

Evaluating cognitive maps for mobile robot navigation behaviour

Verena V. Hafner*

*Artificial Intelligence Laboratory
Department of Information Technology
University of Zurich, Switzerland
vhafner@ifi.unizh.ch

Abstract

This paper examines biologically inspired cognitive map models, which provide an artificial navigating agent with a topological map of places after an exploration and learning phase in a previously unknown environment. The evaluation requires analytical methods that either rate the cognitive map on the basis of the map's features, or that rate the observed behaviour of the navigating agent. Some possible solutions to this problem are investigated, including map optimisation using artificial evolution, and alternative navigation behaviours.

1 Introduction

The navigation strategies of animals can be categorised according to their complexity, ranging from systematic search at the simplest end of the spectrum to complex wayfinding, with intermediate strategies such as local visual homing (Trullier et al. (1997), Franz and Mallot (2000)). The same categories can be applied to the behaviour of mobile robots or simulated agents navigating in their environments, where the choice of navigation strategy is usually both task dependent and dependent on the sensory information available to the agent.

More complex navigation strategies require an internal (neural) representation of the environment and are often referred to as 'cognitive maps' (O'Keefe and Nadel (1978)). Neurophysiological and behavioural experiments suggest that rats and primates use cognitive maps for navigation, whereas most insects have to rely on more basic navigation strategies. The cognitive maps referred to here are based on the following principles: *Place cells* are neurons found mainly in the hippocampus of rats, their activation is dependent on the location of the rat within its environment (O'Keefe and Nadel (1978)). *Head direction cells* are neurons whose activity is dependent on the orientation of the rat's head within the environment (Taube et al. (1990)). *Place fields* are areas in the environment where a particular place cell has the strongest activation. The evidence for cognitive maps in animals has inspired roboticists to implement those biological findings to allow for stable and adaptive navigation behaviour in mobile robots coping with (partly) unknown and ever-changing environments.

Here, we consider biologically-inspired cognitive map models, which provide an artificial navigating agent with a topological map of places (possibly enhanced with additional metric information) after an exploration and

learning phase in a previously unknown environment. Some models and implementations can be found in Schölkopf and Mallot (1995), Hafner (2000), Gaussier et al. (2002), Filliat and Meyer (2002). We are now faced with the problem of how to evaluate these representations with respect to their usefulness for navigation behaviour. Evaluating the cognitive map in a quantitative way is difficult, and there are several levels of analysis that may be applied. The evaluation requires either analytical methods that rate the cognitive map on the basis of the map's features, or that rate the observed behaviour of the navigating agent. Both methods require a large number of different exploration runs in different environments to make a general statement on the usefulness of the navigation strategy.

In section 2, some possible methods are described. Section 3 shows the difficulties of evaluating cognitive maps for navigation on a practical example of a robot navigating in a simulated environment, and shows how these measures can be used for a fitness function to optimise the agent's exploration and learning strategy. In section 4, analysing the behaviour of the agent will be discussed and further possible approaches will be outlined.

2 Methods for Analysis

Methods for the analysis of simulated cognitive maps include the use of statistical properties such as density, standard deviation, shape, and number of place fields; number, covered space, and activity of place cells; number of connections per place cell, graph distance vs. metric distance of the cognitive map. These data can be collected easily in simulation. In principle, the same can be done for a real world mobile robot, but various problems increase the difficulty of the task: The exact robot position has to be known at any time, a sensor information

database with place data for a large number of positions may be necessary, and the number of runs has to be higher than in simulation to make up for ‘real world’ noise. Since two different maps of the same environment can result in similar navigation behaviour, it is necessary to consider the navigation behaviour of the robot in addition to the analysis of the map. One possibility of evaluating the behaviour of a robot with a particular navigation strategy is to answer the following question: How long does it take to get from position A to position B after having explored the environment? To have any statistical significance, this has to be done for a large number of different positions, runs, and environments.

3 Evaluating Cognitive Maps: An Example.

3.1 Exploration and Learning

In this section, an example of an agent learning how to navigate in a virtual environment is shown. During an exploration phase, the agent has the task of learning a map of the environment using its internal neural structure such as place cells and connecting synapses. The usefulness of the cognitive map for a specific environment depends crucially on the exploration and map learning strategy.

The environment consists of a plane with twelve cylinders of different diameter (see figure 1 for a top view). The agent is equipped with a compass and omnidirectional 1D view with a resolution of 90 pixels. The cylinders are black, walls are white. The agent starts

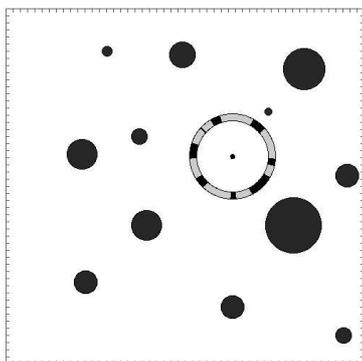


Figure 1: Virtual environment for the robot navigation experiments. The environment consists of a plane with twelve cylinders (shown from above) of different size. The recorded views are taken at the crossing points of a 50×50 grid (indicated by tick marks on the frame). The ring shows an example of an omnidirectional view in this environment.

with a given number of place cells which are initialised at the beginning with random weights to the visual

input. It performs random exploration tours within the environment, avoiding bumping into the obstacles. At each step, the current view is taken, resulting in a certain neural activation in the place cell layer. Since two views at different positions could look similar, the activation of the place cells is in addition dependent on the previous place cell activation and the movement direction of the agent. Information on the connection and the connection heading of two place cells representing adjacent places is - for simplicity - directly encoded, instead of being available through head direction cells.

During exploration and learning, the neural weights are adjusted as follows:

1. The activation a is calculated for the current view v for each place cell, $a = f(vw + c + o)$, where w are the weights connecting the visual input and the place cells. c is a value which enforces place cells with a lateral connection to the previous winning cell to win, o accounts for the heading of the agent, which is also stored in the lateral connections. f is a sigmoid function.
2. The weights w of the cell with the highest activation a are updated and moved towards the input:
 $\Delta w = \eta_1(v - w)$
3. The weights w_- of the cell with the previously highest activation are updated and moved towards the input:
 $\Delta w_- = \alpha \eta_1(v - w_-)$
4. w are normalised for each place cell
5. The lateral connections w_l between the winning cell and the previous winner are updated:
 $\Delta w_l = \eta_2 a_1 a_2 (1 - w_l)$
6. The orientation information γ_l in the lateral connections between the winning cell and the previous winner is updated.
7. A decay factor is applied to all lateral connections:
 $w_l = \delta w_l$

In figure 2, the place field representations of the agent are projected onto the 2D environment. Areas of same colour are represented by the same place cell. The centres of mass of these place fields are shown as small white circles, with variable sized circles indicating the variance of the specific place field. Connections are drawn between the centres of mass of connected place cells. A map is well suited for navigation purposes when it contains many, roughly circular place fields with connections between regions that are adjacent and can be traversed.

One problem which can be seen in figure 2 is that of multiple place fields. Here, one place cell fires at

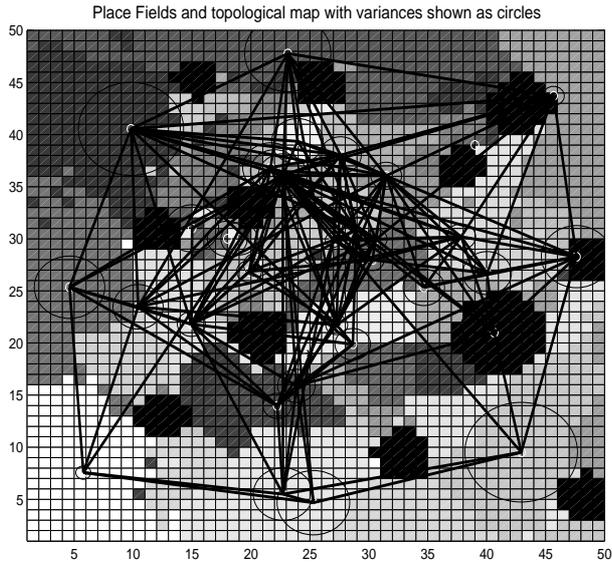


Figure 2: Map of the environment with black areas indicating obstacles, and grey coloured shadings indicating the different place fields after the agent has explored the environment. The centres of the place fields are interconnected as specified by the weights.

two spatially separated regions within the environment. Being unaware of this can result in strange behaviour. We know however that rats have to cope with the same problem. Multiple place fields have also been measured in the rat's hippocampus.

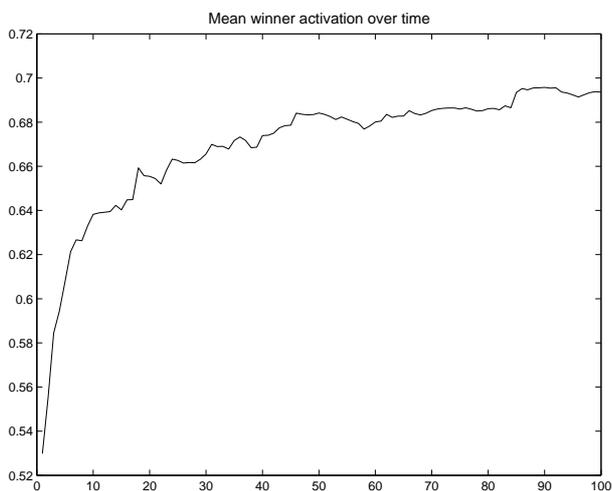


Figure 3: Average activation of the winner neuron for each grid position over time during the exploration and learning phase for one agent.

During the exploration and learning phase, some parameters introduced in section 2 which may be relevant to the quality of the resulting map are plotted over time

for 10,000 steps shown every 100th step. In figure 3, the mean activation of the winner place cell for each grid position is shown. Here, the place cell activation a is calculated purely from the visual input v and the neural weights w .

$$a = wv$$

Figure 4 shows the average variance v of the place fields over time for each place cell with a place field in the environment.

$$v_p = \frac{1}{n} \sum_i (x_i - \mu_p)^2$$

where μ_p is the centre of mass of the place field for place cell p , x_i a position on the grid where p is the winner cell (place cell with highest activation), and n the number of these positions. Figure 5 shows the number of connections between place cells as well as the average connection length using the metric distance between centres of mass of place fields from the plane in the simulation environment.

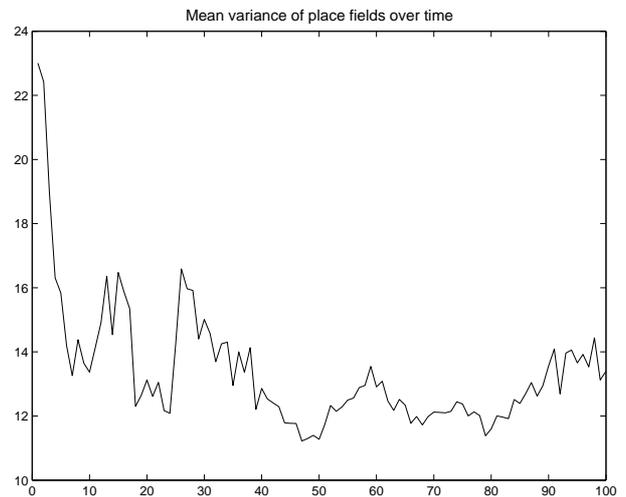


Figure 4: Average variance of the place fields over time during the exploration and learning phase for one agent.

3.2 Optimising Learning Parameters with Evolution Strategies

The learning strategy described in section 3.1 has been optimised to result in 'better' (to be defined) maps for the agent by evolving some of the following parameters:

- η_1 learning rate for weights w between visual input v and place cells
- α factor to learning rate of previous winning cells
- η_2 learning rate for lateral connections w_l between place cells
- β_1 contribution of lateral connections w_l to activation
- β_2 contribution of orientation γ_l to activation
- δ decay factor of map connections

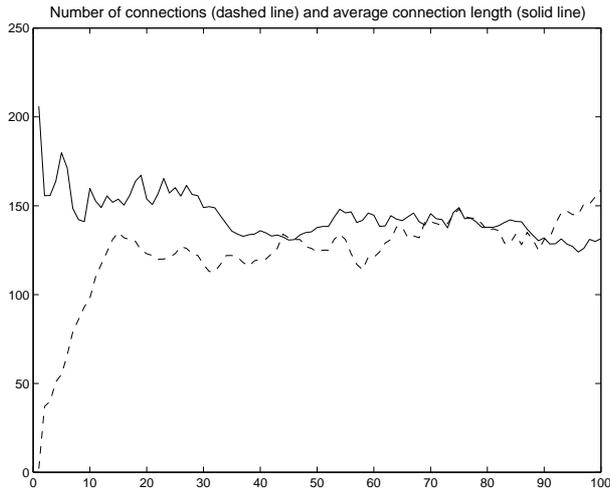


Figure 5: Number of connections and average connections length (10*metric distance of centres of mass of place fields) during the exploration and learning phase for one agent.

Artificial evolution strategies have been used with the following fitness function:

$$f = f_1 * f_2$$

where f_1 is a measure for the regions (the average variance of the place fields), and f_2 is a measure for the connections (between adjacent place fields). f_1 simply contains the mean variance of all place fields, whereas f_2 is a mixture of positive fitness for adjacent place fields whose place cells are connected and negative fitness for connected place cells whose place fields are not adjacent. An alternative fitness function would be the correlation between metric and graph distance for all connections.

The evolution ran on the same randomly generated path through the environment for each individual in a generation. The number of map neurons (place cells) was chosen to be 50, the evaluation ran for 3000 steps each. A detailed description of the evolution strategy and its parameters can be found in Schwitter (2003). The evolved learning parameters of the described cognitive maps are: $\eta_1 = 0.045, \alpha = 0.62, \eta_2 = 0.5, \delta = 0.998, \beta_1 = 0, \beta_2 = 0$. Figures 2-5 are produced using these parameters.

Surprisingly, the evolutionary algorithm came up with a solution for map learning, which only takes into account the current view for the activation of the place fields, and rejects an influence of lateral connections and orientation information to the activation calculation ($\beta_{1,2} = 0$). A possible explanation for this is that the connections are not available at the beginning of the exploration tour, and the orientation error being high. The high error of the orientation information of the connections is due to the

size of place fields. Traversing from one to another place field can be done in a range of different directions.

4 Analysing the Behaviour

We saw from the previous section that evaluating the topological map in a quantitative way is difficult and often needs an external observer, which cannot be the agent itself but someone with a metric map of the environment as well as access to the agent's brain. Analysing the navigation behaviour after exploration and learning is more straightforward, but it is very time consuming and dependent on the specific environment the agent interacts with. One possibility is to assess the behaviour for a high number of randomly chosen place pairs, measuring how often and how fast an agent can navigate from one to another place given the previously learned cognitive map. The success rates in our environment were varying strongly between places. One interesting effect that could be observed is asymmetry: The success rate of navigating from place A to place B can be different than when navigating from place B to place A . The reason is the higher importance of correct connections to the goal place.

Falling back on other navigation strategies like visual homing facilitates the success of the agent. For example, a version of the average landmark vector (ALV) model (Lambrinos et al. (2000)) could be used to navigate between places. The catchment areas in this specific environment are rather large and reliable. An example is shown in figure 6.

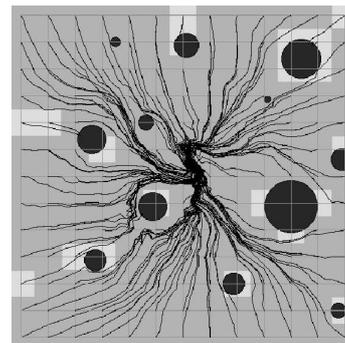


Figure 6: Trajectories of an agent using a learned version of the ALV model to navigate between two places in the environment used for the cognitive map learning.

Although the agent does not have a metric map itself, it is able to extract some metric information from the topological map when it contains some additional information like orientation of the connections. Examples of the application of such algorithms are the spring force model (Hafner (2000)) or the relaxation algorithm (Duckett et al. (2002)).

One has to be careful when drawing conclusions from observations of the behaviour or the cognitive map structure. An interesting example is the observation in rats, that their place field density is higher around obstacles. An outside observer would conclude that these places are ‘more interesting’ and therefore get a better neural representation. However, when the same experiment was conducted with an artificial agent, a higher place field density could be observed around obstacles without the robot having any concept of ‘interesting’ (compare figure 2 or Hafner (2000)). The effect can simply be explained by the sensory information changing more rapidly in the vicinity of obstacles.

5 Discussion

We have shown on the example of learning cognitive maps optimal for navigation that evaluating these is often difficult because of the lack of suitable analysis methods. This applies even stronger for real world mobile robots, and leads us back to the well-known discussion on simulation vs. real-world experiments (Jakobi (1998)). We are either restricted using an unnatural, reproducible toy-world (or simulation), or have a robot in a natural environment (if an office room can be called natural) where the experiments cannot be compared with experiments in other environments. It may be more conclusive to compare the achieved navigation behaviour of a robot not just to other robots and other environments, but directly to the behaviour of navigating animals and humans in the same environment. What remains tricky is the effect of the difference in sensory modalities (e.g. vision, olfactory and somatosensory cues for rat navigation compared with camera and touch sensors on a mobile robot). Therefore, virtual environments would be ideally suited to compare human navigation behaviour with autonomous agent behaviour under the same conditions. There is even a virtual-reality setup for rats in preparation (Dahmen (2003)). A closer collaboration between behavioural robotics and fields like biology, psychology and mathematics would be a first step to create comparable and reproducible results, without having to resort to artificial quantification methods.

Acknowledgements

Many thanks to Christoph Schwitter for implementing the artificial evolution strategy as part of his Semesterarbeit. This research has been supported by grant #20-61372.00 of the Swiss National Science Foundation.

References

- H. Dahmen. Neural Correlates for Spatial Cognition in Virtual Environments. 2003. URL http://www.uni-tuebingen.de/cog/forschung/rodents_place/correlates_spatial_cognition.htm.
- T. Duckett, S. Marsland, and J. Shapiro. Fast, On-Line Learning of Globally Consistent Maps. *Autonomous Robots*, 12:297–300, 2002.
- D. Filliat and J. A. Meyer. Global localization and topological map-learning for robot navigation. *Proceedings of the 7th international conference on simulation of adaptive behavior (SAB)*, 7:131–140, 2002.
- M. O. Franz and H. A. Mallot. Biomimetic robot navigation. *Robotics and Autonomous Systems*, 30:133–153, 2000.
- 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.
- V. V. Hafner. Cognitive Maps for Navigation in Open Environments. *Proceedings of the 6th International Conference on Intelligent Autonomous Systems (IAS-6)*, 2000.
- N. Jakobi. The Minimal Simulation Approach to Evolutionary Robotics. *Proceedings of ER’98, AAI Books, Kanata*, 1998.
- D. Lambrinos, R. Möller, T. Labhart, R. Pfeifer, and R. Wehner. A mobile robot employing insect strategies for navigation. *Robotics and Autonomous Systems, special issue on Biomimetic Robots*, 30:39–64, 2000.
- J. O’Keefe and L. Nadel. *The hippocampus as a cognitive map*. Oxford University Press, 1978.
- B. Schölkopf and H. A. Mallot. View-Based Cognitive Mapping and Path Planning. *Adaptive Behavior*, 3(3): 311–348, 1995.
- C. Schwitter. Learning cognitive maps for navigation - Evolving parameters. *Semesterarbeit AI Lab, Univ. of Zurich*, 2003.
- J. S. Taube, R. U. Muller, and J. B. Ranck. Head-direction cells recorded from the postsubiculum in freely moving rats. *Journal of Neuroscience*, 10:420–447, 1990.
- 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.