Modeling Rodent Head-direction Cells and Place Cells for Spatial Learning in Bio-mimetic Robotics

Angelo Arleo **Wulfram Gerstner**
Centre for Neuro-Mimetic Systems, MANTRA
Swiss Federal Institute of Technology Lausanne,
CH-1015 Lausanne EPFL, Switzerland
angelo.arleo@epfl.ch, wulfram.gerstner@epfl.ch

Abstract

We propose a computational model which is consistent with several neurophysiological findings concerning biological head-direction cells and hippocampal place cells. The model consists of two separate neural systems providing directional and place coding representations, respectively. These two modules are strongly coupled and interact with each other to form a unitary spatial learning system. We stress the importance of correlating idiothetic and allothetic signals to determine the dynamics of the system in order to stabilize head-direction and place representations over time. We give experimental results obtained by implementing the entire model on a real mobile robot.

1. Introduction

Directional sense and place coding are crucial capabilities for solving spatial cognitive tasks. Neurophysiological findings suggest the existence of neural representations of direction and position as a basis for animal spatial behavior (O'Keefe & Nadel, 1978; Taube et al., 1990).

Experimental recordings from freely-moving rodents show the presence of *head-direction cells*, neurons which fire as a function of the animal's allocentric heading in the azimuthal plane (Taube et al., 1990; Blair & Sharp, 1995). A head-direction cell i fires maximally only when the animal's current heading $\theta(t)$ is equal to the cell's *preferred direction* θ_i, regardless of the orientation of the head relative to the body, of the animal's spatial location, or of the animal's ongoing behavior. Thus, the ensemble activity of head-direction cells works as an allocentric compass, providing a neural support for the animal's sense of direction (McNaughton et al., 1991).

Head-direction cells have been observed in several brain regions and in particular in the postsubiculum (PSC), the anterodorsal thalamic nuclei (ADN), the laterodorsal thalamic nucleus (LDN), the lateral mammillary nuclei (LMN), the parietal and cingulate cortices, and the dorsal striatum. The postsubicular cortex (PSC), the anterodorsal nucleus (ADN), and the lateral mammillary nuclei (LMN) seem to be primary structures with respect to the directional selectivity property (see (Redish, 1997) for a review).

Neurons whose activity is tuned to the animal's current location $\mathbf{p}(t)$ have been recorded from the rat hippocampal formation (O'Keefe & Nadel, 1978). These neurons, termed *place cells*, show a high spiking frequency only when the rat is in a specific region of the environment, which defines the *place field* of the cell. Due to the place-coding property of cells in the rat hippocampus, this area of the brain has been thought to play a major role in mammalian spatial navigation.

Place cells have been experimentally recorded from the hippocampus proper (CA1 and CA3 areas), and from other extra-hippocampal regions such as the entorhinal cortex (EC), the dentate gyrus, the subiculum, and the parasubiculum (Redish, 1997). Lesions to the hippocampal formation disrupt rodents' spatial learning, which supports the hypothesis that the hippocampus is primarily involved in space representation and navigation (O'Keefe & Nadel, 1978).

Head-direction and place cells are not two separate circuits. Rather, they are strongly coupled (Knierim et al., 1995), which suggests they might form a single neural substrate for spatial navigation.

In our model, we implement head-direction cells and place cells as two interacting modules. Consistently with experimental data (Knierim et al., 1995), inhibiting either one of these two modules would critically impair the performance of the other and would disrupt the spatial learning capability of the system. Fig. 1 shows a functional overview of the entire navigation system. Dark grey areas form the space representation module (modeling hippocampal place cell activity), whereas light grey areas provide head-direction representation (modeling directional cells). Both systems rely on allothetic signals (e.g., visual cues) as well as idiothetic information (e.g., vestibular and proprioceptive cues) to maintain stable internal representations. Finally, the combined output of both systems is used for motor commands to achieve goal-oriented navigation.

In order to validate the model experimentally, we have implemented it on a real mobile Khepera robot (Fig. 2(a)). The experimental setup consists of a square arena within which the robot can freely move (Fig. 2(b)). The robot's behavior

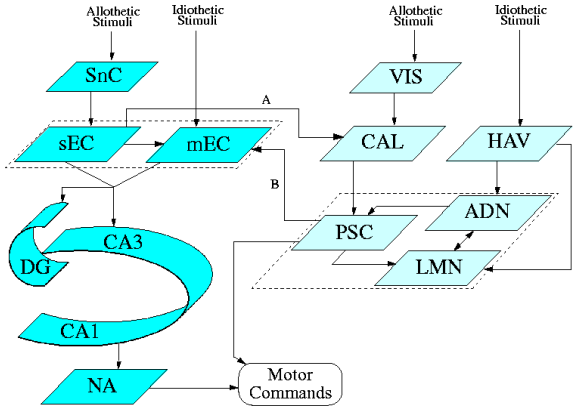

Figure 1: An overview of the entire system. Glossary: SnC: snapshot cells, sEC: superficial entorhinal cortex, mEC: medial entorhinal cortex, DG: dentate gyrus, CA3-CA1: hippocampus proper, NA: nucleus accumbens, VIS: visual bearing cells, CAL: calibration cells, HAV: head angular velocity cells, PSC: postsubiculum, ADN: anterodorsal nucleus, LMN: lateral mammillary nuclei.

is monitored by means of a camera above the arena. The robot's sensory system consists of eight infrared sensors to detect obstacles and measure ambient light, a light-detector, an on-board camera for vision-based self-localization, and an odometer for sensing internal self-motion signals.

2. Modeling Head-direction Cells

We model biological directional cells by means of a neural architecture in which internal and external stimuli are combined to establish a stable head-direction representation over time. The dynamics of the system is primarily controlled by idiothetic signals (e.g., proprioceptive inputs) which determine the directional selectivity property. Allothetic information (e.g., visual input) is used to occasionally modify the system's dynamics and calibrate head-direction cell activity.

Fig. 3 shows a functional overview of the model. The architecture may be understood as consisting of three functional modules: *(i)* an angular velocity integrator, *(ii)* a neural substrate forming the output of the system, and *(iii)* a module incorporating external signals to achieve system calibration.

2.1 Integrating Robot Angular Velocity

In order to integrate the angular velocity ω over time, we consider a directional circuit involving three neural populations, namely HAV, ADN, and LMN (Fig. 3).

Head Angular Velocity Cells. To assess angular displacements, we consider a population HAV$= \{ \Omega_+, \Omega_=, \Omega_- \}$ of cells whose activity is correlated to the sign and magnitude of the angular velocity ω. Cells Ω_+ and Ω_- fire proportionally to $|\omega|$ only during clockwise and counterclockwise turning, respectively. Cell $\Omega_=$ fires only when the robot is not rotating. Biological cells encoding animal's head angular velocity have been observed in the parietal, somatosensory, and visual cortices (McNaughton et al., 1991; Blair & Sharp, 1995).

(a) (b)

Figure 2: (a) The mobile Khepera robot. (b) A view of the $80 \cdot 80$ cm experimental arena as seen by the robot.

LMN Directional Cells. LMN cells model neurons in the lateral mammillary nuclei which are tuned to the animal's heading θ as well as the head angular velocity ω (Leonhard et al., 1996).

We discretize the continuous angular space $\Theta = [0, 360]$ in $S = 180$ steps of 2° each. Then, we define a population LMN$= \{ \{ l_z, l_s, l_m, l_f \} \mid 1 \leq l \leq S \}$ of $S \cdot 4 = 720$ cells with evenly distributed preferred directions θ_l. In particular, we take four cells l_i for each direction θ_l. That is, the magnitude of the angular velocity $|\omega|$ is discretized by LMN cells in four possible states which are ω_z (zero), ω_s (slow), ω_m (medium), and ω_f (fast). We assume that $\omega_z = 0^\circ/sec$, $\omega_s = 100^\circ/sec$, $\omega_m = 200^\circ/sec$, and $\omega_f = 400^\circ/sec$. A LMN cell fires maximally only when the robot's heading is equal to its preferred direction *and* the robot is turning at a specific angular velocity. For instance, the activity of cell $l_m \in$ LMN is maximal only if $\theta = \theta_l$ *and* $|\omega| = \omega_m$.

ADN Directional Cells. Cells in the anterodorsal thalamic nucleus have the key property of predicting future animal's headings during head turning, but encoding the current head direction during no rotation (Blair & Sharp, 1995). Indeed, a ADN cell a shifts its preferred direction θ_a as a function of ω such that it may temporally anticipate the head direction θ by a time delay $\tau_a \geq 0$

$$\theta_a(\omega) = \theta - \omega\tau_a \qquad (1)$$

According to recent experimental findings (Blair, 1996), we consider τ_a as characteristic for each ADN cell a of the model. In other words, we assume that some ADN cells may anticipate head-direction more than others. We consider five distinct time delays $\tau_1 < \ldots < \tau_5$ for each preferred direction θ_a, such that $\tau_1 = 20ms$, $\tau_2 = 40ms$, $\tau_3 = 60ms$, $\tau_4 = 80ms$, and $\tau_5 = 100ms$. Then, we define a population ADN$= \{ \{ a_1, \ldots, a_5 \} \mid 1 \leq a \leq S \}$ in which five ADN cells a_j (each of which has a different time delay τ_j) code for each direction θ_a. Every a_j cell anticipates θ by a different angle when $|\omega| > 0$, but all preferred directions θ_{a_j} are equal to θ when $\omega = 0$.

Idiothetic-based Dynamics. LMN and ADN interact with each other in order to integrate angular motion over time, that is to update the direction representation according to HAV cell activity. In Appendix A we describe the synaptic connections and the neural activity of the circuit in detail.

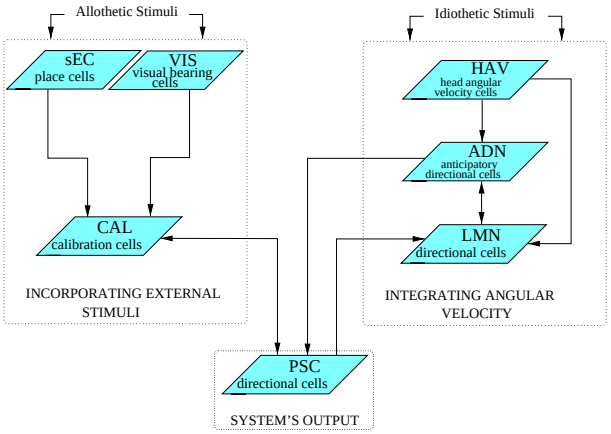

Figure 3: A functional overview of our directional system.

ADN cells play a major role in determining the system's dynamics. They induce the synchronized shifting behavior enabling directional cells to track robot rotations. ADN cells are driven by LMN and HAV cells. In particular, LMN feeds ADN with information about the current heading $\theta(t)$, whereas HAV determines the direction and velocity of the ADN activity shift. In order to induce the ADN anticipatory property (Eq. 1), we employ LMN $\rightarrow$ ADN *matching* and *offset* projections (Skaggs et al., 1995), such that the higher τ_j, the more a cell a_j is informed in advance about the current robot's heading $\theta(t)$, and the more the cell may anticipate the future robot's direction $\theta(t')$. Fig. 4 presents a result obtained by recording the firing activity of a cell a_j in the ADN layer of the model. The time delay of the recorded cell is $\tau_j = 40ms$. When the robot is not rotating, the cell exhibits a preferred direction $\theta_a = 180^\circ$ (dashed line). During counterclockwise turning with $|\omega| = \omega_f = 400^\circ/sec$, the cell predicts the future heading by shifting its mean firing direction (solid curve) of approximately 16°, which is consistent with Eq. 1.

LMN cells are correlated to the heading $\theta(t)$ as well as to the angular velocity ω. Each LMN cell l_i receives afferents from five ADN cells a_j (with, $\tau_1 < \ldots < \tau_5$) and modulates its firing activity according to HAV cell activity. ADN cells inform LMN cells about future headings. Due to the several time delays τ_j associated to each preferred direction, each LMN cell l_i is actually informed about arrival at θ_l by a sequential activation on its ADN inputs. We assume that only when *all* ADN inputs a_j have sequentially predicted arrival at θ_l, cell l_i is enabled to fire. Whenever this is true, activity $r_{l_i}(t)$ is simply given by averaging the afferent ADN activity at time t (see Appendix A). Fig. 5 shows a result obtained by recording the firing activity of a LMN cell l_i in our model.

2.2 *Interpreting the Directional Output*

PSC Directional Cells. We consider a population of cells $PSC = \{\, p \mid 1 \leq p \leq S \,\}$, with evenly distributed preferred directions θ_p, forming the output of the directional system (Fig. 3). Thus, at each time t, the PSC ensemble activity provides the estimate $\overline{\theta}(t)$ of the robot's allocentric heading $\theta(t)$.

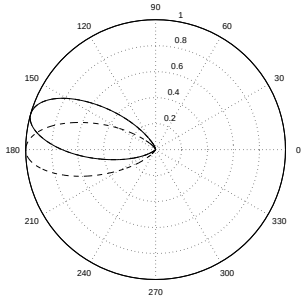

Figure 4: Polar representation of the firing rate $r_{a_j}(\theta)$ of a ADN cell a_j with $\tau_j = 40ms$. When the robot is not rotating the preferred direction is $\theta_a = 180^\circ$ (dashed line). During counterclockwise turning with $|\omega| = 400^\circ/sec$, the cell anticipates the future heading by about 16° (solid line).

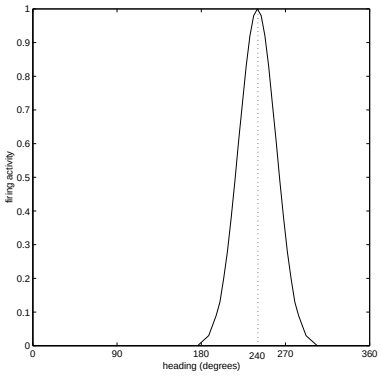

Figure 5: Tuning curve of a LMN cell l_i in our model. The mean preferred direction is $\theta_l = 240^\circ$ (dashed line).

PSC cells are primary driven by ADN cells. In particular, ADN $\rightarrow$ PSC projections are defined according to the same scheme used for ADN $\rightarrow$ LMN connections (see Appendix A). Also, PSC activity is based on the same mechanism used to drive LMN cells. That is, a cell p is enabled to fire only when all its ADN inputs have been sequentially activated.

In order to interpret the PSC ensemble activity as output of the directional system, we apply *population vector coding* (Georgopoulos et al., 1986). Thus, we compute the estimation of the robot's heading $\overline{\theta}(t)$ according to

$$\overline{\theta}(t) = \arctan\big(\sum_p \sin(\theta_p)\, r_p(t) / \sum_p \cos(\theta_p)\, r_p(t)\big) \quad (2)$$

where $r_p(t)$ is the firing rate of PSC cell p, and θ_p is its preferred direction. Fig. 6 shows a population activity as recorded from our PSC layer during robot turning. According to Eq. 2, the ensemble activity codes for the allocentric heading $\overline{\theta}(t) \approx 180^\circ$.

When the robot first enters the arena, the system is initialized with respect to an arbitrary absolute direction Φ. We set the population activity of PSC, LMN, and ADN cells as a Gaussian distribution centered at Φ and with $\sigma = 20^o$. That is, the activity of a cell i is initialized as $r_i(t) = \exp(-(\theta_i - \Phi(t))^2 \,/\, 2\sigma^2)$, where θ_i is the preferred direction of cell i. As the robot starts rotating, the system starts shifting

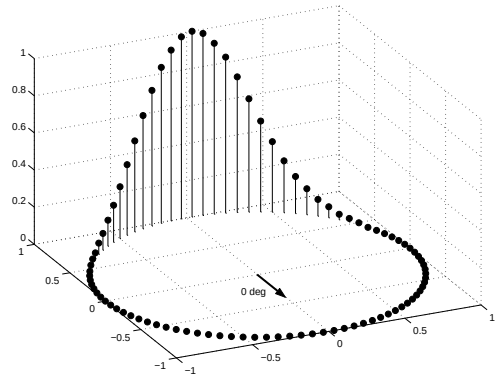

Figure 6: PSC population activity coding for a robot's heading of $180°$.

its internal representation to track angular displacements.

2.3 *Using Allothetic Cues for Calibration*

Head-direction cells which only rely on inertial stimuli induce a directional sense affected by cumulative tracking error (McNaughton et al., 1991). Our directional system incorporates allothetic information (e.g., visual cues) to generate a stable representation.

We let the robot estimate its egocentric bearing α relative to a weak light source L which is located on one of the walls of the arena. The idea is to use the bearing signal α for calibrating the directional system. However, since L is not an infinitely distant cue, it does not provide an absolute directional reference (as the sun does for desert ants and bees). Thus, compared to the absolute framework defined by Φ, the egocentric bearing α is not invariant with respect to spatial location. As a consequence, we need to combine the bearing information α with some place coding information. We consider a calibrating system (Figs. 1, 3) made of:

(i) A population of visual cells (VIS) whose activity encodes the robot's egocentric bearing $\alpha(t)$ at time t.

(ii) A population of place cells coding for the current robot's spatial location $\mathbf{p}(t)$. In particular, we consider place cells in the superficial entorhinal cortex (sEC) of our hippocampal model (projection A in Fig. 1), whose activity depends on visual cues only (see Sec. 3.).

(iii) A population of calibration cells (CAL), driven by VIS and sEC cells, firing as a function of $\alpha(t)$ and $\mathbf{p}(t)$. We call these neurons calibration cells because their activity is directly used to calibrate directional cells. CAL cells project directly onto PSC. Then, PSC propagates the calibration signal to LMN cells, which in turn start driving ADN cells based on recalibrated signals.

Visual Bearing Cells. A VIS cell v fires as a function of the current bearing angle α between the robot and the light cue L. Biological cells which respond maximally only when an external stimulus arrives from a particular egocentric direction have been observed in the inferior parietal cortex, the internal medullary thalamic lamina and the superior collicus (McNaughton et al., 1991; Skaggs et al., 1995; Knierim et al., 1995). For the robotic implementation, we consider a population VIS $= \{\, v \mid 1 \leq v \leq S \,\}$ of $S = 180$ cells with preferred direction evenly distributed over $360°$. To interpret the VIS ensemble activity $\mathcal{R}^{vis}(t) = \{\, r_v(t) \mid 1 \leq v \leq S\}$, we apply population vector coding (Eq. 2).

The robot detects light by means of eight infrared sensors, and one light-sensor facing forward (Fig. 2(a)). These nine sensory inputs $\mathbf{s}(t) = (s_1(t), \ldots, s_9(t))$ have to be interpreted to drive VIS activity $\mathcal{R}^{vis}(t)$. To deal with such a noisy and high-dimensional input, we trained a feed-forward neural network $\mathcal{N}$ to approximate the mapping function

$$\mathcal{M} : \mathcal{S} \rightarrow \mathcal{R}^{vis}(\alpha) \tag{3}$$

where $\mathcal{S}$ is the input sensory space. The idea is to let the robot learn $\mathcal{M}$ only once by off-line supervised learning. Network $\mathcal{N}$ is trained by using gradient descent back-propagation. The network's input consists of the sensory reading vector $\mathbf{s}(t) = (s_1(t), \ldots, s_9(t))$. The output is the VIS cell population activity $\mathcal{R}^{vis}(t)$.

Calibration Cells. CAL cells play a central role in the calibration process. They form the "neural boundary" between the allothetic visual-based representation and the idiothetic movement-related representation. On the one hand, CAL cells receive afferents from both VIS and sEC cells. On the other hand, CAL cells are interconnected with PSC cells. In particular, each CAL cell c receives an input connection from one and only one PSC cell p, such that $\theta_c = \theta_p$. By contrast, cell c projects efferents onto all PSC cells. We consider a population CAL $= \{\, c \mid 1 \leq c \leq S \,\}$ of $S = 180$ cells.

Correlating Allothetic and Idiothetic Signals. We let the robot explore the environment and apply LTP correlational learning to modify synaptic connections VIS $\rightarrow$ CAL, sEC $\rightarrow$ CAL, and CAL $\rightarrow$ PSC. In particular, the aim is to associate the heading $\overline{\theta}(t_1)$ encoded by PSC cells at time t_1, with the bearing angle $\alpha(t_1)$ *and* the spatial location $\mathbf{p}(t_1)$ encoded by VIS and sEC cells, respectively, at time t_1. As a consequence, if at time t_2 the animat is at position $\mathbf{p}(t_2) \approx \mathbf{p}(t_1)$ with bearing $\alpha(t_2) \approx \alpha(t_1)$, it may calibrate its directional cells by recalling the memorized activity pattern corresponding to $\overline{\theta}(t_1)$. In other words, Hebbian learning enables the system to associate allothetic and idiothetic representations based on the agent's experience. CAL cells function as a long-term memory device: they allow the agent to store "snapshots" of PSC cell activity and use the combined signal $(\alpha(t),\ \mathbf{p}(t))$ to recall this activity to achieve head-direction calibration.

During learning, CAL cell activity depends on PSC cells only, whereas all inputs coming from VIS and sEC cells are inhibited. In particular, each CAL cell c is driven by the corresponding PSC cell p, such that $\theta_c = \theta_p$. The $p \rightarrow c$ connection is fixed and has a synaptic strength $w_{cp} = 1$. The activity r_c of each CAL cell c is given by

$$r_c(t) = \begin{cases} 0 & \exists c^* : r_{c^*}(t) > r_c(t) \\ r_p(t) & otherwise \end{cases} \tag{4}$$

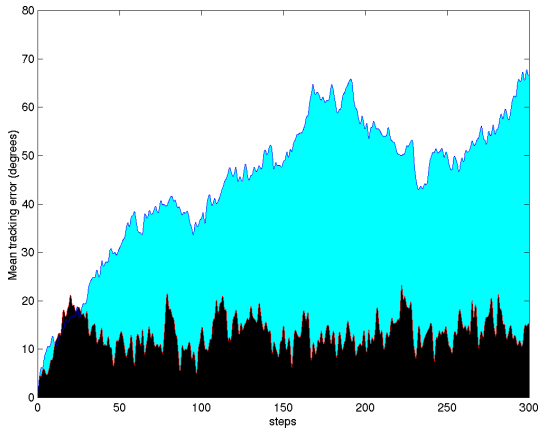

Figure 7: Uncalibrated mean tracking error (grey region) versus calibrated error (black region). The mean is computed over $n = 10$ trials.

Thus, during learning, CAL cells are rigidly coupled with PSC cells and function as directional neurons.

Connections CAL $\rightarrow$ PSC, VIS $\rightarrow$ CAL, and sEC $\rightarrow$ CAL are initialized to zero, and changed on-line by means of one-shot Hebbian learning

$$\Delta w_{ij} = \nu \, r_i \, r_j \, (1 - w_{ij}) \tag{5}$$

where i and j are the post- and presynaptic neuron, respectively, and ν is the learning rate[1]. At time t, learning is triggered only if:

(i) the relation $t - t' \leq \mathcal{T}$ holds, where t' is the last calibration time, and $\mathcal{T}$ is a fixed temporal threshold. This prevents the system from correlating allothetic and idiothetic signals when the error $\theta(t) - \overline{\theta}(t)$ might be large.

(ii) relations $\sigma_v \leq \Sigma_v$ and $\sigma_s \leq \Sigma_s$ hold, where $\sigma_v(t)$ is the VIS activity variance around the center of mass $\mu_v(t)$, and $\sigma_s(t)$ is sEC activity variance around $\mu_s(t)$, respectively. Σ_v and Σ_s are two fixed thresholds. We assume that only when σ_v and σ_s variances fall below a threshold, the informations carried by VIS and sEC are suitable for calibration.

Bearing-place associations $(\alpha_i, \, \mathbf{p}_i)$, which have been established at different learning times t_i, are maintained separate in order to achieve effective calibration. Thus, each CAL cell c has several neuroreceptors nr_i to keep separate those allothetic signal pairs $(\alpha_i, \, \mathbf{p}_i)$ that are temporally distinct.

Calibrating Head-direction Cells. During non-learning, CAL cell activity is determined by external cues via connections VIS $\rightarrow$ CAL and sEC $\rightarrow$ CAL previously created. In particular, the activity r_c of cell c is defined by

$$r_c(t) = \begin{cases} 0 & \exists \, c^* : r_{c^*}(t) > r_c(t) \\ \max_{nr_i} \left(\frac{\sum_j w_{cj}^{nr_i} \cdot r_j(t)}{\sum_j w_{cj}^{nr_i}} \right) & otherwise \end{cases} \tag{6}$$

where j varies over VIS and sEC cells, and $w_{cj}^{nr_i}$ represents an afferent synapse to receptor nr_i.

[1]Results presented in this paper have been obtained by using $\nu = 1$.

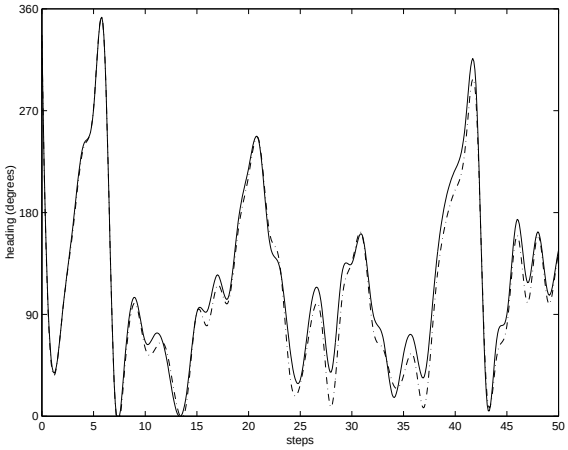

Figure 8: A sample of head-tracking over time. The solid line is the robot's heading θ, whereas the dashed line is the estimated heading $\overline{\theta}$.

CAL cell activity is maximal only when the robot has a specific bearing α relative to L *and* it is at a specific spatial location $\mathbf{p}$. Thus, CAL cells enable the robot to recognize previously learned place-orientation contexts. As a consequence, whenever there exists a maximally active CAL cell (i.e., $\exists c : r_c(t) \geq \epsilon$), shifting the directional cell activity towards CAL cell activity results in calibrating the robot's head-direction system. Calibration is accomplished by means of the learned CAL $\rightarrow$ PSC projections. Indeed, after one-shot Hebbian learning (Eq. 5), synapses encode the PSC ensemble activity as memorized when learning occurred. Thus, if at time t there exists a maximally active cell $c \in$ CAL, PSC activity may be simply calibrated by:

(i) making CAL cell c inhibit all ADN efferents to PSC

(ii) updating PSC activity according to

$$r_p(t) = w_{pc} \cdot \max(r_c(t), \, 1) \qquad \forall p \in PSC \tag{7}$$

Let $\overline{\theta}_c(t)$ and $\overline{\theta}_p(t)$ be the headings encoded by CAL and PSC activity at time t, respectively. Eq. 7 forces a shift in the PSC activity towards CAL activity such that, after calibration, the relation $\overline{\theta}_p(t) = \overline{\theta}_c(t)$ holds.

In order to evaluate the advantages of using external cues to calibrate idiothetic information, we run a series of n experiments all starting at $t = 0$. At each step the robot updates its orientation randomly and the head-direction system is used to estimate its allocentric heading. The above learning scheme is employed to correlate external and internal cues on-line. In each trial i we monitor the deviation between the actual robot's heading $\theta_i(t)$ and the estimation $\overline{\theta}_i(t)$ encoded by the directional system (Eq. 2) over time. The mean tracking error at time t is defined as

$$e(t) = \frac{1}{n} \sum_{i=1}^{n} \mid \theta_i(t) - \overline{\theta}_i(t) \mid \tag{8}$$

The robot turns with a constant angular velocity $\omega \approx 100 \; ^\circ/sec$, and rotates, at each time step, by an angle $\Delta\theta$ drawn from $[-90^\circ, \, 90^\circ]$. For simplicity, the robot turns on

the spot during the whole trial at a location where the necessary learning condition $\sigma_s \leq \Sigma_s$ holds. The experiment includes $n = 10$ sessions, each of which consists of 300 steps.

In Fig. 7 we compare the averaged errors $e(t)$ and $e_c(t)$ corresponding to the non-calibrated and the calibrated system, respectively. The uncalibrated error $e(t)$ grows continuously over time because of its cumulative nature. On the contrary, the calibrated error $e_c(t)$ remains bounded over time.

In Fig. 8 we present a result concerning the head-tracking capability of the system when using allothetic-based calibration. According to the above experimental setup, we let the robot randomly update its orientation during trials of 50 time steps. Solid lines represent the robot's nominal heading θ, whereas dashed lines represent the heading $\overline{\theta}$ as encoded by the PSC population activity. Due to the combination of inertial and visual stimuli, the directional system tracks the robot heading rather effectively over time.

3. Modeling Hippocampal Place Cells

The above head-direction system relies on hippocampal place cell activity (i.e., sEC) to maintain the heading representation consistent over time. Place-coding information is used to map the egocentric bearing α within an allocentric framework.

In this section, we present a neural system for space representation based on a computational model of the rat hippocampus (dark grey areas in Fig. 1). Superficial entorhinal place cells (sEC) are a part of this model and, in particular, provide place identification based on visual stimulation only. The model has been already explained exhaustively in (Arleo & Gerstner, 2000). Thus, here we do not describe the hippocampal place coding system in detail. In this section, we rather outline the whole space representation model. In Sec. 3.1 we focus on the visual pathway enabling the robot to self-localize by extracting spatio-temporal properties of the environment from real visual data. We present new experimental results obtained by replacing the linear visual module[2] used in (Arleo & Gerstner, 2000) by a CCD camera (Fig. 2).

Hippocampal place fields are determined by a combination of environmental cues whose mutual relationships code for the current animal location (O'Keefe & Nadel, 1978). We propose a functional model in which bimodal sensory information is used to establish and maintain a population of place cells (Arleo & Gerstner, 2000). External cues and internal self-generated information are integrated in order to generate a stable space representation encoding robot positions. Ethological and neurophysiological data suggest that allothetic signals (in particular visual cues) exert a strong influence on hippocampal place cell activity (Knierim et al., 1995). On the other hand, experiments on rodents show that place cells can maintain stable receptive fields even in absence of reliable allothetic cues (e.g., in the dark) (Quirk et al., 1990). This suggests that idiothetic information also plays an important role in place representation.

[2] A one-dimensional array of 64 photo-receptors covering 36° on the azimuthal plane.

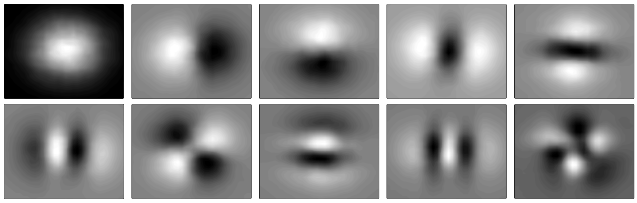

Figure 9: The 10 filters f_i correspond to the first 10 principal components (numbered from left to right, top to bottom) learned by applying Sanger's neural network technique (Sanger, 1989).

The implemented system consists of a multi-layer neural architecture (dark grey part in Fig. 1) modeling high-dimensional continuous sensory input by means of localized overlapping place fields. Starting with no prior knowledge, the system grows incrementally and on-line based on the agent-environment interaction. Since the robot moves in a two-dimensional space, all visual data lie on a low-dimensional manifold within the high-dimensional input space. We apply unsupervised Hebbian learning to detect the low-dimensional view manifold representing the visual input space. In particular, visual stimuli are interpreted by means of neurons that only respond to combinations of specific visual patterns. We call these neurons snapshot cells (SnC) (Fig. 1). The activity of these neurons implicitly represents properties like agent-landmark distance and egocentric orientation to visual cues. In order to achieve spatial discrimination, the activity of these neurons is combined by unsupervised Hebbian learning. This allows the system to create a population of visually driven place cells incrementally (sEC cells in Fig. 1).

Since distinct spatial locations might provide identical visual stimuli, the low-dimensional view manifold might be singular. To remove such singularities we consider idiothetic information along with visual signals. An extra-hippocampal path integrator drives Gaussian-tuned neurons modeling internal movement-related stimuli. We suppose that the idiothetic place representation takes place in the medial layer of the entorhinal cortex (mEC region in Fig. 1) and is environment independent. A fundamental contribution to build the idiothetic space representation in mEC comes from the head-direction system (projection B in Fig. 1).

The spatial representation in the hippocampal CA3-CA1 areas is obtained by integrating the above idiothetic and allothetic representations (Fig. 1). Correlational learning is applied to combine visual cues and path integration. On the one hand, unreliable visual data are compensated for by means of path integration. On the other hand, reliable visual information can calibrate the path integrator system and maintain the odometric error bounded over time (Arleo & Gerstner, 2000).

Finally, to accomplish its functional role in goal-oriented navigation, our hippocampal model must incorporate the knowledge about relationships between the environment, its obstacles and specific target locations. The proposed approach results in an incrementally learned coarse coding representation suitable for applying reinforcement learning for continuous input state spaces. In particular, we apply a Q-

learning scheme to acquire navigational maps. Our place fields work as a basis function approximator that can be utilized to learn a parameterized form of the Q-function (Watkins, 1989). We focus on a specific neural pathway, namely the fornix projection, connecting the hippocampus (in particular the CA1 region) to the nucleus accumbens (NA region in Fig. 1). The latter is an extra-hippocampal structure that seems involved in reward-based goal memory and in locomotor behavior (Redish, 1997). In our model, CA1-CA3 place cells drive a population of action neurons in NA. Synaptic efficacy between CA1-CA3 cells and action cells is modified as a function of dopaminergic target-related reward signals. This results in an ensemble activity of the action neurons that provides a navigational map to support spatial behavior (Arleo & Gerstner, 2000).

3.1 Modeling Visual Input

In this section, we focus on the visual pathway of our hippocampal model. We apply a simple computational strategy to emulate the feature-extraction mechanism observed in the visual cortex (Hubel & Wiesel, 1977). Visual input is provided by a black and white CCD camera having a view field of approximately 90° in the horizontal plane (Fig. 2(a)). Image processing involves histogram equalization and resolution reduction. Thus, working images $\mathbf{I}(x, y)$ have a resolution of $L_x \cdot L_y = 115 \cdot 86$ pixels. In addition, the 256 grey levels of original images are mapped into the $[-1, 1]$ range, that is $-1 \leq \mathbf{I}(x, y) \leq 1$. The aim is to model spatio-temporal relationships between visual cues by means of neural activity.

We take a class P of 10 two-dimensional filters f_i, each of which responds maximally to a specific visual pattern. In order to define the set of basic filters f_i, we take the first 10 principal components extracted from the robot's visual data input. In particular, we apply Sanger's neural network algorithm (Sanger, 1989) to approximate principal components through off-line learning. We let the robot collect 100 images $\mathbf{I}$, and we build a training set by randomly selecting 10000 square samples of $32 \cdot 32$ pixels from these images. Finally, we train a network consisting of $32 \cdot 32 = 1024$ input units x_j and 10 output units y_i according to (Sanger, 1989)

$$\Delta w_{ij} = \eta y_i \Big(x_j - \sum_{k=1}^{i} w_{kj} y_k \Big) \tag{9}$$

After training, each i_{th} weight vector $\mathbf{w}_i$ corresponds to the eigenvector of the input correlation matrix which has the i_{th} maximal eigenvalue. In other words, the receptive field of output unit y_i corresponds to the i_{th} principal component of input data (Sanger, 1989). Fig. 9 shows the 10 first principal components extracted by applying the above learning scheme. These 10 receptive fields form our class P of basic $32 \cdot 32$ filters f_i.

For each filter $f_i \in P$, we consider 5 different scales f_{is}, which endows the system with a distance discrimination property. In particular, if $l_i \cdot l_i$ is the size of the original f_i filter, we take 5 filters of size $l_{is} \cdot l_{is}$ where $l_{is} = l_i \cdot s/100$, with $s \in \{100, 60, 50, 40, 25\}$.

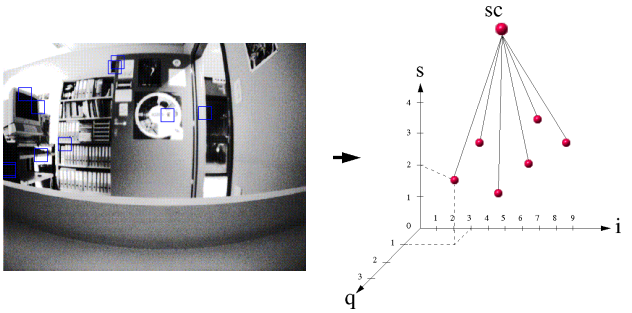

(a) (b)

Figure 10: Visual data are interpreted by mapping images (a) into a filter-activity space (b). Given an image $\mathbf{I}$, the cluster of maximally active filters is used to generate a neural unit sc which is the internal representation of $\mathbf{I}$.

The response a_{is} of a filter f_{is} to an input image $\mathbf{I}$ is computed by shifting the filter across the image to look for the location of maximal matching between f_{is} and $\mathbf{I}$. That is,

$$a_{is} = \max_{x,y} \Big(\sum_{i=0}^{l_{is}-1} \sum_{j=0}^{l_{is}-1} f_{is}(i, j) \cdot \mathbf{I}(i + x,\ j + x) \Big) \tag{10}$$

where $-1 \leq \mathbf{I}(x, y) \leq 1$, $-1 \leq f_{is}(i, j) \leq 1$, $0 \leq x \leq L_x - l_f$, and $0 \leq y \leq L_y - l_f$. In addition, the activity of a filter f_{is} is normalized with respect to its potential maximal response

$$a_{is} = \frac{a_{is}}{\sum_{i=0}^{l_{is}-1} \sum_{j=0}^{l_{is}-1} |f_{is}(i, j)|} \tag{11}$$

Let $0 \leq x^* \leq L_x - l_f$ and $0 \leq y^* \leq L_x - l_f$ identify the location of best match. We segment images in 4 quadrants $\mathbf{I}_q$ and we utilize the spatial information (x^*, y^*) to characterize each filter f_{is} by the region of the image in which it has detected its preferred visual pattern, that is f_{is}^q.

In other words, we interpret visual data by mapping images $\mathbf{I}$ into a three-dimensional filter-activity space $\mathcal{F}$ consisting of $F = 10 \cdot 5 \cdot 4 = 200$ elements f_{is}^q (Fig. 10). Let k be an index over $\mathcal{F}$, that is $0 \leq k \leq F$. Thus, each element f_k identifies a unique neural filter $f_{is}^q \in \mathcal{F}$ responding to a specific localized pattern within $\mathbf{I}$.

In order to detect more complex features, we consider a layer of visual cells one synapse downstream the neural filter layer. We call these neurons snapshot cells (SnC in Fig. 1). The idea is to represent each image by the cluster of filters with the highest activation value, defined by Eq. 10. Given an image $\mathbf{I}$, the set of active filters projects one layer forward to form a snapshot cell $sc \in$ SnC

$$sc = \{ f_k \mid a_k \geq C_k \} \tag{12}$$

where a_k is the response of filter f_k computed according to Eqs. 10 and 11, and C_k is the threshold above which a filter f_k is considered maximally active. The firing activity r_{sc} of cell $sc \in$ SnC is given by

$$r_{sc} = \frac{\sum_{k \in sc} \mathcal{H}(a_k - C_k)}{N_{sc}} \tag{13}$$

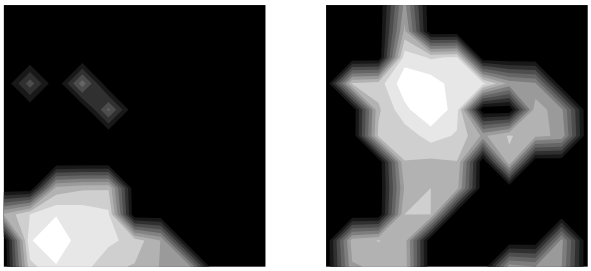

(a) (b)

Figure 11: Two examples of sEC receptive fields. The lighter a region, the higher the firing rate of the cell when the robot is in that region. (a) The visual input is reliable, so that the maximal activity is confined to a localized spot in the environment. (b) The receptive field has multiple peaks indicating that similar visual stimuli occur at different locations.

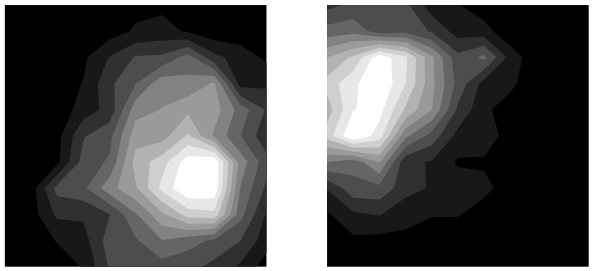

(a) (b)

Figure 12: Two place fields recorded from the CA3-CA1 layer of the model. When the robot is in the region of the white spot the firing rate of the cell is maximal. Notice the Gaussian-like tuning curve, which is compatible with single cell recordings from real place cells (O'Keefe & Nadel, 1978).

where $\sum_{k \in sc}$ sums over all the N_{sc} filters projecting to cell sc, and $\mathcal{H}$ is the Heaviside function. The normalization has been chosen so that $0 \leq r_{sc} \leq 1$.

The robot takes four snapshots corresponding to the north, east, south, and west views at each location visited during exploration (Burgess et al., 1994). To do this, it relies on the allocentric compass information provided by the head-direction system (Sec. 2.). Thus, for each visited location the robot creates four snapshot cells, which are bound together to form a non-directional local view. Given the growing population of snapshot cells, unsupervised Hebbian learning is applied to combine SnC receptive fields and create a population of sEC neurons. Thus, sEC cell activity implicitly depends on a combination of several visual cues, which results in directionally independent place fields.

Figs. 11, 12, and 13 show some results obtained with the experimental setup of Fig. 2. In particular, Fig. 11 shows two receptive fields as recorded from the sEC layer of the model, whereas Fig. 12 represents two typical place fields of CA3-CA1 cells. Finally, Fig. 13 shows a typical population activity of approximately 1000 CA3-CA1 place cells[3] created by the robot through environment exploration. The ensemble activity shown in Fig. 13 has been recorded when the robot was visiting the upper-right corner of the arena.

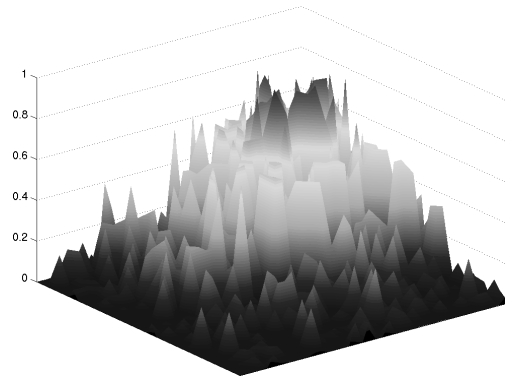

Figure 13: The ensemble activity of approximately 1000 CA3-CA1 place cells created by the robot during environment exploration. The population activity codes for the upper-right corner of the arena.

4. Discussion

In this paper we address the problem of spatial learning and navigation for neuro-mimetic systems. In particular, we model two neural biological systems which are crucial for solving spatial navigation tasks: head-direction cells and hippocampal place cells. We propose a unitary architecture in which directional and place coding systems are strongly coupled, interacting with each other to provide consistent representations. We stress the importance of combining internal and external signals to determine the dynamics of both systems. This allows us to maintain stable representations over time. We present experimental results obtained by implementing the entire system on a mobile robotic platform. The head-direction system is used to track robot rotations, whereas a population of place cells allows the robot to self-localize within its environment. Recording data from our directional and place cell layers are consistent with several neurophysiological findings.

A review of existing hippocampal models for spatial representation (Burgess et al., 1994; Redish, 1997; Gaussier et al., 1997; Trullier & Meyer, 1997; Mallot et al., 1997) has been given in (Arleo & Gerstner, 2000). Here we focus our discussion on the head-direction representation.

Ever since earlier neurophysiological results showed the existence of directional cells in the rat brain, several approaches have been proposed to model the orientation selectivity property of these neurons. (McNaughton et al., 1991) propose the first plausible theory explaining the update mechanism underlying head-direction cells and the influence of extrinsic signals. Nevertheless, they address the problem at a rather abstract level, without accounting for neural connections and dynamics. (Skaggs et al., 1995) propose a refinement of the previous model mainly focusing on the neural architecture of the directional system. They first introduce the idea of attractor dynamics in a head-direction system. However, they do not report any experimental implementation of

[3] This diagram was obtained by discretizing the two-dimensional environmental space in $40 \cdot 40$ subregions, and then plotting the mean firing activity characterizing each subregion.

the model into a concrete directional system. (Zhang, 1996) stresses the central role of the intrinsic dynamics of the head-direction cell ensemble. He reports the first simulation results of a directional system based on a pure attractor scheme. Still, he does not transpose his model into an anatomically plausible framework (i.e., he does not identify the role of distinct areas involved in biological head-direction representation). (Blair, 1996) proposes the first model in which PSC and ATN (anterior thalamic nuclei) form an angular velocity dead-reckoning circuit. In particular, this is the first model stressing the importance of anticipatory cells in ATN and in which the ATN representation leads the PSC representation. In addition, the model fits some firing properties of ATN and PSC biological cells. Nevertheless, the model produces directional tuning curves which are much narrower than those recorded from rodents' head-direction cells. Likewise, (Blair, 1996) simulates tracking performances of the system only on a simple case (one single 60° back and forth rotation), but does not report results on more realistic data. (Redish et al., 1996) further develop the hypothesis of attractor network formulated by (Skaggs et al., 1995) and simulated by (Zhang, 1996). They propose an ATN $\leftrightarrow$ PSC circuit in which both ATN and PSC consist of coupled attractor networks. Simulation results reported by (Redish et al., 1996) show plausible directional tuning curves and effective head-tracking capabilities. However, (Redish et al., 1996) do not report any results concerning calibrating the system by means of extrinsic cues. Also, anatomical lesion data suggest that impairing PSC functionality does not disrupt ATN direction selectivity, whereas (Redish et al., 1996) predict that lesions to PSC cells would disrupt the entire directional system.

Our approach is similar to the above models in that it postulates the primary role of inertial stimuli in determining the directional cell dynamics. Nevertheless, we stress the importance of calibrating the system by incorporating extrinsic signals. This is necessary to track the animat's heading in real sensory-environment contexts. We propose a circuit formed by PSC, ADN, and LMN. ADN and LMN form the angular velocity integrator, whereas PSC is more related to allothetic cues. Lesions to our PSC cells do not disrupt ADN and LMN directional tunings, but make ADN and LMN insensitive to external cues. As (Blair, 1996), we emphasize the anticipatory property of ADN cells. However, we assume that some ADN cells may anticipate head-direction more than others (Blair & Sharp, 1995).

Acknowledgments

Supported by the Swiss National Science Foundation, project nr. 21-49174.96. The authors thank Stéphane Hug for implementing the PCA-based visual interpretation.

APPENDIX A:

Synaptic Connections and Neural Activity in the Angular Velocity Integrator Module of our Head-direction System

HAV Cell Activity. Cells $\Omega_i \in$ HAV are correlated to the sign and magnitude of the angular velocity ω according to

$$r_{\Omega_+} = \begin{cases} \frac{\omega}{\omega_{max}} & \omega > 0 \\ 0 & otw. \end{cases}, \quad r_{\Omega_-} = \begin{cases} \frac{|\omega|}{\omega_{max}} & \omega < 0 \\ 0 & otw. \end{cases}, \quad r_{\Omega_=} = \begin{cases} 1 & \omega = 0 \\ 0 & otw. \end{cases} \tag{14}$$

Thus, $r_{\Omega_+} \in [0,1]$, $r_{\Omega_-} \in [0,1]$, and $r_{\Omega_=} \in \{0,1\}$.

LMN Cell Activity. Each cell $l_i \in$ LMN receives inputs projections from ADN cells and modulates its spiking frequency according to HAV cell activity. Let l and a be LMN and ADN cells, respectively, such that $\theta_l = \theta_a$. Each cell $l_i \in \{ l_z, l_s, l_m, l_f \}$ receives five input connections w_{ij}, one from each cell a_j, where $1 \leq j \leq 5$ (Fig. 14). Connections w_{ij} are fixed and are defined according to

$$w_{ij} = \begin{cases} W_{strong} & j = 1 \\ W_{weak} & 2 \leq j \leq 5 \end{cases} \tag{15}$$

where W_{weak} and W_{strong} are two constant terms such that $W_{strong} \gg W_{weak}$, and $\sum_j w_{ij} = 1$. The activity of a LMN cell l_i is given by

$$r_{l_i}(t) = \begin{cases} \sum_{j=1}^{5} w_{ij} \cdot r_{a_j}(t) & |\omega| = \omega_i \\ 0 & otherwise \end{cases} \tag{16}$$

where $\omega_i \in \{\omega_z, \omega_s, \omega_m, \omega_f\}$. Since a_j cells anticipate the current heading θ by different angles during robot turning, the LMN cell l_i receives a sequential activation of its input. We assume that only when *all* ADN cells a_j have been sequentially activated, cell l_i is enabled to fire according to Eq. 16. Note that, due to Eq. 15, $r_{l_i}(t) \approx r_{a_1}(t)$. That is, cell l_i needs the sequential activation of all anticipatory inputs a_j to be able to fire, but its firing rate is mostly determined by the most recent input a_1.

ADN Cell Activity. Each cell $a_j \in$ ADN is driven by LMN and HAV cells. In the current implementation, all ADN afferent connections have synaptic weight equal to 1. In particular, each ADN cell a_j has three neural receptors nr_+, $nr_=$, nr_-, which correspond to positive, negative, and zero ω, respectively. Each receptor nr_i receives inputs from LMN cells and it is gated by the activity of cell $\omega_i \in$ HAV (Fig. 15(a)).

Synaptic projections LMN$\rightarrow$ADN (Fig. 15(b)) are such that ADN cells a_j may predict future headings during robot turning, and code for the present heading when the robot is not rotating. Let l and a be LMN and ADN cells such that $\theta_l = \theta_a$. Each a_j neuron, with $1 \leq j \leq 5$, receives:

(i) one matching input from cell l_z on its receptor $nr_=$;
(ii) three left-offset inputs on its nr_+ neuroreceptor from cells $(l-j)_s$, $(l-2j)_m$, and $(l-4j)_f$, respectively;
(iii) three right-offset inputs on its nr_- neuroreceptor from cells $(l+j)_s$, $(l+2j)_m$, and $(l+4j)_f$, respectively.

The activity $r_{a_j}(t)$ of a ADN cell is determined by its LMN input. The sign of the angular velocity ω decides which receptor nr_i is currently open, while the magnitude of ω determines which LMN cell is currently firing through nr_i. For instance, if $\omega > 0$ and $|\omega| = \omega_m = 200^\circ/sec$, then cell a_j will listen to its nr_+ input and will fire according to

$$r_{a_j}(t) = r_{(l-2j)_m}(t) \cdot w_{jl} \tag{17}$$

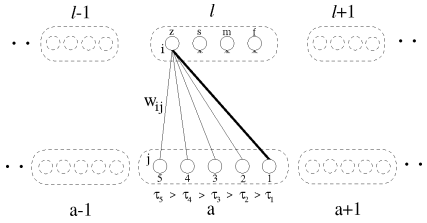

Figure 14: Synaptic projections from ADN to LMN in our model.

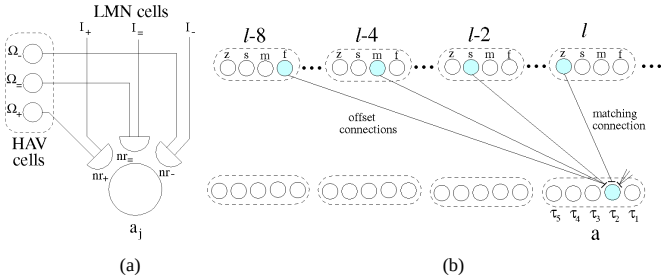

Figure 15: (a) Each ADN cell receives afferents through three distinct receptors nr_+, $nr_=$, nr_-. (b) Synaptic inputs of a ADN cell a_j, with $j = 2$.

where w_{jl} is the connection from cell $(l-2j)_m$ to cell a_j.

According to the above LMN→ADN connecting scheme, the rotational speed of the system depends on the anticipatory behavior of ADN cells and then on the angular velocity ω as discretized by LMN activity. Moreover, since receptors nr_i are gated by the continuous HAV cell activity, the larger $|\omega|$, the more frequently cell a_j will listen to its LMN afferents inputs. Then, the velocity of the shift in ADN activity is proportional to the continuous $|\omega|$.

References

Arleo, A., & Gerstner, W. (2000). Spatial cognition and neuro-mimetic navigation: A model of hippocampal place cell activity. *Biological Cybernetics, (In Press)*.

Blair, H. (1996). A thalamocortical circuit for computing directional heading in the rat. In Touretzky, D., Mozer, M., & Hasselmo, M. (Eds.), *Neural Information Processing Systems 8*, pp. 152–158. MIT Press.

Blair, H., & Sharp, P. (1995). Anticipatory head direction signals in anterior thalamus: Evidence for a thalamocortical circuit that integrates angular head motion to compute head direction. *J. of Neuroscience, 15(9)*, 6260–6270.

Burgess, N., Recce, M., & O'Keefe, J. (1994). A model of hippocampal function. *Neural Networks, 7*, 1065–1081.

Gaussier, P., Joulain, C., Revel, A., Zrehen, S., & Banquet, J. (1997). Building grounded symbols for localization using motivation. In *Fourth European Conf. on Artificial Life*, pp. 299–308.

Georgopoulos, A., Schwartz, A., & Kettner, R. (1986). Neuronal population coding of movement direction. *Science, 233*, 1416–1419.

Hubel, D. H., & Wiesel, T. N. (1977). Functional architecture of macaque monkey visual cortex. *Proc. of the Royal Soc., London, series B, 198*, 1–59.

Knierim, J., Kudrimoti, H., & McNaughton, B. (1995). Place cells, head direction cells, and the learning of landmark stability. *J. of Neuroscience, 15*, 1648–1659.

Leonhard, C., Stackman, R., & Taube, J. (1996). Head direction cells recorded from the lateral mammillary nuclei in rats. *Soc. for Neuroscience Abstr., 22*, 1873.

Mallot, H., Franz, M., Schölkopf, B., & Bülthoff, H. (1997). The view–graph approach to visual navigation and spatial memory. In Gerstner, W., Germond, A., Hasler, M., & Nicoud, J. (Eds.), *Artificial Neural Networks - ICANN'97*, pp. 751–756, Switzerland. Springer Verlag.

McNaughton, B., Chen, L., & Markus, E. (1991). Dead reckoning, landmark learning, and the sense of direction: A neurophysiological and computational hypothesis. *J. of Cognitive Neuroscience, 3*, 190.

O'Keefe, J., & Nadel, L. (1978). *The Hippocampus as a cognitive map*. Clarendon Press, Oxford.

Quirk, G., Muller, R., & Kubie, J. (1990). The firing of hippocampal place cells in the dark depends on the rat's recent experience. *J. of Neuroscience, 10(6)*, 2008–2017.

Redish, A. (1997). *Beyond the cognitive map*. Ph.D. thesis, Dep. of Computer Science, Carnegie Mellon University, Pittsburgh, PA.

Redish, A., Elga, A., & Touretzky, D. (1996). A coupled attractor model of the rodent head direction system. *Network, 7(4)*, 671–685.

Sanger, T. (1989). Optimal unsupervised learning in a single-layer linear feedforward neural network. *Neural Networks, 2*, 459–473.

Skaggs, W., Knierim, J., Kudrimoti, H., & McNaughton, B. (1995). A model of the neural basis of the rat's sense of direction. In Tesauro, G., Touretzky, D., & Leen, T. (Eds.), *Neural Information Processing Systems 7*, pp. 173–180 Cambridge, MA. MIT Press.

Taube, J., Muller, R., & Ranck, J. (1990). Head direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *J. of Neuroscience, 10*, 420–435.

Trullier, O., & Meyer, J. (1997). Place sequence learning for navigation. In Gerstner, W., Germond, A., Hasler, M., & Nicoud, J. (Eds.), *Artificial Neural Networks - ICANN'97*, pp. 757–762, Switzerland. Springer Verlag.

Watkins, C. (1989). *Learning from delayed rewards*. Ph.D. thesis, University of Cambridge, England.

Zhang, K. (1996). Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: A theory. *J. of Neuroscience, 16(6)*, 2112–2126.