2015 Special Issue

Goal-oriented robot navigation learning using a multi-scale space representation

CrossMark

M. Llofriu [a,c,*], G. Tejera [c], M. Contreras [b], T. Pelc [b], J.M. Fellous [b], A. Weitzenfeld [a]

[a] University of South Florida, United States
[b] University of Arizona, United States
[c] Universidad de la Republica, Uruguay

ARTICLE INFO

ABSTRACT

There has been extensive research in recent years on the multi-scale nature of hippocampal place cells and entorhinal grid cells encoding which led to many speculations on their role in spatial cognition. In this paper we focus on the multi-scale nature of place cells and how they contribute to faster learning during goal-oriented navigation when compared to a spatial cognition system composed of single scale place cells. The task consists of a circular arena with a fixed goal location, in which a robot is trained to find the shortest path to the goal after a number of learning trials. Synaptic connections are modified using a reinforcement learning paradigm adapted to the place cells multi-scale architecture. The model is evaluated in both simulation and physical robots. We find that larger scale and combined multi-scale representations favor goal-oriented navigation task learning.

© 2015 Elsevier Ltd. All rights reserved.

1. Introduction

We present a model for spatial cognition based on the multi-scale organization of the dorsal–ventral axes of the hippocampus. The theoretical concept of a cognitive map in the brain was first proposed by Tolman (1948) as the essential module responsible for estimating the rat's position in the environment. Through an extensive review of the literature at the time, O'Keefe and Nadel proposed that the cognitive map laid within the brain's hippocampus (O'Keefe & Nadel, 1978). This suggestion was in part based on the discovery of O'Keefe and Dostrovsky (O'Keefe & Dostrovsky, 1971) place cells in the rat's hippocampus, termed as such due to the high correlation between their firing and the rat location in the environment. The region of the environment specific to each cell, is called their place field.

This property turns these cells into a rich source of information for navigational purposes, as they provide an encoding of the rat location. Even though no individual place cell provides accurate location information by itself, it has been shown that the location of the animal can be predicted with an error of 1 cm based solely on the activity levels of an ensemble of cells (Guger et al., 2011; Jensen & Lisman, 2000; Wilson & McNaughton, 1993; Zhang, Ginzburg, McNaughton, & Sejnowski, 1998).

The discovery of grid cells in the rat's Medial Entorhinal Cortex (MEC) was first published by Fyhn, Molden, Witter, Moser, and Moser (2004), it suggested that grid cell firing signaled the rat's changing position in the environment. Hafting, Fyhn, Molden, Moser, and Moser (2005) presented findings that grid node spacing increased in a modular fashion in MEC in overall correspondence with the gradual increase in place field size along the dorsoventral axis of the hippocampus (Brun et al., 2008; Jung, Wiener, & McNaughton, 1994; Maurer, Vanrhoads, Sutherland, Lipa, & McNaughton, 2005). For a review of the multi-scale representation of space, we refer the reader to Geva-Sagiv, Las, Yovel, and Ulanovsky (2015).

The MEC also contains head-direction cells whose activity is related to the current head orientation of the rat (Blumberg, 2015). Altogether, these discoveries suggest that spatial navigation may originate in a fundamental multi-scale representation of space along the dorso-ventral axis of the hippocampal formation, modulated by self-motion and external cues that includes, although not limited to, visual detection of distal and proximal landmarks. Recent work has shown a gradient of head direction cells tuning width along the dorso-ventral axis of the MEC (Giocomo et al., 2014) in the upper layers (inputs to the hippocampus), but not in the deep layers. While there is some theoretical evidence that grid

cells may help the accuracy of spatial navigation (Guanella & Verschure, 2007), there is little evidence that they do so experimentally (Hales et al., 2014).

There is extensive research on spatial cognition models inspired by place cells coding in the rat's hippocampus used to evaluate goal-oriented spatial navigation with simulation and with real robots (Arleo, Smeraldi, & Gerstner, 2004; Barrera & Weitzenfeld, 2008; Brown & Sharp, 1995; Burgess, Recce, & O'Keefe, 1994; Caluwaerts et al., 2012; Dollé, Sheynikhovich, Girard, Chavarriaga, & Guillot, 2010; Filliat & Meyer, 2002; Gaussier, Revel, Banquet, & Babeau, 2002; Guazzelli, Bota, Corbacho, & Arbib, 1998; Krichmar, Nitz, Gally, & Edelman, 2005; Milford & Wyeth, 2007, 2009; Pata, Escuredo, Lallée, & Verschure, 2014; Recce & Harris, 1996; Redish & Touretzky, 1997; Sukumar, Rengaswamy, & Chakravarthy, 2012; Tejera, Barrera, Fellous, Llofriu, & Weitzenfeld, 2013). However, few of them incorporate some aspects of multi-scale representation of space.

Chen, Jacobson, Erdem, Hasselmo, and Milford (2013) implement an array of support vector machines on video segments to recognize places, where different segment lengths represent different scales of representation. The introduction of larger scales of representation improves recall in their classification system. Pata et al. (2014) develops a model of the hippocampus to explore the functional differences of DG and CA3 while accounting for the differences in scale across the dorsoventral axis. These two have a strong focus on place and grid cell formation. Our work, in contrast, focuses in goal directed navigation, on how to use the information provided by these cell to reach a desired goal.

Erdem and Hasselmo (2014) adds multiple scales of representation to a previous model of spatial navigation based on forward lookup probes, which resemble short-wave ripple (SWR) activity in the hippocampus. The addition of larger scales of representation improves the effective distance of forward lookup probes, improving navigational performance. We believe our work complements this approach. Short-wave ripple activity have been suggested to guide navigation (Johnson & Redish, 2007; Pfeiffer & Foster, 2013), but it occurs during sleep or when the rat is still (Foster & Wilson, 2006). Thus, while Erdem and Hasselmo (2014) work focuses on high level planning during key decision points, our model focuses on the decision making that takes place while the rat is in motion. Our working hypothesis is that this decision making can also benefit from different scales of representation.

We based our study on the biological role of dorsal and ventral hippocampal place cells and contrast the roles of the small, medium and large place fields represented across the septo-temporal (dorsoventral) axis of the hippocampus. We develop a neural architecture of multi-scale hippocampal place cells to be evaluated during a goal-oriented robot navigational task. This task involves learning to locate a fixed goal in a circular arena, inspired by the Morris' water maze task, where instead of a submerged platform, the goal can only be recognized when the robot is very close to it (de Jong, Gereke, Martin, & Fellous, 2011). Distal cues are set on the perimeter of the arena to facilitate localization. The task involves multiple trials during which navigation to the goal is reinforced by applying a Q-Learning algorithm (Sutton & Barto, 1998), inspired by the neuro-modulatory effects of dopamine (Cox & Krichmar, 2009), and adapted to the multi-scale nature of place fields. The task is evaluated using computer simulations and physical robots. Section 2 describes the spatial cognition model, Section 3 presents the goal-oriented task and the experimental layout, Section 4 presents simulated and robot experimental results, Section 5 includes a discussion of the results and Section 6 provides conclusions and discussion of future work.

2. The spatial cognition model

The spatial cognition model is comprised of six main modules, described below and shown in Fig. 1.

It has been proposed that navigation involves the interaction of four components: place cells, head direction cells, local view and path integration (Redish & Touretzky, 1997; Touretzky, 2002). We consider our path integration and local view components as solved. Namely, place cell firing values are derived from sources of location information directly, rather than computing them from path integration and visual information, as will be explained in the Experiments section. Thus, we focus in this work on the place cell and head direction cells components and their contribution to learning using multiple scales. Our model uses this multi-scale representation as the information source for a reward driven learning system (Krichmar & Röhrbein, 2013).

2.1. Modules

Place Cell Module. This module calculates the firing of a population artificial place cells. They take the current position x of the robot as input and calculate the firing rate as Eq. (1).

$$f_i = exp\left(-\frac{(x - c_i)^T \Sigma_i^{-1}(x - c_i)}{2}\right). \tag{1}$$

Where f_i is the firing rate of cell i, c_i its preferred location and Σ_i its covariance matrix. Namely, each cell fires according to a 2D gaussian function with a center on each place cell preferred position, as modeled by O'Keefe and Burgess (1996).

The key of this work involves the use of different scales of place cells, which we map to choosing different Σ. The covariances matrix are always of the form $\sigma^2 I$, where σ^2 models de specificity and I is the identity matrix.

Head Direction Module. This model computes the firing of a population of artificial head direction cells. This module takes the current heading θ of the robot and compute the firing rate of each cells as Eq. (2).

$$f_i = exp\left(-\frac{(\theta - \theta_i)^2}{2\sigma^2}\right). \tag{2}$$

Where f_i is the firing rate of the ith head direction cell, σ^2 its variance and θ_i its preferred orientation. Thus, this cells are also computed as a gaussian function with the peak in the cell's preferred value.

Multi-Scale QL Module. This module performs Q-Learning on the information provided by the place cells and head direction cells. Place and orientation information is obtained by selecting all possible pairs from both sets and computing the resulting activity as the product of both the place cell and head direction cell. This combined source of information is passed onto the QL module, which will be explained in detail below.

Taxic Behavior Module. This behavior moves towards a visible goal. It works cooperatively with the QL learning module by assigning a fixed value to the action that will take the robot to the goal. In the framework proposed by Guazzelli et al. (1998), this module corresponds to the execution of the affordance of going to a visible goal.

Exploration Behavior. This module promotes exploration in early phases of an experiment. The exploration value is calculated as shown in Eq. (3), where episode is the episode number, maxReward is the maximum reward possible given to the robot, and β is a given parameter that models how fast the exploration