

# Navigating With a Rat Brain: A Neurobiologically-Inspired Model for Robot Spatial Representation

Maja J Mataric  
MIT Artificial Intelligence Laboratory  
545 Technology Square, #721  
Cambridge, MA 02139, USA  
maja@ai.mit.edu

## Abstract

This paper presents a neurobiologically-feasible spatial representation model. The model was implemented and tested on a physical autonomous mobile robot. It was shown to be both computationally simple and physically robust.

The described model is a possible interpretation of the organization and function of the rat hippocampus. The paper presents relevant biological, psychological, and neurobiological data, and gives a detailed set of comparisons between the physical hippocampus and our "synthetic" rat implementation. The implications of the many similarities are described. Finally, areas for future study in both biology and robotics are suggested.

## 1 Introduction

Most animals, including humans, spend much of their waking time in transit from one place to another [Waterman 89]. Purposefully moving about requires a system for spatial modeling integrated with the mechanisms for handling navigation, locomotion, and motivation. These systems have evolved to perform with impressive robustness. Understanding their function has long been a goal of cognitive scientists, biologists, and neuroscientists. More recently, this goal has been adopted by members of the Artificial Intelligence and robotics communities interested both in simulating biological systems and designing better artificial ones.

The question asked by both communities is: "What kind of spatial information is stored?" In order to answer it, experiments are designed to test where on the qualitative-to-quantitative scale the representation lies, and whether it is centralized or distributed. This paper describes a qualitative, distributed spatial representation tested empirically on a mobile robot.

## 2 Cognitive Maps

A *cognitive map* is a generic term for an internal representation of spatial information. The term has come to connote a very analytical, centralized representation. In this paper, we will use the term in its generic meaning, and analyze its variants.

A cognitive map is usually assumed to represent space with a set of landmarks, each of which is an element (object or feature) serving as a point of reference [Presson and Montello 88]. According to Piaget, a landmark is a spatial primitive, and thus a basic building block of spatial representations [Piaget and Inhelder 67]. Although most landmark studies concentrate on visual cues, the concept generalizes to any perceptible feature. Animals construct landmarks from auditory, olfactory, and tactile cues as well [Gould 82], taking advantage of their different characteristics [O'Keefe 89].

In this paper, we will analyze cognitive maps along two dimensions: 1) what information they encode and 2) how they encode it. The "what" dimension can vary from completely qualitative or topological to very quantitative or metric. The "how" dimension varies from totally global or centralized to entirely distributed or decentralized.

### 2.1 How Qualitative?

The nature of the representation determines the type and number of landmarks required for localizing. In a qualitative representation, an object can be remembered as being proximate to a landmark, defined within a radius around it. On the other end of the spectrum, the position of an object can be computed precisely from the known locations of three landmarks [Pick, Montello and Somerville 88]. The question is how much metric information is recorded.

The psychological literature is divided on this issue. Studies testing response times in object position recall

indicate that adults are certainly capable of reconstructing Euclidian distance and constructing metric maps. While sufficient information is gathered and recorded to make metric inference possible, the computation appears to be relatively slow and performed only when needed. Experiments have shown that the produced metric inferences tend to be distorted based on local feature density and importance.

[Piaget and Inhelder 67] proposed, and later research supported, an almost exclusively topological theory of early spatial knowledge. The knowledge of children in the preschool period is believed to be fundamentally non-metric, and rely on the presence of landmarks. More specifically, the landmarks must be close, so that children can form relative distance relations among them. This introduces a paradox: while children do not seem to record metric information, they must be using it to form relations between landmarks. These studies are especially interesting in testing the limits of qualitative representations.

### 2.1.1 Experiments With Rats

Acquiring spatial information is one of the most highly developed cognitive abilities in rats [McNaughton and Nadel 90]. The following experiments were aimed at testing the types of information encoded in the rats' cognitive maps.

Maze-running experiments were used to test the animals' ability to construct internal maps [Gallistel 89]. After allowing the rats to familiarize themselves with a maze, the lengths of some corridors were altered. On subsequent trials, the rats ignored their sensory input and ran into the walls of the shortened corridors, and stopped before the ends of the lengthened ones, at locations corresponding to their previous length.

A more challenging experiment allowed the rats to learn a maze and then blocked the known path to the goal. The rats were able to find alternate routes and even generate short-cuts to the goal. They succeeded in doing so in the dark as well, presumably using an entirely different set of non-visual cues.

A well known series of experiments employed a radial maze. After allowing the rats to explore some of the arms, the maze was rotated. Subsequently, the rats entered already sampled arms without realizing the redundancy of their actions. This evidence points toward a metric representation of space, utilizing absolute angles and distances, as opposed to a topological representation utilizing adjacency relations between landmarks, in this case the arms of the maze. Instead of selecting an arm to explore based on whether its neighbor was already explored, the rats apparently made the selection based on the arm's absolute position. In order to find that position and establish global references, the rats used

landmarks external to the maze, such as the position of overhead lights.

Another experiment showed that even when given perceptible sensory features, rats seem to prefer to find the goal using the cognitive maps they have developed. The experiment involved placing a rat into a circular vat of opaque liquid and allowing it to locate a submerged platform. In subsequent trials, even if the platform was moved to a new but visible location, the rat preferred to locate it based on previously established landmarks [Morris 81]. The rat was able to find the platform even if local cues were eliminated by moving the vat to a different room. This indicates the use of both local and global information about the environment.

The rat's spatial representation has proven to be very well adapted to its navigational tasks. Elaborate sensory systems are designed for functional redundancy. Special-purpose computational hardware performs various routine tasks, such as position and angular displacement from a reference point. Both topological and metric information appears to be represented and accessible.

The task of the navigation system, biological or artificial, determines the level of qualitiveness of the representation. In this paper we present a topological model containing only the metric information necessary for the task.

## 2.2 How Distributed?

The "hardware" of biological systems is fully distributed, but it is not yet known how the "software" maps onto this structure. In the case of cognitive maps, we are interested in the level on which the spatial representation is distributed over the neurons computing the relevant information.

Neurophysiological data gathered from rats indicate a distributed spatial representation on the level of landmarks in the hippocampus. Hippocampal neurons are believed to encode non-egocentric cognitive maps. The firing of these neurons, or *place cells*, is associated with the animal's absolute location in the environment. A *place field* is the part of the environment represented by maximal firing of a neuron. Rather than having specific neurons or network patches correspond to particular locations, the entire hippocampal network participates in representing various aspects of the mapped region [Eichenbaum and Cohen 88]. Place fields appear to be distributed evenly around the mapped environment. Consequently, a single neuron often has multiple place fields and codes for multiple locations in the environment [O'Keefe 89]. [Eichenbaum and Cohen 88] suggest that such a fully distributed, non-topological representation lends the system additional generality, allowing it to possibly store semantic as well as spatial relationships.

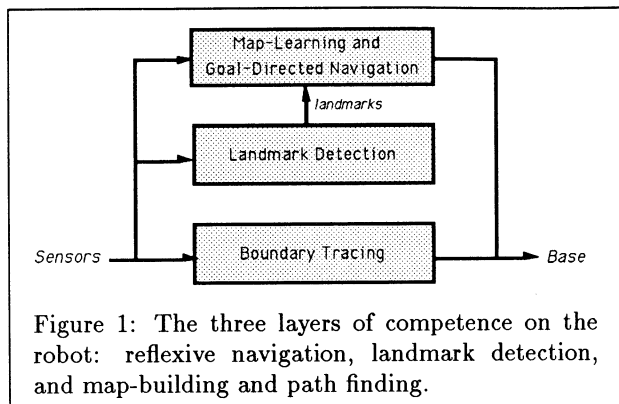


Figure 1: The three layers of competence on the robot: reflexive navigation, landmark detection, and map-building and path finding.

The biological data is hotly debated, but the benefits of such a decentralized representation are worth exploring. In contrast, in robotics, the prevalence of single-processor von Neumann computers has so far overwhelmingly mandated the adoption of centralized models for spatial representation. We introduce a model for a very different, distributed representation, which was implemented and tested on a fully autonomous mobile robot.

### 3 A Robot With an Artificial Rat Hippocampus

The robot is equipped with sonar sensors and a compass. The sonars provide range information in twelve directions around the robot, and the compass provides roughly sixteen heading directions. The robot is programmed with a parallel, distributed control system based on the subsumption architecture. The subsumption architecture is an approach to robot control which proposes simple, tight reflex-like rules coupling the robot's sensory inputs and effector outputs [Brooks 86]. This organization allows the robot to act in real-time without having to perform time consuming processing. An important feature of this approach is that the reflexes are active in parallel, providing distributed, decentralized control of the robot. The robot's behavior is made more complex by incrementally adding more rules, combinations of which form *layers of competence*. Each layer can take advantage of those below.

Our robot consists of three layers of competence (figure 1). The bottom layer combines a set of simple reflexes which, when acting in parallel, result in following boundaries of objects while avoiding static and dynamic obstacles. This is the basic, reflexive and exploratory behavior of the robot.

The middle layer implements a proprioceptive system which detects landmarks in the environment by monitoring the way the robot is moving. The landmarks are




Figure 2: An example of the robot's reflexive navigation behavior in a cluttered office environment. The landmark labels include the landmark type and compass bearing (LW8 = left wall heading South, C0 = corridor heading North, J = long irregular boundary).




Figure 3: The map the robot constructed in the shown environment. Topological links between landmarks indicate their physical spatial adjacency.

defined as combinations of the robot's motion and its sensory input. For example, a corridor is a combination of moving straight and receiving short lateral distance readings. Figure 2 illustrates the robot's typical behavior in a cluttered office environment. Landmark labels consist of the type (LW = left wall, C = corridor, J = long irregular boundary, etc.) and the associated compass direction.

As they are discovered, the landmarks are sent to the third layer of control, which constructs a map of the environment while the robot is moving. Figure 3 shows the map the robot constructed for the shown environment.

The nature of the map representation and the path finding algorithm implemented on the robot will be compared to some interpretations of biological data gathered from analogous processes performed in the rat hippocampus.

### 3.1 The Distributed Spatial Representation

In our implementation, the distributedness of the representation lies at the level of landmarks. We use a network of landmarks as the basic representation structure. Physically adjacent landmarks in the environment are represented as immediate neighbors in the network.

Making a biological analogy, each node in the network corresponds to a place field of neurons in the hippocampus. Each node is implemented as a behavior, or a collection of real-time rules, corresponding to a unique landmark [Mataric and Brooks 90]. All network nodes are processing information in parallel. They receive inputs from the landmark detector, as well as from the sensors on the robot. Additionally, each landmark can communicate with its neighbors by *spreading activation*.

Our approach models a place field as a single landmark in the map. A more neurobiological model would distribute each landmark over a field of neurons. It would be interesting to explore the potential advantages of such an approach for both biological and artificial systems. In the hippocampus, multiple neurons may be necessary to represent a single landmark, but how a landmark is mapped over a place field is not known. The distributedness of the representation may utilize redundancy for increased fault tolerance.

The robot's location in the environment is indicated by an *activated* network node, analogous to a hippocampal place cell with high potential. In order to assure uniqueness of activation, the active node performs lateral inhibition. Additionally, the activated node spreads *expectation* to its neighbor in the direction of motion. Consequently, the next node to be activated is primed, or pre-activated. This is based on experiments showing that proximal landmarks are experienced together in the hippocampus in that they may share activation and are recalled simultaneously. In our implementation, expectation is a method of preserving minimal context which is useful for disambiguating similar landmarks. It also bears similarity to one of the functions of the head-direction neurons which respond to particular orientations of the rat's head [McNaughton and Nadel 90].

Neurobiological data indicates that, in rats, the environment is represented by the changes in the activation levels of the neurons associated with motion, rather than their absolute, instantaneous firing rates [O'Keefe 89]. Our implementation simulates the activation levels by propagating different types of messages in the landmark network. The complexity of the analog signals is only roughly approximated by the binary signals for lateral inhibition and expectation. The messages spreading motivation to pursue a particular goal transmit eight bits

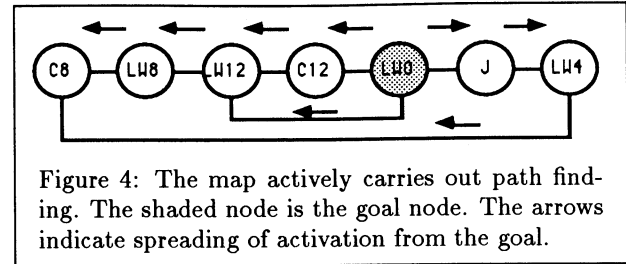


Figure 4: The map actively carries out path finding. The shaded node is the goal node. The arrows indicate spreading of activation from the goal.

of distance information.

### 3.2 Building a Map and Finding Paths

The process of building a cognitive map consists of the robot wandering through the environment and storing the encountered landmarks in a distributed network. This is analogous to spontaneous growth of the hippocampal potential resulting from exposure of animals to new environments [Barnes 88].

The landmark map is used for localizing, i.e. determining the robot's position. Once localized, the robot can find and execute paths to desired goal locations. The robot's goals are currently selected by the user, but could be supplied by other reasoning modules associated with fulfilling tasks or satisfying the robot's needs. An analogous process is found in biology. According to [O'Keefe and Nadel 78], goal locations in the hippocampus are selected by an external mechanism which acts in response to the animal's needs (hunger, thirst, etc.) and desires (shelter, mate, etc.).

The activity of the hippocampal place cells need not be stimulated by explicit sensory inputs. Cognitive activity, such as visualizing or "thinking about" the goal location, is sufficient [Eichenbaum and Cohen 88]. This process is modeled by path finding. Given a goal and a current location, path finding is implemented by *spreading of activation* from the goal landmark throughout the map as illustrate in figure 4. The activation is eventually propagated to all locations in the map. The direction from which activation arrives into a network node specifies locally the direction in which the robot should travel next from that landmark. This eliminates the need for replanning if the robot strays from the desired path since regardless of where the robot is located, it knows the optimal direction to pursue toward the goal [Mataric 90b].

The process of activation spreading is equivalent to performing a network search. Each landmark in the map stores an estimate of its physical length. As the activation wave is propagated through the network, it accumulates the length values of the landmarks it passes. Consequently, at each landmark, the shortest path towards the goal corresponds to the direction with the smallest accumulated length. The robot takes the opti-

mal direction locally, at each landmark, which produces the globally shortest path to the goal.

Very similar, purely local approaches to navigation are suggested by [McNaughton 89]. He identifies three possibilities. The first, most similar to our implementation, utilizes a transition matrix linking local views with movements connecting them. Another uses spatially invariant landmarks and forms permanent associations between them and the goal. This is equivalent to the result of activation spreading in our system, but has a fixed goal location. The third approach proposes a built-in method for computing motion sequences explicitly.

## 4 Using Motion as a Representation Tool

In [McNaughton and Nadel 90], McNaughton describes experiments in which a rat's legs were bound to prevent motion. Such restraint prevented all activation change in the hippocampus as the rat was carried through both a known and a novel environment. When motion of the legs was allowed, neuronal activity resumed, even if the rat was not free to walk about. These experiments stress the importance of integration between the locomotion mechanism and the cognitive map, suggesting that motion is crucial in generating a map.

Similarly, the connections in our network representation do not carry any information. All of the data is contained in the landmarks, and in the robot's reflexive navigation behavior. Consequently, the landmarks are connected by motion directives. At each location, the robot knows where to move in order to reach the next landmark. This may model exactly the mechanism used in the hippocampus. [McNaughton 89] presents data suggesting that relationships between landmarks in the hippocampus are most likely encoded in terms of movement. [Foster, Castro and McNaughton 89] suggests that hippocampal activity reflects the association of movements with their spatial consequences. The relationship between two places is represented by the movements executed in getting from one to the other. As in our model, the rat knows what movement will take it from one landmark to another.

## 5 Evaluation

Computationally, our method is attractive for a number of reasons. The parallel implementation allows for very fast (constant time,  $O(1)$ ) localization within the map, regardless of its size. Shortest path finding to a goal location consists of a parallel network search from the goal to the activated node. The robot makes a greedy,

locally optimal choice at each landmark which results in a globally optimal path to the goal. The algorithm runs in worst case linear time ( $O(n)$ ) in the size of the map [Mataric 90a].

Aside from using landmark length estimates to augment path finding, the method we present relies almost entirely on topological information. It could be made more general and biologically feasible by taking advantage of the available metric information for making geometric inferences. For example, discovering physical proximity between two landmarks can generate new neighborhood connections in the network. This would allow the robot to create shortcuts through reasoning only, without having to physically traverse the paths. Rats are known to possess this ability [McNaughton 89].

Most robotics planning systems use some kind of a central control module which consults the map to make decisions about the robot's motion. In contrast, our approach does not use such a separate, centralized reasoning engine outside the network. All spatial reasoning necessary for building and using the map is in the map itself. This is also biologically rooted: the hippocampus performs all of the necessary computation independently, using its specialized hardware. In general, the notion of many specialized modules acting in parallel seems to be more biologically justified than the classical, sequential planning approach.

The controversy between global versus distributed spatial coding in the hippocampus remains to be resolved through further neurobiological study. The approach presented here, while distributed according to the standards of the robotics community, is nonetheless global in the biological sense. The structure of the distributed landmark network is isomorphic to the topology of the environment, and is thus a centralized structure, which can be accessed sequentially. Whether such an isomorphism exists in the hippocampal coding is not yet known.

## 6 Conclusions and Future Directions

The interpretation of experimental data usually suffers from the bias of the interpreter. Rat hippocampus studies remain controversial as the field is divided on the fundamental issues concerning the nature of the spatial representation involved.

In this paper we have taken a particular set of interpretations of the function of the rat hippocampus and compared it to our physical implementation of a parallel distributed algorithm for robot spatial mapping and path finding. While they do not prove the correctness of the interpretation, our results serve as an empirical

proof of its feasibility and efficiency. A number of important questions remain:

How is motion information used to encode the relationships between different landmarks or place fields? How are they linked together into sequences generating shortest paths? The answers to those questions might resolve the distributed versus global map dilemma.

What are the benefits of distributing landmarks over multiple, overlapping place fields? How are multiple maps overlaid onto the same area? In the hippocampus, this mechanism allows for encoding many environments in finite space. Any hints toward the solution of this problem would certainly be useful in representation research in Artificial Intelligence.

Which tasks and architectures are well suited for the use of distributed spatial representations? When are they better suited than global maps? These questions can be answered by building more distributed systems, and evaluating their performance.

Rather than propose a theory of either rat or robot navigation, we presented a working alternative to the standard planning approach usually employed in Artificial Intelligence. Additionally, we suggested the usefulness of biological models in designing artificial systems, as well as the utility of computational models in analyzing biological ones.

Our approach is an example of a specialized subsystem performing self-contained computation in parallel. Biological systems demonstrate that collections or societies of such computational agents are worth exploring in order to address global issues in Artificial Intelligence [Minsky 86] and robot control.

## 7 Acknowledgements

Geoffrey Hinton originally pointed out the work on the rat hippocampus and its relevance to the research presented here. Rodney Brooks suggested the idea of a distributed representation and provided comments on an earlier draft of this paper. Bruce McNaughton provided helpful information about current hippocampus studies.

The research reported here was done at the MIT Artificial Intelligence Laboratory. Support for this research was provided in part by the Hughes Research Laboratories Artificial Intelligence Center in Malibu, in part by the University Research Initiative under Office of Naval Research contract N00014-86-K-0685, and in part by the Advanced Research Projects Agency under Office of Naval Research contract N00014-85-K-0124.

## 8 References

- [Barnes 88] "Spatial Learning and Memory Processes: The Search for Their Neurobiological Mechanisms In the

Rat", C. A. Barnes, *TINS*, Vol. 11, No. 4, 1988.

[Brooks 86] "A Robust Layered Control System for a Mobile Robot", Rodney A. Brooks, *IEEE Journal of Robotics and Automation*, RA-2, April 1986, 14-23.

[Eichenbaum and Cohen 88] "Representation in the Hippocampus: What Do Hippocampal Neurons Code?", Howard Eichenbaum and Neal J. Cohen, *TINS*, Vol. 11, No. 6, 1988.

[Foster, Castro and McNaughton 89] "Spatial Selectivity of Rat Hippocampal Neurons: Dependence on Preparedness for Movement", Tom C. Foster, Carl A. Castro, Bruce L. McNaughton, *Science*, June 30, 1989, 1589-1582.

[Gallistel 89] "Animal Cognition: The Representation of Space, Time and Number", C. R. Gallistel, *Annual Review of Psychology 1989*, Vol 40, Palo Alto, California, 155-189.

[Gould 82] "Ethology: The Mechanisms and Evolution of Behavior", James L. Gould, *W. W. Norton and Company*, New York, 1982.

[Mataric 90a] "A Distributed Model for Mobile Robot Environment-Learning and Navigation", Maja J Mataric, *MIT AI Lab Tech Report #1228*, May 1990.

[Mataric 90b] "Environment Learning Using a Distributed Representation", Maja J Mataric, *Proceedings of 1990 IEEE International Conference on Robotics and Automation*, May 1990, 402-406.

[Mataric and Brooks 90] "Learning a Distributed Map Representation Based on navigation Behaviors", Maja J Mataric and Rodney A. Brooks, *Proceedings of 1990 USA-Japan Symposium on Flexible Automation*, Kyoto, July 1990.

[McNaughton 89] "Neuronal Mechanisms for Spatial Computation and Information Storage", Bruce L. McNaughton, *Neural Connections, Mental Computation*, Lynn Nadel, Lynn A. Cooper, Peter Culicover, R. Michael Harnish, eds., MIT Press, Cambridge, Massachusetts, 1989.

[McNaughton and Nadel 90] "Hebb-Marr Networks and the Neurobiological Representation of Action in Space", Neuroscience and Connectionist Theory, Mark A. Gluck and David E. Rumelhart, eds., *Lawrence Erlbaum Associates, Publishers*, Hillsdale, New Jersey, 1990.

[Minsky 86] "The Society of Mind", Marvin L. Minsky, *Simon and Schuster*, New York, 1986.

[Morris 81] "Spatial Localization Does Not Require the Presence of Local Cues", R. G. M. Morris, *Learning*

and *Motivation*, Vol. 12, 239-261.

[O'Keefe 89] "Computations the Hippocampus Might Perform", John O'Keefe, *Neural Connections, Mental Computation*, Lynn Nadel, Lynn A. Cooper, Peter Culicover, R. Michael Harnish, eds., MIT Press, Cambridge, Massachusetts, 1989.

[O'Keefe and Nadel 78] "The Hippocampus as a Cognitive Map", John O'Keefe and Lynn Nadel, *Clarendon Press*, Oxford, 1978.

[Piaget and Inhelder 67] "The Child's Conception of Space", J. Piaget and B. Inhelder, *Norton*, New York.

[Pick, Montello and Somerville 88] "Landmarks and the Coordination and Integration of Spatial Information", Herbert L. Pick, Jr., Daniel R. Montello, and Susan C. Somerville, *British Journal of Developmental Psychology*, 1988:6, The British Psychological Society, 372-375.

[Presson and Montello 88] "Points of Reference in Spatial Cognition: Stalking the Elusive Landmark", Clark C. Presson and Daniel R. Montello, *British Journal of Developmental Psychology*, 1988:6, The British Psychological Society, 378-381.

[Waterman 89] "Animal Navigation", Talbot H. Waterman, *Scientific American Library*, New York, 1989.