black, and the walls are white with red tops. The configuration of the walls is revealed only after all robots have been handed in to the judges. The robot dimensions are limited to 250 x 250 mm, allowing the robot to look at the top of walls for navigation purposes. The robot may not jump over or knock down walls. The robots must be completely self contained with no external sources of power or computation. The robot handlers may not provide strategies to the mouse during the starting procedure.

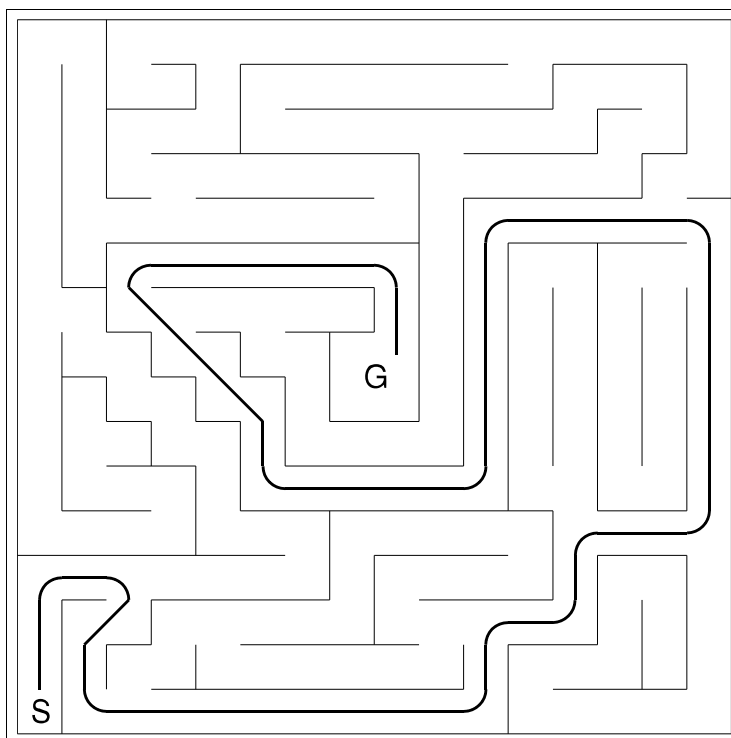


Figure 2: The maze used in the 1996 Australian micromouse championships.

2.2 Overview of the robot hardware

CUQEE III interacts with the maze through its sensors and actuators (Figure 3). The actuators are a pair of drive wheels arranged in wheelchair fashion on either side of the robot. The sensors are arranged to hang over the walls to detect the distance from the body of the robot to walls on either side and straight ahead. The sensors also detect whether or not a wall is present under the sensor. The side sensors have a range of 80 mm and the front sensor has a range of 100 mm. Odometric sensors are also attached to each wheel to detect the distance travelled and current velocity of the robot. It must be noted that the odometers can only supply estimates, as the robot often slips as it moves with high velocities and accelerations. As odometers are on both wheels, both location and orientation information can be derived by path integration. There is no compass sensor to explicitly define orientation. The robot weighs 260 g, and fits in the palm of your hand.

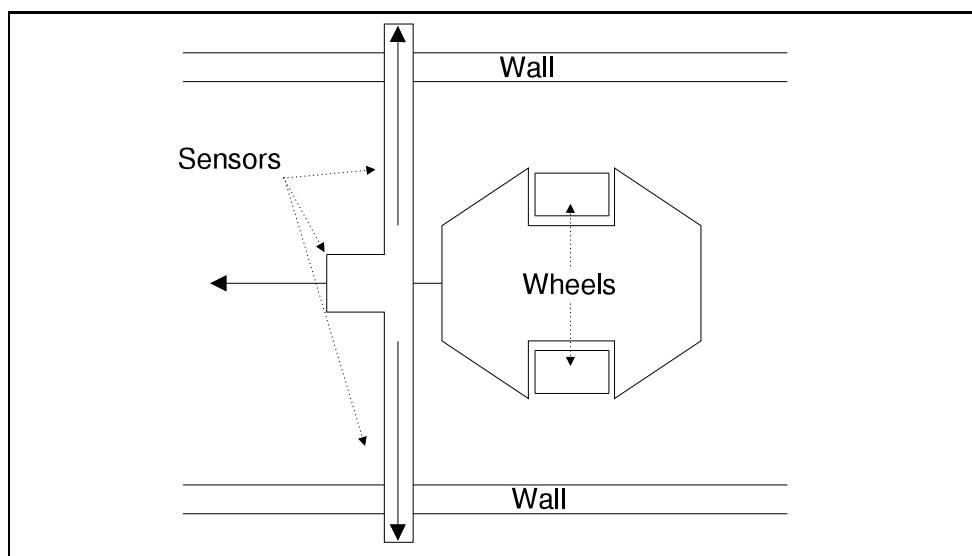


Figure 3: The mechanical layout of CUQEE.

The on-board microprocessor is not computationally powerful by modern standards (around 1 MIPS (Million Instructions Per Second)). The processor is not only responsible for the cognitive processes involved in navigating and solving the maze, but must also perform low level functions such as sensor reading and motor control. The low level functions consume about 50% of the processor time leaving 0.5 MIPS to perform the cognitive processes. This lack of processing power combined with the high speed at which the robot moves has lead to a parsimonious approach to algorithm design. This parsimony in turn has produced an elegant solution to the spatial navigation problem.

2.3 Cognitive architecture of the robot

In this section we present the cognitive architecture of the robot in a way that is consistent with our knowledge of the operation of the underlying algorithms and processes. We use the term *cognitive architecture* to refer to the complete set of processes involved in combining perception and memory of perception into the correct action to perform the maze navigating task. The architecture (Figure 4) has three levels of competence: a schema level, a cognitive level and a motivational level.

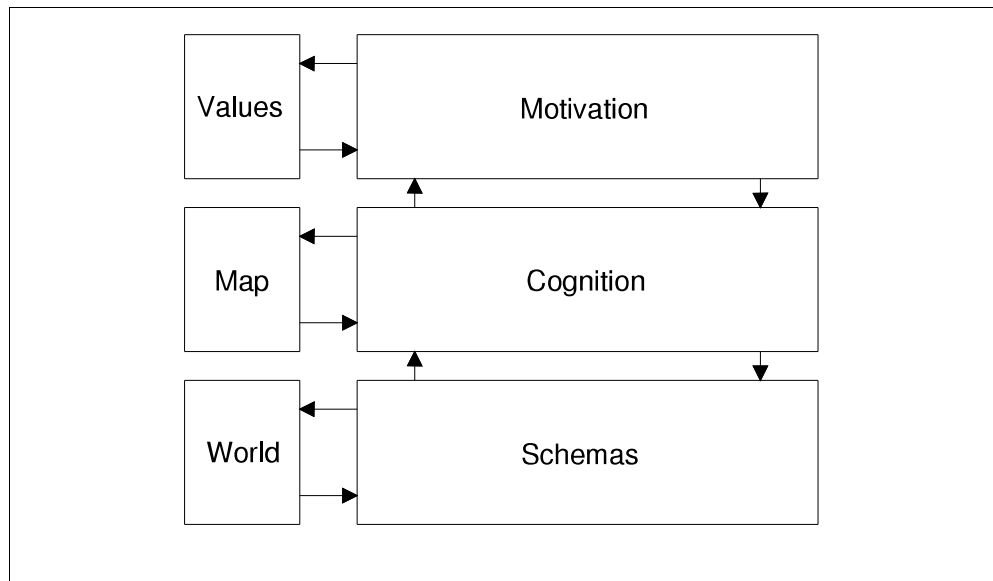


Figure 4: The cognitive architecture of CUQEE has three levels. The lowest level is implemented as schemas which interface in a reactive manner with the world. The cognitive level instantiates schemas to perform the spatial navigation task. The cognitive level operates virtually with a cognitive map of the maze. A motivational level generates goals, and determines the contest winning strategy.

The lowest level of competence is the schema level. The schemas used in this level bear resemblance to those used in psychology, neurology and brain theory. The term *schema* is probably overused: we use a definition that relates strongly to its use in robotics [Arkin, 1989], but we add some restrictions to the definition that are important to our architecture.

A schema is an adaptive reactive controller that uses primitive interactions with the environment to generate a pattern of action.

The use of the term *reactive* here implies that schemas have only limited state information [Gat, 1994], and do not generate or access memory. Schemas rely on interaction with the environment and one another to produce emergent behaviour, as is the case for so-called behaviour-based robots ([Brooks, 1990], [Payton, 1986], etc). Schemas are instantiated by the more competent cognitive level. A typical schema instantiation might be "travel down a corridor for three squares". The schema is then responsible for keeping the robot centred in the corridor, and observing kinaesthetic values and sensor readings to determine when the schema should terminate. Note that a schema could not be instantiated by "go to square 12, 10" as this would require knowledge of the internal representation of the world and the robot's position in the world: that is, the action would not be reactive.

The cognitive level manages the issues of representation and planning. It is responsible for maintaining the map of the maze and using that map to plan action by selecting schemas. Action is planned by generating solutions to the maze with respect to the current goal. The cognitive level maintains a virtual presence of the robot in the internal representation of the maze. By evaluating the solutions at the point of virtual presence the robot can plan action to achieve the goal. The cognitive processes act at a tactical level, generating plans to meet the immediate needs.

The definition of the goal and the speed at which the robot travels are strategic issues. The strategy is defined by the highest level of competence, which we have dubbed the motivational level. A key factor in the success of the robot is its ability to choose when to quit exploring and generate a fast run, and its ability to choose an appropriate speed. To quit exploring too early may lose the robot valuable seconds in its fast run, but to continue exploring may be wasteful and cause a profitless increase in exploration penalty. Similarly the

robot must judge its speed as does a racing driver: too much speed may cause the robot to crash, but if the robot is too conservative it will lose anyway. Speed is determined by the tyre force applied during that run. The robot is always accelerating, decelerating or cornering so that force is always being applied to the tyres. A larger force increases the likelihood of slip and hence accidents; lower forces increase the margin for error.

To best interpret the current situation the robot must maintain a set of values that it can use to judge the most appropriate strategy for a given maze. The motivation level employs its strategy by instantiating the appropriate cognitive processes².

2.3.1 Schema Level

There are five fundamental motor schemas (*turn*, *straight*, *diagonal*, *spot turn* and *wait*) and two perceptual schemas. The motor schemas are competitive, with only a single schema being active at one time. These schemas run in parallel with the perceptual schemas and the higher level cognitive and motivational processes. This ability allows CUQEE to move and solve the maze at the same time. CUQEE never stops moving in the maze while exploring, unless it reaches a dead end. This constant motion is in contrast to most other micromice which stop every time they need to solve the maze. The following summarises the behaviour of each motor schema:

Turn: This schema operates in three phases, a constant velocity phase as the robot rounds the bend at the determined velocity and radius, and two transition phases as the robot enters and exits the turn. The transition phases are required to ensure that a constant force is applied to each tyre as the robot goes from pure translational to combined translational and rotational motion. The schema uses odometric information to control the turn profile, and sensor information in the final phase to align with the new corridor.

Straight : When travelling a corridor the robot uses its sensor information to maintain its central position, and uses odometric information to decide whether the robot should be accelerating or braking. Sensor information is also used to calibrate the odometers: the spacing of any openings that the robot passes is well defined allowing update of odometric estimates.

Diagonal : Many maze designers add a staircase pattern to the maze (Figure 5), intended to slow the mouse down as it executes the consecutive turns. Most world class micromice are designed to negotiate the staircase pattern in a procedure known as "running the diagonal" [Otten, 1990]. To execute this schema the robot must first turn 45° and then negotiate the diagonals of the consecutive squares. This procedure requires a different approach to corridor sensing and odometric calibration, as the robot encounters a completely different corridor geometry.

²It may be argued that under many interpretations of a schema, the processes in the cognitive level are themselves schemas. For the purposes of this paper, we adopt the definition that schemas are reactive processes that do not access internal representations.

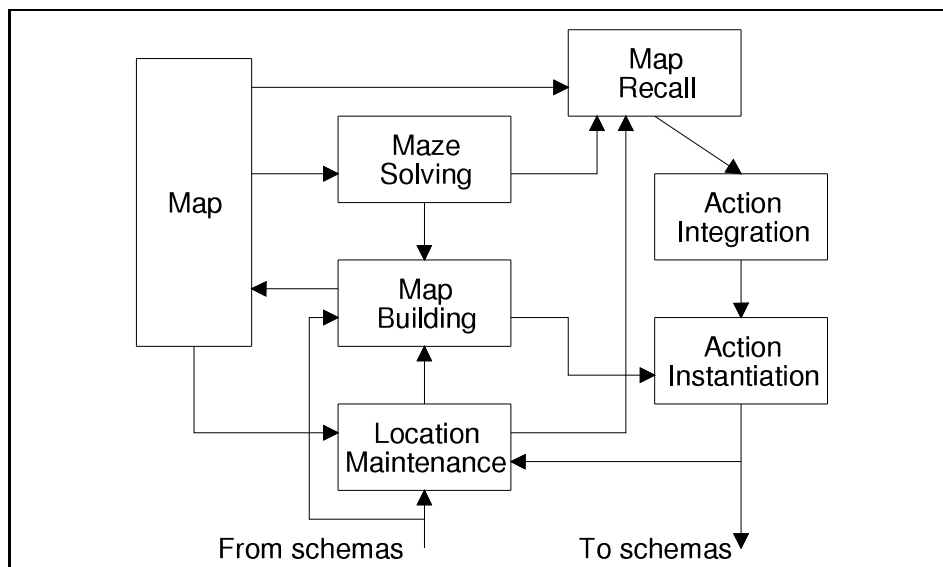


Figure 6: The flow of information between different cognitive processes. The key resource is the map which stores the maze. This map is constructed by a map building process, and can be recalled by a map recall process. The recall module plans paths using the solutions to the maze calculated by the maze solver. Action is generated by an action instantiation module, with action integration during recall. Any physical action of the robot is accompanied by virtual movement in the location maintenance module, providing that the low level schemas indicate a satisfactory execution of action.

Paths to the goal may be found using a flood fill technique. The basic flood fill algorithm is best described in terms of a distance calculation. The aim of the algorithm is to flood the maze with the distance (number of cells) from each square to the current goal. The goal square is given a value of zero, and joining squares are given a value of one. The list of joining squares is recorded, and called the tails. The algorithm then runs through the list of tails and places a two in each square that is adjacent to a tail. Adjacency is determined by checking for the presence of walls in the map. The tail list is then updated to reflect the new adjacent squares and the algorithm reiterates until no tails are left in the list. CUQEE uses a modified version of this algorithm which computes time instead of distance by taking account of the acceleration capabilities of the robot. This ensures that CUQEE takes the path that is shortest in time rather than shortest in distance.

The mouse uses the time from each location in the map to the current goal for the purposes of map construction and maze recall. In order to construct and recall the maze model, the robot requires a sense of its own current location and orientation in the map. The location and orientation of the robot is maintained by performing the actions in the virtual maze as well as the physical maze. Virtual action is performed by reflecting the effect of motor schema instantiation in the map, and ensuring that the schema was effected by verifying information from the perceptual and motor schemas.

Map construction involves not only using perceptual information to build the map, but also involves directing the search through the maze to areas that may contain a faster route to the goal. In areas where no wall information is recorded, the solver assumes that there are no walls. This optimistic assumption directs the mouse towards squares that possibly contain a faster path, while maintaining a minimal explore pattern. As the new maze information comes in, CUQEE re-solves the maze. In fact, CUQEE solves the maze before it enters the next square so that it can continue searching without hesitation. This leaves about 300 ms to complete the flooding algorithm. When CUQEE enters a square that has a choice, it looks at the solution and chooses the square with the lowest flood value, which represents the square that is closest in time to the goal.

Similarly, maze recall uses the flood values to determine the fastest path through the maze. Maze recall uses another feature that recognises that several actions may be integrated into a single action that may be executed more efficiently. For example, if the maze recall process recognises consecutive straight moves, it will continue to evaluate the path through the maze until it encounters a turn. The consecutive straight moves will be integrated into a single straight move that takes advantage of the uninterrupted straight to build up more speed. Similarly, two turns in the same direction are combined to produce a single U-turn, or two turns in opposing directions are used to initiate a diagonal manoeuvre.

2.3.3 Motivation Level

At the highest level, the robot must evaluate its strategy in winning the contest. The primary function of the motivational level is in defining the current goal. A simple strategy is to initially set the centre as the goal, and when that goal is achieved set the start as the goal. This cycle is repeated until the robot notes that it did not perform any exploration on its way to the current goal. This means that the robot has found the fastest path and should increase its speed to attempt to achieve a better time.

This type of strategy is wasteful of explore time as the robot will spend a lot of time in sections of the maze that have already been explored. CUQEE's strategy is somewhat more sophisticated: it generates sub-goals that prevents excessive running through explored sections of the maze. CUQEE may decide to attempt a fast run before it has completely explored all possible options, by noting that the time to explore will probably add more penalty to the score than the benefits of having a marginally shorter path. CUQEE also makes decisions about how fast to run based on the complexity of the maze and the length of the path.

The effort that has been placed at the motivational level has paid off handsomely for CUQEE. In the highly challenging APEC championships, the shortest path still covered over 100 squares in the maze and contained many treacherous combinations of turns. Robots that used the same strategy as they would in a smaller simpler maze either failed to reach the centre with a fast run, or failed on the first increment of speed. CUQEE recognised the difficulty of the maze and proceeded with caution, allowing the robot to reach its top speed to win the contest. Not only did CUQEE have the fastest run, it also explored the maze in about half the time of other robots. This was partially due to the use of sub-goals in exploration, and partially due to the efficiency of running navigation schemas and cognitive processes in parallel.

2.4 Similarities to Rat Behaviour

The behaviour of the robot in the maze, while not neuroethological in nature, shows some similarities to real rodent behaviour. One obvious similarity is that both solve mazes and can recall them for fast runs through the maze. The more specific behavioural similarities between CUQEE and real rodents are of greater interest.

In [O'Keefe & Nadel, 1978] it is proposed that the rodent (and indeed human) hippocampus is the location of a cognitive map. This proposal is based upon their work on the place units within the hippocampus. Their main argument is that locale navigation (point-to-point navigation in an learnt environment) is performed using place codes stored in the hippocampus. They propose that the spatio-cognitive process of locale navigation is built on lower levels of taxon and praxic navigation. Taxon and praxic navigation are reactive behaviours, and as such fall within the definition we use for schemas. The similarity between this hierarchy of biological spatial navigation and the architecture presented for CUQEE is readily apparent.

The arguments of O'Keefe and Nadel add specific details at both a neurological and functional level to the model presented by [Tolman, 1932]. Tolman's argument for a cognitive map was in contrast to popular models of his time in that it argued for an allocentric cognitive map that stores the relations between objects, not so much the relations between the objects and the agent. CUQEE's allocentric map of the maze fits this model of navigation.

Cognitive map based models argue that the allocentric map must contain sufficient information to perform dead reckoning, also called path integration. Path integration is the ability to move in the environment based solely upon kinaesthetic information, without external sensory input. It has been shown ([Tolman, 1932], [Gallistel, 1990], [Redish, 1997]) that rodents are capable of using dead reckoning. Gallistel argues that dead reckoning is a critical feature to any navigational strategy. CUQEE uses dead reckoning in its motor schemas, relying on odometric information to judge when to start decelerating to enter a corner, or when to shift from transition phase to constant velocity cornering when executing a turn schema. The dead reckoning parameters come from the cognitive level based on metrics derived from the map.

The robot behaviour that can result from errors in dead reckoning, bears remarkable similarity to the behaviour observed in rodents. In the case of the rodent experiments, a number of corridors were shortened or lengthened after the animals had learnt to run the maze "fast" ([Gallistel, 1990], [Redish, 1997]). The results were that the rats would make major mistakes (such as running full tilt into walls) that indicated that the rodent was using limited external sensory information, and relying on the internal model to navigate. CUQEE makes similar mistakes when walls are changed from the learnt maze. In such circumstances, the robot will career into the unexpected walls, relying on its internal model of the world rather than its sensors. Similarly, CUQEE ignores unexpected openings and will continue to navigate around non-existent walls as if they are still present. Similar behaviour has been noted in biological agents [Gallistel, 1989].

CUQEE's mapping system stores the same type of information that must be stored in any navigational map structure that is to be successful. That is, the map stores enough information to enable the robot to navigate the environment in an efficient manner. The map effectively stores the landmarks at each place in the maze along with the spatio-temporal relationships between them. In order to do this the map needs to be multi-modal in nature, this bears strong correlation with the observations of rodents ([Gallistel, 1990], [Redish, 1997], [O'Keefe & Nadel, 1978]). Multi-modal sensory data includes both external and internal sensory information. The external information is used to store information about the walls within the environment. The kinaesthetic information helps develop the spatio-temporal relationships within the map. It is this kinaesthetic information, which is used for dead reckoning, that forms the *metric* of the map.

Gallistel argues that cognitive maps have a metric nature, and that metric information is part of the rat's cognitive map. In contrast, the use of kinaesthetic information provides metric information implicitly. Furthermore, this information is directly related to the movement that can be generated by the robot. In the case of CUQEE, the kinaesthetic data at the schema level is represented by how far the wheels have travelled. At the higher levels this is represented by the movement of the current position throughout the map (the set grid size of the cells in the map).

The enforced correspondence of the physical motion and virtual motion that takes place in CUQEE also occurs in rats. In [McNaughton and Nadel, 1990], McNaughton showed that rats that were prevented from moving did not register activity in the hippocampus, even if the rat was carried through the environment. In other words, if the motor schemas do not carry out successfully the rat does not update the cognitive representation. Similarly, CUQEE will

not register an action in the location module unless confirmation of that action returns from the appropriate motor schema.

CUQEE's method of maze solution involves generating local information at each cell about the time required to reach the goal square. These times are generated by starting at the goal and spreading activity to all parts of the cognitive map. [Mataric, 1990] proposes that spread of activation, as used in her path finding algorithm, is biologically plausible and is similar to methods proposed by [McNaughton, 1989]. It should be noted that CUQEE finds the shortest path (in time) to the goal based on its cognitive map. There is evidence that rats find the shortest path [Gallistel, 1990], though it is not made clear whether this is shortest temporally or spatially.

According to [O'Keefe and Nadel, 1978] a rat chooses its goal location based on internal values which are generated in response to the animal's needs and desires (hunger, thirst, shelter, mate). These goal locations are reflected in the hippocampus, but the mechanism that produces them is located elsewhere. Similarly, CUQEE chooses its goal location based on a separate mechanism that evaluates current needs in terms of producing the fastest possible run from the start to the centre of the maze.

2.5 Dissimilarities to Rat Behaviour

The similarities just discussed are not intended to demonstrate that CUQEE's cognitive architecture is an accurate reflection of the rat's actual cognitive model. Rather it highlights that the overall structure of the architecture is *very* successful at solving the mapping problem in the given maze; successful to the extent that it can be argued to be optimal. In fact this optimum is a point of difference between the two. For example, the robot requires only a single pass of the environment to perform learning, while rats typically require several passes to learn the structure of the maze. The robot's search is highly directed and near optimal, whereas rats perform a combination of random and directed search. The robot solved the maze illustrated in Figure 2 in less than one minute, and performed the fast run in 12 seconds. While there are no results for a real rodent in this maze, clearly the robot is at least an order of magnitude faster, if the rodent could learn this maze at all. These are results that are not of particular concern as robot designers, but illustrate some differences in terms of using CUQEE as a cognitive model.

Another condition where the robot demonstrates behaviour quite different to its biological counterpart is disorientation. In this situation, the robot's virtual location disagrees with its actual location in the real world. Real rodents can compensate and then recover from this condition. CUQEE is capable of detecting disorientation but is limited in its ability to recover.

Generally, the robot will become disoriented by losing traction at high speed, which in turn leads to an increase in navigational error. The motor schemas attempt to compensate for these errors, but can not always react sufficiently to prevent a collision with a wall. Collisions cause greater loss of traction which causes the robot to slip again. Frequently, such collisions are fortuitous as they tend to push the robot in the opposite direction of the initial error and remove some of the speed. The motor schemas take advantage of this inadvertent correction and bring the robot back in line. However if a collision is sufficiently large, the schema will recognise that all is lost and that further attempts at navigation are pointless or risk damage to the robot. In this case, the motor schemas will report the status to the cognitive processes which in turn prompt the motivational level. The motivational level strategy in this situation is to remove power from the motors and light an LED that prompts for user intervention. Processes for recovering from disorientation are difficult to implement, and are of little value in the context of the contest.

The robot can also be disoriented by blocking passages along the known shortest path. In this situation, a rat will recover and find a new path. CUQEE has no such provision, as the rules of the contest preclude the changing of the maze. A remapping facility would be trivial to add to the motivational level, providing that the robot could recover from the initial collision.

The most significant difference between the robot and the rat is the robot's reliance on the well defined structure of the maze. The robot cannot navigate in environments that are not readily represented as an orthogonal grid. The rat, on the other hand, can clearly handle any physical arrangement. This discrepancy stems from the heart of the cognitive processes of the animal and the animat - the cognitive map.

The robot's cognitive map is made feasible by the grid structure of the maze. Having a fixed number of cells arranged in a grid makes the storage of topological relationships simple. Having fixed cell widths and orthogonality makes the inference of distance and orientation metrics simple. Having a reliable and accurate internal representation of the environment makes the codification of the navigating processes simple. Take away the regular structure of the maze and the problem requires a quite different solution.

This difference is disturbing from both a cognitive modelling perspective and a robot building perspective. The core of the model is shown to be lacking in both biological plausibility and robotic useability. The solution must be a change in the underlying representation of the cognitive map. A solution based on neuroethological principles lends itself to cognitive modelling and shows promise for revealing a superior robotic engineering approach.

3 Cognitive Models of Spatial Navigation

The previous section shows that CUQEE is a complete maze navigating agent, but that its similarity to real rodents is limited by its internal representation, and hence its applicability as a basis for a cognitive model is limited. The aim of this section is to review the literature in search of solutions or partial solutions in spatial navigation that could be used to form a navigating animat that displays more of the characteristics of navigation in rats. This review focuses on models that can be implemented using neural techniques.

In general, researchers of these models have not set out to build complete systems, or systems that can operate in the framework of a complete system. The intention of these models has been to illustrate specific aspects of spatial modelling and navigation. Nevertheless, as robot builders, we must look to these systems for inspiration in the construction of complete animats. The comments made in this review are based on our perspective as constructionists. As such we are bound to comment on the suitability of the model for use in a complete agent even if this was not the author's intended use of the work.

The review is presented at each of the levels that were identified in CUQEE's cognitive architecture: the schema level, the cognitive level and the motivational level. The review concludes with current models that span two or more levels.

3.1 Schema Level

This review of the schema level of navigation relates to reactive or stimulus-response (SR) models of navigation that are based upon biological evidence. The extraction of behaviour from neuroanatomy is a difficult problem (uphill analysis), and progress in this area mostly relates to animals that use SR as a basis for navigation. An oft-quoted example is the work on the sea slug *Aplysia* [Kandel, 1976] which describes neural mechanisms for various reflex behaviours; the behaviour of the slug is described in terms of neural activity. This idea is taken further by [Beer *et al.*, 1990] with their simulated cockroach *Periplaneta*

Computatrix, based on the neural pathways of the American cockroach *Periplaneta Americana*. The simulated cockroach shows that complex motor schemas such as running with a variable gait can be achieved using neural mechanisms. [Cliff, 1990a] similarly modelled visually guided behaviour in the hoverfly *Syrirta pipiens* based on neuroethological evidence. Cliff dubs this process of building complete simulations of a neuroethological process as computational neuroethology (CNE) [Cliff, 1990b].

These and other CNE examples (for example [Webb, 1996]) are currently limited to the schema level; they do not model cognitive processes based on internal representations. There is no doubt that reactive models can be used to create useful robots. For example, [Connell, 1990] has shown that complex behaviour can emerge from the combination of many schemas (or behaviours) with his soda can collecting robot, Herbert. However there is also no doubt that we are looking for something more in our simulation of the rat navigation behaviour than SR behaviour ([Tolman, 1932], [Gallistel, 1990], [Redish, 1997]).

The CNE models of navigation provide useful building blocks as the perception and action tools of higher cognitive processes. In [Wyeth, 1997] it is shown that simple models based on Braitenberg vehicles [Braitenberg, 1984] can perform these elements of spatial navigation task in structured environments. More complex schemas based on CNE models provide an interface to environments with less structure, by providing better or more relevant perception and more sophisticated motor techniques. However, they do not solve the problem of rat-like navigation which requires an internal representation of the environment, and processes that can build and interpret that representation to instantiate the relevant schemas.

3.2 Cognitive Level

The review of the cognitive level shows a variety of approaches. The groupings shown here are based in the choice of representation used in building the spatial map.

3.2.1 Self Organising Maps

Of the systems that can form spatial maps of an environment based on exploration, the majority use self organising maps (SOMs) [Kohonen, 1995] to generate a distributed neural representation of the environment. [Smart, 1994] and [Nehmzow, 1990] present such models embedded in real robots. The architecture is shown to be capable of dealing with actual sensor data and of generating a representation. However, the representation is limited to distinguishing different landmarks, with no notion of topology or spatio-temporal relationships between landmarks. As such, it can not be used as the basis for path planning or any meaningful cognition. The representation could be used in a reactive fashion. For example, a landmark could be associated with a goal, so that if it were perceived some goal directed response may ensue. However, as we have stated earlier, rats use more than SR behaviour. Rats tend to have "place learning" rather than the "response learning" of the SR style models ([Tolman, 1932], [Redish, 1997], [Hampson, 1990]).

[Krose & Eecen, 1994] show that topological relationships may be generated in a SOM by using richer sensory input. Their approach is to develop a representation in sensor space rather than any notion of real space. The system produces a representation, but the authors admit that the topological information generated is not useable for navigation purposes. A complete system based on SOMs was constructed by [Tani & Fukumara, 1994]. They, too, built a sensor based representation of the environment and used that representation to build a relationship between sensor space and motor action. The relationship between the SOM and motor action was one-to-one, making this another SR model. It might be noted that a similar effect can be generated by the simplest of training algorithms and networks (gradient descent

of linear units) without the need for intermediate representations, as illustrated by [Nehmzow, 1995] and [Wyeth, 1997].

None of these systems form maps that contain the necessary spatial information to generate path plans. This makes them unsuitable for use in a complete animat, and is contrary to the observed biological data. An exception is [Najand *et al.*, 1992] who also use a SOM to represent the environment, but rely on the Cartesian coordinates of the robot rather than sensory input to generate the topology. Here, the representation is useable for path planning but the model of the sensory input lacks credibility from a biological perspective and useability from a robotics perspective.

[Scholkopf & Mallot, 1995] argue that the mapping performed by SOM style networks, which is to cluster similar features together, is not appropriate for developing a cognitive map of the environment. Rather information should be clustered according to spatio-temporal similarity. They argue that the SOM's incorrect emphasis on the nature of the mapping is due to the lack of consideration for the output representation of the map. While the input representation is considered in detail, and with a view to realistic data, there is often no consideration for what kind of output representations are required for the maze recall processes. As a result of this, the maps are strongly linked in a sensory sense, but are weakly linked in a spatio-temporal and an action sense. As robot builders, we agree that the lack of representational power of the map creates difficulties in using the map in a complete system.

3.2.2 Recurrent Networks

[Hetherington & Shapiro, 1993] simulate hippocampal place fields using a network structure based around recurrent structures [Elman, 1990]. The agent is located in a square arena represented by an 8x8 grid. The inputs to the network are a set of five inputs representing the location of the goal cue around the wall (one at each corner and one at the centre), and the angle subtended to each of the 12 cues around the environment (on a 360 degree retina with a constant orientation). The hidden layer of units was trained to remember the trajectories from each possible start position to each possible goal position. The hidden layer of units shows place-field-like properties to various areas within the environment and responds to the constellation of visual stimulus rather than individual stimulus [O'Keefe & Nadel, 1978], and [Redish, 1997]. Finally the recurrent connections allow the persistence of the place representation after the visual stimulus is removed.

The authors present the model to produce some predictions about the nature of place units. These are:

1. place fields should not persist in unfamiliar environments, because the recurrent connections are experience dependent modifications ,
2. place cells with overlapping place fields should have stronger recurrent connections than non-overlapping place cells (on average), and
3. the location of place fields should not be modulated by the goals, rather the presence or absence of place fields should be altered when different goals are made available.

The study shows place fields alone do not give relationships in environment (distances orientations and actions), but connections between place cells (recurrent connections in this case) can store such information and they argue that this can be used to plan paths and actions.

The authors clearly state that their model is, in its present form, unsuitable for mapping and path planning. There is also no consideration of orientation, nor is there kinaesthetic information included in the inputs to the map. This limits the direct applicability of the model to a complete system. It does however illustrate that neural representations can contain at least

a representation of topology and distance that is independent of sensory stimulus or motivational goal.

3.2.3 Association of Places and Views

The models proposed by Schmajuk [Schmajuk & Thieme, 1992] and Scholkopf [Scholkopf & Mallot, 1995] both take the approach that the maze can be broken up into places and views. The cognitive map developed is then one that represents the relationships between the different places and views within the maze and not the distances or orientations between them. In both models the maps are used by higher level cognitive processes, to plan and execute paths to the goal. The models differ in their structure and behaviour, due to the assumptions made about the nature of the relationships between places and views. Schmajuk assumes that there is a one-to-one correspondence between the places and views within the maze. In contrast, Scholkopf argues that in the real world there may be more than one view of the same place (for instance turning on the spot).

Schmajuk's cognitive architecture consists of two main parts, as first suggested by [Tolman, 1932]; a cognitive map and an action system. The map constructed is topological in nature in that it stores adjacencies between places but not distances or orientations. The cognitive map is responsible for both the real-time predictions of the view to be seen from a place, and also the fast-time predictions of where to go to reach the goal. The important point is that the map is not goal driven in nature, as suggested first by Tolman. In various simulations, the architecture is shown to be capable of reproducing latent learning and the detour behaviour observed in real rats.

The key to the operation of this model is the dual nature of the network. High pass filters are placed on recurrent connections in the place-view association network. The network is able to generate fast predictions about the path to the goal using these recurrent connections. The speed of these predictions prevents modification of the place-view associations. The slower real-time predictions are filtered out and do not cause recurrent activity. The feed-forward operation of the network maintains the topological integrity of the map.

The model proposed by [Scholkopf & Mallot, 1995] uses similar principles in that the map relates places to views. Scholkopf and Mallot extend Schmajuk's ideas by recognising that in real environments there is generally more than one view for a given place. A cognitive map is then developed from a directed graph relating the adjacencies of the views within the maze-like environment. In the neural architecture, proximity in the map relates to connectedness of the views in the world, not of similarities between the feature vectors of the views. The distance measure of the map is also different in that it is the minimum number of synapses that must be passed between the two units. The architecture also includes knowledge of "how" to get from one node to the next by using afferent connections from motor representations such as "forward" or "back".

In the simulation of the architecture, the agent's position and orientation in the virtual maze is represented by the ordered pair of the current place and view. The cognitive map is built by allowing an agent to randomly "walk" through the maze environment. In contrast to Schmajuk's model, there is not a specific neural construct to perform maze recall. Instead an algorithm is presented in which a "higher" cognitive system interacts with the map in order to make predictions about where to go. This higher cognitive system actually changes the activations within the network, thus the dual nature of Schmajuk's model is lost.

The main difficulties with the use of these models in a complete animat are the assumptions made about the place-view system. These assumptions include the requirements for uniqueness of views and the one-to-one correspondence of views to places (although Scholkopf attempts to improve the place-view relationships). This issue is especially critical in the low sensory resolution systems found in robots such as CUQEE. With only three

sensors, every T-junction looks identical. Schmajuk's model does not recognise that junctions may appear completely different depending on the angle of approach. The model should produce different place fields for the different views, an argument that is supported by biological evidence [McNaughton *et al.*, 1983]). Lack of attention to realistic sensory input makes the use of these models impractical for a physically embodied system.

These problems aside, it is difficult to see benefits in the cognitive maps presented in place-view models over CUQEE's cognitive map. The models' place cells represent distinct positions in the maze, as does CUQEE, and distance metrics are derived from topological information, as does CUQEE. As such the model of the place cells are at least as limited in applicability and plausibility as CUQEE's.

3.3 Motivational Level

We do not expect to find biological inspiration at a motivational level. CUQEE does not get hungry, thirsty or desire sex. Of more interest are issues in interfacing a symbolically programmed motivational level with the sub-symbolic computation at the cognitive level. These issues are beyond the scope of this paper. The reader may consult [Donnart & Meyer, 1994], [Werner, 1994], and [Toates & Jensen, 1990] for some results in the study of motivation.

3.4 Systems that Span Levels

The robot *Toto* [Mataric, 1990] represents a system that spans the cognitive and schema levels and has some of the features we have identified as important in cognitive maps. The robot is not based on a neural network architecture, but bears similarities to connectionist models in that it uses parallel competing processes and graph based representations of space. The map is a graph of landmarks with connectedness representing closeness. Each node in the graph represents a particular landmark encountered in the environment. In contrast to the Schmajuk and Scholkopf models, there is metric information stored in the map which aids in path planning. Furthermore there is action information stored which gives the relationships between nodes and how to get to the connected nodes. Mataric argues that metric information is required, not only to plan the shortest path, but also to distinguish between similar landmarks that are in different locations. The graph connections are dynamic in nature allowing multiple connections. This means that a particular landmark may be a branching point for more than two paths (consider a 5 way intersection for example).

The path planning system spreads activation from the goal node, throughout the network. Navigating to the goal is then an emergent behaviour, where the robot decides which landmark to go to next based upon which leads to the shortest path to the goal. This is determined by the original spread of activation. Mataric's results show that the robot can work successfully in an office environment with low resolution, noisy sensors which reveal a handful of distinct landmarks.

This complete model, built to operate in an unstructured environment, has many desirable features as a model for navigation. As a robot, the chief limitation is in the exploration strategy which is limited to following boundaries. This is an inefficient maze solver, and will not solve mazes where the goal is not connected to the outside wall. The key argument against biological plausibility is the choice of representation. Nevertheless, a connectionist architecture that displayed many of *Toto's* properties would be a strong candidate for a cognitive model for maze navigation.

Complete models such as *Toto* and CUQEE quickly reveal their limitations as they are forced to cope with real sensors and real action. The difficulties shown in using partial neural solutions in a real animat presents a need for complete connectionist models of spatial

navigation. The interfaces connecting the animat to the real world, the sensory and motor interfaces, are the perhaps the most important in the entire system. Inadequate or inappropriate representation of data at these points determines the usefulness of the architecture.

Other researchers [Yeap, 1990], [Cliff, 1990b], and [Prescott, 1994], have argued for development of complete animat systems. The arguments centre on issues of representation. Incomplete systems can choose arbitrary representations of data at both the inputs and the outputs, yet these representational issues should be considered first and foremost. The development of the architecture becomes a matter of course from that point. While this is understating the difficulties in developing a neural structure that has the necessary representational abilities for spatial navigation, it does highlight that the determination of these representations is much more critical than the development of the architecture. If the representations of data are not realistic, or don't properly represent the data, then the architecture will not work in a complete animat.

3.5 In Summary

The spatial navigation abilities of CUQEE are limited in applicability and plausibility by the internal representation of the environment: the cognitive map. While our review shows that systems that address the schema level abound, existing models of the cognitive level vary greatly depending on their implementation of cognitive maps. All have apparently insurmountable difficulties for the robot builder. SOMs map sensor space and not allocentric space making path planning impossible. Current implementations of recurrent neural structures do not address issues of orientation or kinaesthetic information. Models based on the associations of places to views make unreasonable assumptions about the nature of the world and lack the metric information needed for path planning. In short, the existing implementations of cognitive maps do not offer a way forward for the robot building community.

4 Cognitive Model Wish List

This paper has presented a robot that can learn to run through a maze, in similar fashion to a maze solving rat. In the process of building this robot from an engineering perspective, we have inadvertently created a machine that displays characteristics of rodent behaviour in the maze. It is proposed then that the models used for the development of the robot are applicable to the behaviour of the rodent. In saying this, we would hope to be able to build a better robot by closer studies of biological models of spatial navigation. Our review of the biological models reveals a significant disparity between the characteristics of those models and the needs of the robot builder. It also follows that many of the existing cognitive models of spatial navigation are lacking in plausibility, as they cannot be used to form a complete system.

The study of CUQEE and the review of the literature has revealed some key aspects of cognitive models of spatial navigation. In this section we highlight the aspects that are highly desirable for the purposes of building better robots, and also, we believe, for constructing more plausible models. Certainly, as robot builders, we envisage a plethora of applications for robots that can navigate as reliably and swiftly as CUQEE, but in unstructured environments rather than the confines of the maze. This work cannot proceed at present without a better model for spatial navigation. While recognising that navigation research is progressing with different agendas in mind, we have compiled a cogent wish list for navigation models that suit the robots of the future.

1. Navigation must be considered as a whole.

Models that split navigation into small pieces, and consider only singular aspects of the navigation problem, run the risk of getting the interfaces wrong. Models of spatial navigation should reflect all levels of navigational competence from taxon (reactive) navigation to locale (cognitive) navigation, and make it clear how the systems interface. Similarly, models of navigation should place equal importance on exploration and recall. Building a map that can't be recalled is of no value. While it is possible to recall a map built by others, or supplied by the user, maps that are built internally, or at least updated internally, more accurately represent the environment within the animat's ability to perceive it.

2. Navigation models must interface with a real world

There is little point in building models that operate in toy worlds but fail in real worlds. Toy worlds bear little resemblance to the rich, real worlds in which biological navigation systems evolved. While simplified environments present opportunities for isolating specific behaviours for evaluation, models also require a reality check to see whether they can cope with natural environments. As Wish 1 has already highlighted, this necessarily involves providing a world interface so that the model can be tested beyond a single problem.

3. Maps should be maps

There is little value in a street map that tells you only, "You will know you are in Ann Street because you can see the bank, and you will know you are in George Street because you can see the bakery." . Clearly such a description does not allow the map user to navigate efficiently from the Ann Street to George Street. This is the argument against maps that operate in sensor space, or rely solely on place-view associations. Maps must illustrate topology and spatio-temporal metrics, or provide methods of inference for these properties, to be useful as navigational aids.

4. Use biological components to build biological structures

This paper has shown that the representation of the cognitive map is a key issue in the determining the functionality of a navigating agent. If the aim of the research is to mimic a biological process, it is clear that the representation should be biologically sound. Given the importance that Wishes 1 and 2 place on treating the process as a whole, and on providing realistic contact with the world, it follows that biological components should be used throughout the cognitive structure. Biological components refer to acceptable models of neuron function.

This wish list presents a challenge for spatial navigation researchers as we develop the next generation of biologically plausible robots.

5 Acknowledgements

We gratefully acknowledge the assistance of Janet Wiles for her comments on the early drafts of this paper, and her guiding thoughts in its construction. We thank Mark Schulz for his comments on the later drafts and his efforts in proof reading.

6 References

Arkin, R.C. (1989). Motor schema-based mobile robot navigation, *The International Journal of Robotics Research*, vol. 8, no. 4, pp. 92-112.

- Beer, R.D, Chiel, H.J. and Sterling, L.S. (1990). A Biological Perspective on Autonomous Agent Design. *Robotics and Autonomous Systems*, v. 6, pp 169-186.
- Braitenberg, V. (1984). *Vehicles: Experiments in Synthetic Psychology*, MIT Press, Cambridge, MA.
- Brooks, R.A. (1990). Elephants Don't Play Chess. *Robotics and Autonomous Systems*, vol. 6, pp. 3-15.
- Cliff, D. (1990a). The Computational Hoverfly. *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behaviour*, The MIT Press: pp. 87-96.
- Cliff, D. (1990b). Computational Neuroethology: A Provisional Manifesto. *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behaviour*, The MIT Press: pp. 29-39.
- Connell, J. (1990). *Minimalist Mobile Robotics*. San Diego: Academic Press, Inc.
- Donnart, J., & Meyer, J. (1994). A hierarchical classifier system implementing a motivationally autonomous animat. *From Animals to Animats: Proceedings of the Third International Conference on Simulation of Adaptive Behaviour*, The MIT Press: pp. 144-153.
- Eichenbaum, H. & Cohen, N. (1988). Representation in the hippocampus: What do hippocampal neurons encode? *Trends in Neuroscience*, 11, pp. 244-248.
- Eichenbaum, H., Wiener, S., Shapiro, M., & Cohen, N. (1989). The organization of spatial coding in the hippocampus: a study of neural ensemble activity. *Journal of Neuroscience*, 9, pp. 2764-2775.
- Elman, J. (1990). Finding structure in time. *Cognitive Science*, pp. 179-212.
- Gallistel, C.R. (1989). Animal Cognition: The Representation of Space, Time and Number. *Annual Review of Psychology 1989*, Vol 40, Palo Alto, California, pp. 155-189.
- Gallistel, C. R. (1990). *The organization of learning*. Cambridge, MA: MIT Press.
- Gat, E. (1994) Behaviour Control for Robotic Exploration of Planetary Surfaces, *IEEE Transactions on Robotics and Automation*, August 1994.
- Hampson, S. (1990). *Connectionistic Problem Solving*. Birkhauser, Boston.
- Hetherington, P. & Shapiro, M. (1993). A simple network model simulates hippocampal place fields: II. Computing goal-directed trajectories and memory fields. *Behavioral Neuroscience*, 107, No. 3: pp. 434-443.
- Kandel, E.R. (1976). *Cellular Basis of Behavior*, W.H. Freeman and Company.

- Kohonen, T. (1995). *Self-organizing maps*. New York: Springer.
- Krose, B., & Eecen, M. (1994). A self-organizing representation of sensor space for mobile robot navigation. *Proceedings of the IEEE/RSJ/GI International Conference on Intelligent Robots and Systems IROS'94*. Munchen: pp. 9-14.
- Mataric, M. (1990). *A Distributed Model for Mobile Robot Environment-Learning and Navigation*. Masters Thesis. MIT Artificial Intelligence Laboratory, MIT.
- McNaughton, B. (1989). Neuronal mechanisms for spatial computation and information storage. In L. Nadel, L. Cooper, P. Culicover, & R. Harnish (Eds.), *Neural connections, mental computation*. London: MIT Press.
- McNaughton, B., Barnes, C., & O'Keefe, J. (1983). The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely moving rats. *Experimental Brain Research*, 52, pp. 41-49.
- McNaughton, B., & Nadel, L. (1990). Hebb-Marr networks and the neurobiological representations of action in space. In M. A. Gluck & D. E. Rumhart (Eds.), *Neuroscience and Connectionist Theory*. Hillsdale, NJ. Lawrence Erlbaum Associates, pp. 1-63.
- Muller, R. & Kubie, J. (1989). The firing of hippocampal place cells predicts the future position of freely moving rats. *Journal of Neuroscience*, 9, pp. 4101-4110.
- Najand, S., Lo, Z., & Bavarian, B. (1992). Application of self-organizing neural networks for mobile robot environment learning. *Neural Networks in Robotics*. Ed. Bekey, G., Goldberg, K. Kluwer Academic Publishers, Dordrecht. pp. 85-96.
- Nehmzow, U. & Smithers, T. (1990). Mapbuilding using self-organising networks in "Really Useful Robots". *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behaviour*, The MIT Press, pp. 152-159.
- Nehmzow, U. (1995). Flexible control of mobile robots through autonomous competence acquisition. *Measurement and Control*, vol. 28, pp. 48-54.
- O'Keefe, J. & Nadel, L. (1978). *The Hippocampus as a Cognitive Map*. Oxford, Clarendon Press.
- O'Keefe, J. & Speakman, A. (1987). Single unit activity in the rat hippocampus during a spatial memory task. *Experimental Brain Research*, 68, pp. 1-27.
- Otten, D. (1990). Building MITEE Mouse III. Part 2. *Circuit Cellar Ink*. pp. 40-51.
- Payton, D.W. (1986) An Architecture for Reflexive Autonomous Vehicle Control, *IEEE International Conference on Robotics and Automation*, pp.1838-1845.

- Prescott, T. J. (1994). Spatial learning and representation in animats. *From Animals to Animats: Proceedings of the Third International Conference on Simulation of Adaptive Behaviour*, The MIT Press: pp. 164-173.
- Redish, A. D. (1997). *Beyond the Cognitive Map: Contributions to a Computational Neuroscience Theory of Rodent Navigation*. Phd Thesis. Computer Science Department & Center for the Neural Basis of Cognition, Carnegie Mellon University. Pittsburgh.
- Roitblat, H. L. (1994). Mechanisms and process in animal behavior: models of animals, animals as models. *From Animals to Animats: Proceedings of the Third International Conference on Simulation of Adaptive Behaviour*, The MIT Press: pp. 12-21.
- Schmajuk, N. A., & Thieme, A. D. (1992). Purposive behavior and cognitive mapping: A neural network model. *Biological Cybernetics*, 67, pp. 165-174.
- Scholkopf, B. & Mallot, H. (1995). View-based cognitive mapping and path planning. *Adaptive Behaviour*, 3, No. 3:, pp. 311-348.
- Smart, W. & Hallam, J. (1994). Location recognition in rats and robots. *From Animals to Animats: Proceedings of the Third International Conference on Simulation of Adaptive Behaviour*, The MIT Press: pp. 174-178.
- Tani, J. & Fukumura, N. (1994). Learning goal-directed sensory-based navigation of a mobile robot. *Neural Networks*, 7, No 3, pp. 553-563.
- Toates, F., & Jensen, P. (1990). Ethological and psychological models of motivation-towards a syntehsis. *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behaviour*, The MIT Press: pp. 194-205.
- Tolman, E. C. (1932). *Purposive Behavior in animals and men*. New York: The Century Co.
- Webb, B. & Hallam, J. (1996) How to Attract Females: Further Robotic Experiments in Cricket Phonotaxis. *From Animals to Animats: Proceedings of the Fourth International Conference on Simulation of Adaptive Behaviour*, The MIT Press: pp. 75-83.
- Werner, G. (1994). Using second order neural connections for motivation of behavioral choices. *From Animals to Animats: Proceedings of the Third International Conference on Simulation of Adaptive Behaviour*, The MIT Press: pp. 154-161.
- Wyeth G.F. (1997) Neural Mechanisms for Training Autonomous Robots. *Proceedings of Mechatronics and Machine Vision in Practice Conference*, publ. IEEE Computer Society Press, in press.
- Yeap, W. K. & Handley, C. C. (1990). Four important issues in cognitive mapping. *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behaviour*, The MIT Press: pp. 176-183.