Robots as Tools for Modelling Navigation Skills - A Neural Cognitive Map Approach

Verena V. Hafner

Humboldt-Universität zu Berlin, Institut für Informatik, Berlin, Germany
`hafner@informatik.hu-berlin.de`

Summary. This chapter attempts to show how cognitive map models can be combined with robotic navigation strategies. A neural cognitive mapping strategy that is inspired by place cells but still abstract enough to be interpreted in a meaningful way is implemented in different experiments with both mobile robot and simulation experiments.

1 Introduction

The ability to navigate in a complex environment is one of the most challenging skills for every animal on this planet. It is crucial for survival, and is even seen as the evolutionary pressure to develop brains. The reason why plants do not have a brain is that they do not have to move [Wol02].

There are different amounts of cognition involved in navigation when considering different species and different goals. This starts with simple aiming and obstacle avoidance already present in one-cell organisms, over route-following for example in insects, to high-level survey navigation which includes both topological and metric information. A useful categorisation of these different skill levels of navigation is presented by Franz et al. in this volume.

The approaches to navigation involving cognitive mapping that this book focuses on, are mainly found at the levels of topological and survey navigation. These approaches require some kind of internal representation of the space the animal is navigating in. However, complex navigation, such as finding back home from various locations on paths never travelled before, is also possible without a full representation of the environment, for example in insect visual homing. Here, the animal can reach the goal position from its current position by comparing visual snapshots in memory without knowing about the location of the goal position nor the current position [Haf04].

There is also evidence from the neurosciences for an internal spatial map: In these experiments (see for example [O'K76], cells have been found in the rat's brain

that represent a certain place within the environment, and may therefore be part of a so-called cognitive map for the animal. These cells are called 'place cells' and are located mainly in the hippocampus, an area in the brain which is also responsible for memory [RH96]. Such cells have been found in rats, mice, and even primates including humans. However in primates, additional cells, called 'view cells', that are related both to a particular place and view have been found.

Biologists [WR79] and neuroscientists [O'K76] have been studying navigation behaviour in animals for several decades, coming up with different hypotheses of how the navigation skills are acquired and implemented in the animal's brain and body. With the advent of behavioural [Bro91] and biomimetic [Web01] robotics, a new field of research got interested in navigation behaviour. These subfields of robotics are inspired by biology, in particular using behavioural experiments. In these experiments, the interaction between the agent and the environment plays a major role. The morphology of the agent with the arrangement of different sensors along its body is also very important for the interaction [PB06]. Traditional or industrial robotics in contrast is only interested in fulfilling a task in a controlled and predictable way using methods such as planning. Along with biorobotics came several advantages which promised to strengthen or falsify the hypotheses of biologists by repeatable experiments using mobile robots. A good review on bio-inspired robot navigation can be found in [TWBM97]. Recently, the robotics community is focussing more and more on hybrid approaches [AC04, TS05] to find a good balance between engineering and biological plausibility.

In this chapter, a neural cognitive map model inspired by place cells, implemented on a mobile robot is presented. In chapter 2, the neuroscientific foundations of place cells are explained and some particular properties are discussed. Chapter 3 presents the neural cognitive map algorithm and structure, and different experiments on a mobile robot and in simulation are shown and their results discussed.

2 Neuroscientific Foundations: Place Cells

Place cells were first discovered in the hippocampus of rats [O'K76]. They are cells whose firing activity depends on the spatial position of the animal in its environment. This implies some sort of internal representation of the outside environment in the brain. There is similar evidence of cognitive maps in humans. We present some of the experiments measuring cells in the human hippocampus in the next subsection. Spatial information is an important part of long-term memory, therefore it is also interesting that the place cell activity seems to be transferred into long term memory during sleep, which is presented in the subsection after that.

2.1 Place Cells in the Human Hippocampus

Place cells have also been found in the hippocampi of humans recently [EKC$^+$03, ECH$^+$05]. Cells in the hippocampal and parahippocampal region of patients with epilepsy were recorded with implanted clinical depth electrodes. During the recording, the subjects were exploring and navigating a virtual town in a taxi driver computer game, searching for passengers and delivering them to fixed target locations.

Out of 317 recorded neurons, 26% responded to a place, 12% responded to a view, and 21% to a goal. Eleven per cent of the cells were true place cells, which only responded to a place. From the 67 neurons measured in the hippocampus, 24% were true place cells. These findings can be seen as evidence for a neural code of human spatial navigation, a cognitive map.

2.2 Place Cells during Sleep

Hippocampal place cells have also been recorded during rat's rapid eye movement (REM) sleep[1]. Wilson and McNaughton [WM94] discovered that cells that fired together when the rat occupied a particular location exhibited an increased tendency to fire together during subsequent sleep, in comparison to sleep episodes preceding the behavioural tasks. In these experiments, 50 to 100 single cells in area CA1 were recorded. They suggest that the neural states encoded within the hippocampus are "played back" as part of a consolidation process by which information is gradually transferred to the neocortex.

# 3 Robotic Experiments: Cognitive Maps

This section describes experiments performed with navigating artificial agents. The experiments are inspired by the findings of place cells in rats, and aim to produce a cognitive map of an environment during exploration. The question is, whether a neural representation can be found that explains the findings and properties of place cells for navigation in rats.

3.1 Experimental Setup

The experimental setup consists of an artificial agent (mobile robot or simulated agent) that performs random exploration tours within a newly encountered open environment with objects functioning as obstacles. The agent has omnidirectional sensory stimulation, either panoramic vision or distance sensors. A benefit of omni-directional vision is that it approximates more closely the rat's field of view spanning $320°$, making the sensory input more comparable between rat and robot. The agent gets proprioceptive feedback about its heading direction, but does not have access to any exact metric distance information, nor does it know its position within the environment. The learning of a cognitive map is purely based on neural plasticity (changing weights) within the agent's brain, using a variant of a self-organising map (SOM). The number of recruitable place cells, i.e existing cells that could function as place cells, is fixed at the beginning of the experiment. During exploration, the agent gets a new visual input every few time steps. The exploration of the environment is open-ended without being task-specific. There is no reward during learning nor is there a goal state. This setup is comparable to rats exploring a newly encountered environment before eating even if they are hungry[Pre96]. The attributed motivation for this behaviour is curiosity.

[1] REM sleep: periods of mental activity during sleep
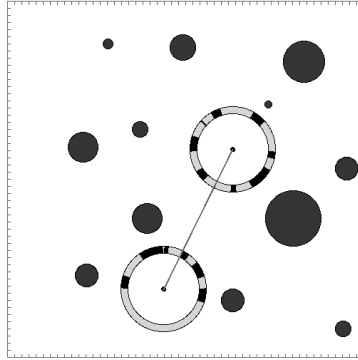
Fig. 1. Cylinder environment and visual information of the agent. In this figure, the omnidirectional binary view of an agent at two different positions is shown (connected by the line). Filled circles are obstacles of different sizes.
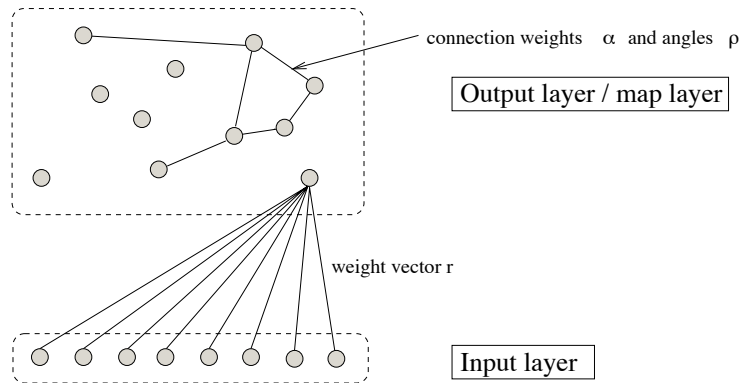
Fig. 2. Neural Network structure for learning of place cells and their connections. The cells in the output layer are place cells. The input layer consists of the sensory input from the robot.

The neural network representing the cognitive map is similar to Kohonen's [Koh82] self-organising map (SOM), where the map layer neurons represent the place cells. As in Kohonen's SOM, there is a winner neuron for each visual input with the strongest activation in the map layer whose connections to the input layer will get strengthened. The main difference, however, is that the neighbourhood relationship in this cognitive map model is not fixed. This follows from properties of neural place cells: they do not have a geometric connectivity as standard SOMs do have. During learning, the connections between the current winner cell w_t and the previous winner cell w_{t-1} are strengthened, resulting in a topological map of place cells with a variable number of connections per cell. One of the reasons for this choice is that there is no apparent relation between the spatial positions of place cells within the hippocampus and their corresponding place fields within the environment. Since the movement of the agent is continuous in space over time, place

cells representing adjacent places get connected. The topological map represents a relationship within the sensory space of the agent, and does not explicitly map the two-dimensional Cartesian space (we only consider agents moving on a plane). The map learning parameters were selected empirically in one experiment [Haf00b], and evolved using evolutionary strategies in another one [Haf03]. The parameters are the learning rates of both connection and input weights.
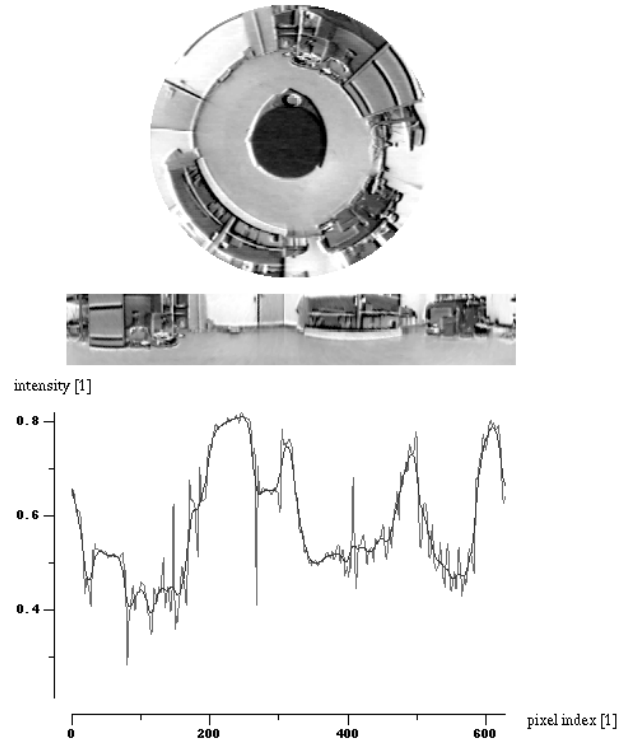

Fig. 3. Visual processing of the omnidirectional camera data from the mobile robot. The three steps are projection, horizontal averaging and low-passfiltering.

The available sensory information varied between the different experiments. In the case of the mobile robot that navigated in a standard office room using a compass and an omnidirectional camera, a 16-dimensional transformation of the omnidirectional camera image serves as sensory input to the neural network (see figure 3). These are equidistant visual input features based on an angular resolution of 22.5 degrees, horizontal averaging and low-pass filtering. In the simulation environment, 90-dimensional binary input from the cylinder world (see figure 1) has been used. Here we use only the information whether there is an obstacle in view at a certain angle or not. The problem of choosing the right sensory information is directly related to two complementary problems of reliable place recognition [KH02]: The first

is perceptual aliasing, which means that different places may have similar or even identical sensory information. The second is image variability. The same position and orientation may have different sensory information at different points in time. Possible reasons are sensory noise, motor noise, or simply change in illumination.

In contrast to the sparse topological representation of this approach, Arleo and Gerstner [AG00] use a population of place cells with overlapping place fields. A similar approach has also been taken by Gaussier et al. [GRBB02].
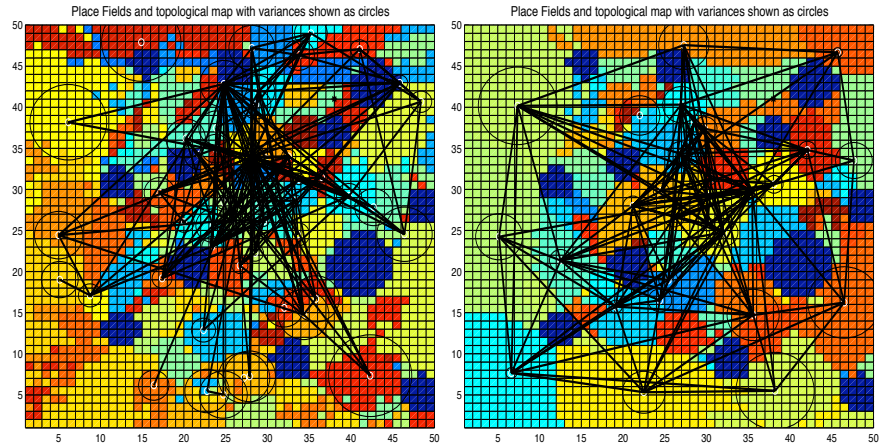
Fig. 4. Place fields after an exploration tour in a virtual environment with random learning parameters (left) and evolved learning parameters (right). The cylindrical objects (dark blue) are obstacles of different size. The coloured regions represent the different place fields. The black lines connect the centres of mass (the most sensible definition of a centre for an area with undefined shape) of the connected place cells within the environment.

3.2 Density of Place Fields

After an exploration tour in the simulated environment of figure 1, we can see that the same place cell is firing in certain restricted areas. We call these areas 'place fields'. An interesting aspect of the learned place fields after the exploration tour is that their number is significantly higher in the vicinity of objects (see figure 4 right, or other place field figures in [Haf00b]). This property has also been observed during electrophysiological recordings in rats, and tends to be explained by rats using a higher proportion of their place cells for 'interesting' places (for a review see [HDMM01]). Since neither the simulated nor the physical mobile robot have any concept of what counts as interesting, the explanation is simple: The visual input (sensory information) is changing more rapidly when moving close to convex obstacles, and therefore more place cells will be recruited for this space. This effect has

been explicitly encoded as a threshold for building new nodes in the graph algorithm by Franz et al. [FSMB98].

3.3 Extracting Metric Information

The map of the agent built from the place cell information and their connections is purely topological. However, it can be extended by using the additional information on heading directions between two place cells. Physiologically, this additional information might be accomplished by intersynapse connections between head direction cells and place cell connections. To enable the extraction of metric information from a topological map, a theoretical force model [Haf00a] is introduced, which is ideal for energy minimisation and assumes place cells as repulsive charges c_i and connections as springs s_{jk} connecting the cells (see figure 5). The spring constant for all springs is set to the same value based on the assumption that all directly connected places have the same distance. This is a generalisation since we have seen that there are place fields of different shape and size, however it is also clear that the distances of connected places are not too different, since two far-away places are linked by places between them. The initial position of the cells is random in $\mathbb{R}^2$. By repeatedly applying forces to the charges, their position converges to an energy minimum. The forces consist of an attractive spring force, a repulsive force and a third force caused by the preferred orientation of the connecting weight. Duckett et al. [DMS02] proposed a slightly different algorithm called the 'relaxation algorithm' which is based on similar principles, but additionally assumes distance information between the nodes, and assigns a position likelihood to them. The algorithm is computationally cheap, and ensures that an optimal solution will be found.
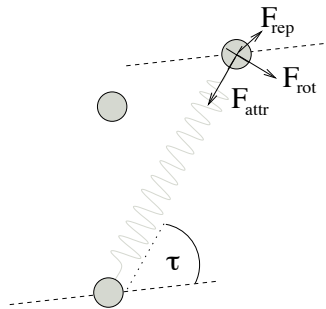
Fig. 5. Force model based on different forces applied on the place cells: an attractive spring force, a repulsive force and a third force caused by the preferred orientation of the connecting weight.

3.4 Evaluation Methods

Evaluating the usefulness of the learned cognitive map for robot navigation is difficult. In principle, we can distinguish three main approaches to the problem of
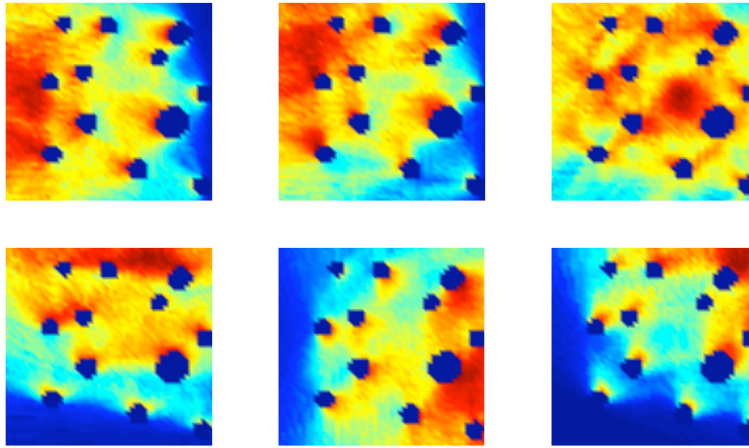
Fig. 6. Activation of six different place cells after an exploration tour in the virtual environment. The dark blue circles represent the obstacles.

evaluating the cognitive maps. First, the evaluation can be based on the properties of the map itself, such as the shape of place fields or the properties of connections between place cells (see figure 4). Second, the learning strategy of the cognitive map can be evaluated by assessing the behaviour of a navigating agent after exploration. And third, both properties of the cognitive map and the resulting navigation behaviour can be directly compared with those of a navigating animal. Let us first consider the method of analysing properties of the cognitive map: Statistical properties such as density, shape, and number of place fields; activity and number of connections per place cell, or metric versus graph distance in the topological map can be collected and analysed easily in simulation. On a mobile robot, additional difficulties arise since these data can only be collected having an accurate tracking system and having access to the sensory input for every position of the robot in space, ideally requiring a large image database.

The activity shapes of single place cells after an exploration tour in the virtual environment using the optimised cognitive map learning strategy can be seen in figure 6. The place field shapes are very similar to place fields of cells in the rat hippocampus, also showing a general exponential decay of activity away from the centre of the place field. Place fields near walls also have a tendency to be more elongated than place fields in the centre of the area, which tend to be more circular.

Assessing the behaviour of an agent to evaluate the cognitive map gives a good fitness measure, since it is focused on the behaviour, but has additional difficulties, since the number of exploration runs and navigation runs between two arbitrarily chosen places within the environment has to be huge in order to be meaningful. The problem with comparing robot and animal behaviour is, that both the environment

and the available sensory information should be comparable. An approach for evaluating a cognitive map learning strategy where the parameters have been evolved using evolutionary strategies can be found in [Haf03].

4 Conclusions

This chapter has given an example of how neural cognitive maps can be implemented in robotic experiments. For simplicity, some of these experiments were performed in simulation. One of the important features of the resulting neural cognitive maps is that it includes both topological and metric information about places. The difficulty is to realise the integration of these informations with a restricted amount of memory (number of place cells and their connections), and without building a geometric world model.

On the one hand, research on cognitive maps for robots has the advantage of exploiting biological principles of navigation for building more reliable robots, and on the other hand, robots are ideal tools to test specific hypotheses on the underlying functions of navigation behaviour in animals. The first point is important in particular for robot navigation in dynamic environments, and in environments, where GPS is not applicable, such as in many indoor environments, underwater and extraterrestrial. To test biological hypotheses, the important argument is the repeatability of robotic experiments and the easy access to internal states.

References

[AC04]   A. Abreu and L. Correia. Place energy and forget in maps for mobile robots. In *Proceedings of the IEEE International Conference on Robotics and Biomimetics - ROBIO2004*, 2004.

[AG00]   A. Arleo and W. Gerstner. Spatial Cognition and Neuro-Mimetic Navigation: A Model of Hippocampal Place Cell Activity. *Biological Cybernetics*, 83:287–299, 2000.

[Bro91]   R. A. Brooks. Intelligence Without Representation. *Artificial Intelligence Journal*, 47:139–159, 1991.

[DMS02]   T. Duckett, S. Marsland, and J. Shapiro. Fast, On-Line Learning of Globally Consistent Maps. *Autonomous Robots*, 12:297–300, 2002.

[ECH+05] A. D. Ekstrom, J. Caplan, E. Ho, K. Shattuck, I. Fried, and M. Kahana. Human hippocampal theta activity during virtual navigation. *Hippocampus*, 15:881–889, 2005.

[EKC+03] A. D. Ekstrom, M. J. Kahana, J. B. Caplan, T. A. Fields, E. A. Isham, E. L. Newman, and I. Fried. Cellular networks underlying human spatial navigation. *Nature*, 425:184–187, 2003.

[FSMB98] M. O. Franz, B. Schölkopf, H. A. Mallot, and H. H. Bülthoff. Learning View Graphs for Robot Navigation. *Autonomous Robots*, 5:111–125, 1998.

[GRBB02] P. Gaussier, A. Revel, J. P. Banquet, and V. Babeau. From view cells and place cells to cognitive map learning: processing stages of the hippocampal system. *Biological Cybernetics*, 86:15–28, 2002.

[Haf00a] V.V. Hafner. Cognitive Maps for Navigation in Open Environments. *Proceedings of the 6th International Conference on Intelligent Autonomous Systems (IAS-6)*, 2000.

[Haf00b] V.V. Hafner. Learning Places in Newly Explored Environments. *From Animals to Animats. SAB2000 Proceedings Supplement Book*, pages 111–120, 2000.

[Haf03] V.V. Hafner. Evaluating cognitive maps for mobile robot navigation behaviour. *Proceedings of the AISB'03 Symposium on "Scientific Methods for the Analysis of Agent-Environment Interaction"*, 2003.

[Haf04] V.V. Hafner. *Adaptive Navigation Strategies in Biorobotics: Visual Homing and Cognitive Mapping in Animals and Machines*. Shaker Verlag, 2004.

[HDMM01] S. A. Hollup, J. G. Donnett, M.-B. Moser, and E. I. Moser. Accumulation of Hippocampal Place Fields at the Goal Location in an Annular Watermaze Task. *Journal of Neuroscience*, 21(5):1635–1644, 2001.

[KH02] B. Kuipers and P. Beeson Horswill. Bootstrap Learning for Place Recognition. In *Proceedings of the 18th National Conference on Artificial Intelligence (AAAI-02)*, pages 824–829. AAAI/MIT Press, 2002.

[Koh82] T. Kohonen. Self-organizing formation of topologically correct feature maps. *Biological Cybernetics*, 43(1):59–69, 1982.

[O'K76] J. O'Keefe. Place units in the hippocampus of the freely moving rat. *Experimental Neurology*, 51:78–109, 1976.

[PB06] R. Pfeifer and J. C. Bongard. *How the Body shapes the Way We Think*. MIT Press, Cambridge, Mass., USA, 2006.

[Pre96] T. J. Prescott. Spatial representation for navigation in animats. *Adaptive Behavior*, 4(2):85–123, 1996.

[RH96] M. Recce and K. D. Harris. Memory for places: A navigational model in support of Marr's theory of hippocampal function. *Hippocampus*, 6:735–748, 1996.

[TS05] A. Tapus and R. Siegwart. Incremental robot mapping with fingerprints of places. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2005.

[TWBM97] O. Trullier, S. I. Wiener, A. Berthoz, and J.-A. Meyer. Biologically based artificial navigation systems: review and prospects. *Progress in Neurobiology*, 51:483–544, 1997.

[Web01] B. Webb. Can robots make good models of biological behaviour? *Behavioral and Brain Sciences*, 24(6), 2001.

[WM94] M. A. Wilson and B. L. McNaughton. Reactivation of Hippocampal Ensemble Memories During Sleep. *Science*, 265:676–679, July 1994.

[Wol02] L. Wolpert. Love is a many-moleculed thing. *The Observer*, 2002.

[WR79] R. Wehner and F. Räber. Visual spatial memory in desert ants: Cataglyphis bicolor (Hymenoptera: Formicidae. *Experientia*, 35:1569–1571, 1979.