

Adaptive sensory processing for efficient place coding

Denis Sheynikhovich*, Ricardo Chavarriaga, Thomas Strösslin, Wulfram Gerstner

Laboratory of Computational Neuroscience, EPFL, CH-1015 Lausanne, Switzerland

Available online 2 February 2006

Abstract

This work presents a neural model of self-localisation implemented on a simulated mobile robot with a realistic visual input. A population of modelled place cells with overlapping receptive fields is constructed online during exploration. In contrast to similar models of place cells, parameters of neurons in the sensory pathway adapt online to the environments statistics in order to maximise information transmission. The robot's position can be decoded from the population activity with high accuracy. The information transmission rate of the cells is comparable to the information rate of biological place cells.

© 2006 Elsevier B.V. All rights reserved.

Keywords: Self-localisation; Information maximisation; Place cells

1. Introduction

A large body of experimental data suggests that rats are able to build a spatial representation of the environment they are located in. Such a representation may reside in spatially tuned neurons (i.e. place cells) found in the rat hippocampus [7]. External sensory input (primarily visual) plays a major role in controlling the formation of such a representation, along with self-motion information [8]. While properties of self-motion signals do not change from one environment to another, external stimuli can vary significantly. In neural models of place cells, the external input is often represented as a response of filters applied to visual images [1], as bearing and distance of landmarks encoded by coarse fields of activations [5] or as distances to walls [6]. In these models, the location sensitivity of modelled hippocampal neurons is achieved by comparing previously stored and currently perceived sensory inputs. This comparison is usually performed by a set of sensory neurons whose activity reflects the similarity between the inputs. In most models so far, parameters of the sensory neurons remain constant during learning, i.e. they must be tuned by the user in order to be valid in environments with potentially different statistics.

However, experimental data suggest that neurons are able to adjust their input/output relation (i.e. transfer function) taking into account the stimuli statistics [4,12]. Moreover, the principle governing such adaptation seems to be the maximisation of information that a neuron transmits [2,4]. Here, we apply this principle in a model of place cells that adapt online to the environment statistics and allow for the decoding of the current location with a small error.

2. Model

Our self-localisation model is implemented and tested on a simulated mobile robot receiving realistic two-dimensional visual input. While the robot explores a novel environment, two populations of artificial neurons are incrementally created. (i) *Sensory cells* store low-level visual features extracted from the input. A simple and fast online algorithm is used to adapt the responses of these neurons to transmit maximum information by taking into account the distribution of the inputs. (ii) *Visual place cells* extract location-dependent information from the sensory activity. Their transfer function depends on a single parameter which is chosen to maximise information rate and location specificity of the cells.

During exploration, a snapshot of the environment is taken by the video camera mounted on the robot (200°

*Corresponding author.

E-mail address: denis.sheynikhovich@epfl.ch (D. Sheynikhovich).

horizontal view field) at each time step. The input image is sampled with a uniform rectangular grid of 96×12 points where at each point of the grid we place a set of 8 two-dimensional Gabor filters with different orientations [1]. The set of responses $\mathcal{I} = \{r_k^G\}_{k=1}^K$ of $K = 8 \times 96 \times 12 = 9216$ filters constitutes a *local view* and serves as an input to the layer of sensory cells (SC).

A new SC is recruited for each local view. The SC receives inputs from all K filters. Synaptic connections between SC i and filter k are set by one-shot Hebbian learning rule as $w_{ik}^{SC} = r_i^{SC} r_k^G$, where $r_i^{SC} = 1$ at creation. Weights w_{ik}^{SC} store the local view \mathcal{I}_i extracted from the image. At this and subsequent time steps the synaptic input to the SC is calculated as a normalised difference $\Delta_i = \|\mathcal{I} - \mathcal{I}_i\|_2$ between the newly observed (\mathcal{I}) and the stored (\mathcal{I}_i) local views (at creation time $\Delta_i = 0$). Before calculating Δ_i , the local views are aligned with an arbitrarily chosen directional frame, and only the responses corresponding to overlapping sectors of the visual field are taken into account. We assume that directional information necessary for such an alignment is available from the head direction system of the rat [10].

We consider each SC as a rate-coding neuron with a non-linear transfer function and firing rate $r_i^{SC} \in [0, 1]$. Such a neuron transmits maximum information if its output firing rate takes all values in the range equally often, i.e. the firing rate distribution is uniform (i.e. maximum entropy distribution) [3]. The transfer function that makes the output distribution uniform corresponds to the integral of the input distribution. An online calculation of this optimal transfer function is implemented at each sensory cell by keeping a history of its inputs and calculating an approximation of the integral for each new input. In particular, we define the firing rate of SC i at time step t to be

$$r_i^{SC}(t) = N^{-1} \sum_{n=1}^N \mathcal{H}(\Delta_i(t-n) - \Delta_i(t)), \quad (1)$$

where $\Delta_i(t)$ is the input to cell i at time step t , N is the memory size and \mathcal{H} is the Heaviside step function.

The activity of a single SC reflects how well a single local view is recognised. In order to combine information from several local views in a *location-sensitive* unit, all simultaneously active SCs are connected to a newly created visual place cell (VPC). A new VPC is created at each time step unless a sufficient number of VPCs are highly active, i.e. the location is already encoded by the VPC population. Connection weights VPC→SC are set by the one-shot Hebbian learning rule as before. A VPC activity is defined by a piecewise linear transfer function with an activity threshold θ_a

$$r_i^{VPC} = \begin{cases} 0 & \text{if } \kappa_i h_i < \theta_a, \\ 1 & \text{if } \kappa_i h_i > 1, \\ \frac{\kappa_i h_i - \theta_a}{1 - \theta_a} & \text{otherwise,} \end{cases} \quad (2)$$

where $h_i = \sum_j w_{ij} r_j^{SC}$ is the input potential of cell i , and $\kappa_i = 1/h_i^0$ determines its saturation potential, with h_i^0 standing for the input potential at the time when the cell was recruited.

The rate at which the activity of a cell transmits information about the locations of the robot (or a rat) can be calculated as

$$I = \int_x \bar{r}(x) \log_2 \frac{\bar{r}(x)}{\bar{r}} p(x) dx, \quad (3)$$

where x is the spatial location, $\bar{r}(x)$ is the mean firing rate at that location, $\bar{r} = \int_x \bar{r}(x) p(x) dx$ is the overall mean firing rate and $p(x)$ is the probability density for the robot being at location x [11]. The information rate I measures how much information in bits a neuron transmits per unit time, whereas the normalised information I/\bar{r} measures the specificity of the cell in bits per unit of firing rate (bits/spike in real neurons). We choose θ_a such that the location sensitivity of the VPC neurons was as high as possible without loss in the rate of information transmitted (see below).

Online adaptation in the SC population ensures that the input to the VPC population is independent from the visual stimuli statistics. A natural parameter adaptation in the VPC population can thus be performed over a much longer time-scale, whereas here we use an offline procedure for this purpose.

3. Simulations and results

The model was tested in three different simulated environments: an open arena (Box 1) of 80×80 cm placed in a room with standard office equipment, a rectangular box (Box 2) of 100×50 cm with high black walls with distinct visual patterns in the corners, and a cylinder (Box 3) of 80 cm in diameter with a panoramic photograph of Swiss mountains as wall texture. We let the robot explore each of the boxes for several thousand time steps, such that about 30 VPCs were highly active ($r^{VPC} > 0.8$) in any location inside the corresponding box. To illustrate the adaptation in the sensory population, we plot normalised histograms of the inputs Δ_i in Box 1 and Box 2 in Fig. 1a. In each sensory cell, its history is a sample of size $N = 100$ from a distribution corresponding to one of the histograms. The adaptation rule Eq. (1) results in a transfer function with a sigmoid-like form whose steepest part corresponds to the area of most frequently occurring differences (marked by squares and triangles for Box 1 and 2, respectively), such that the resulting output distribution is close to uniform and information transmission is maximised [4].

The firing rate of the VPCs (Eq. (2)) depends on the activity threshold θ_a . Values of θ_a close to 1 lead to a high sensitivity of the VPCs to the exact positions where they were recruited. At low values of θ_a the VPCs will be as sensitive to position as the sensory cells ($\bar{r} = 0.5$ due to the uniformity of the output distribution). This is illustrated in

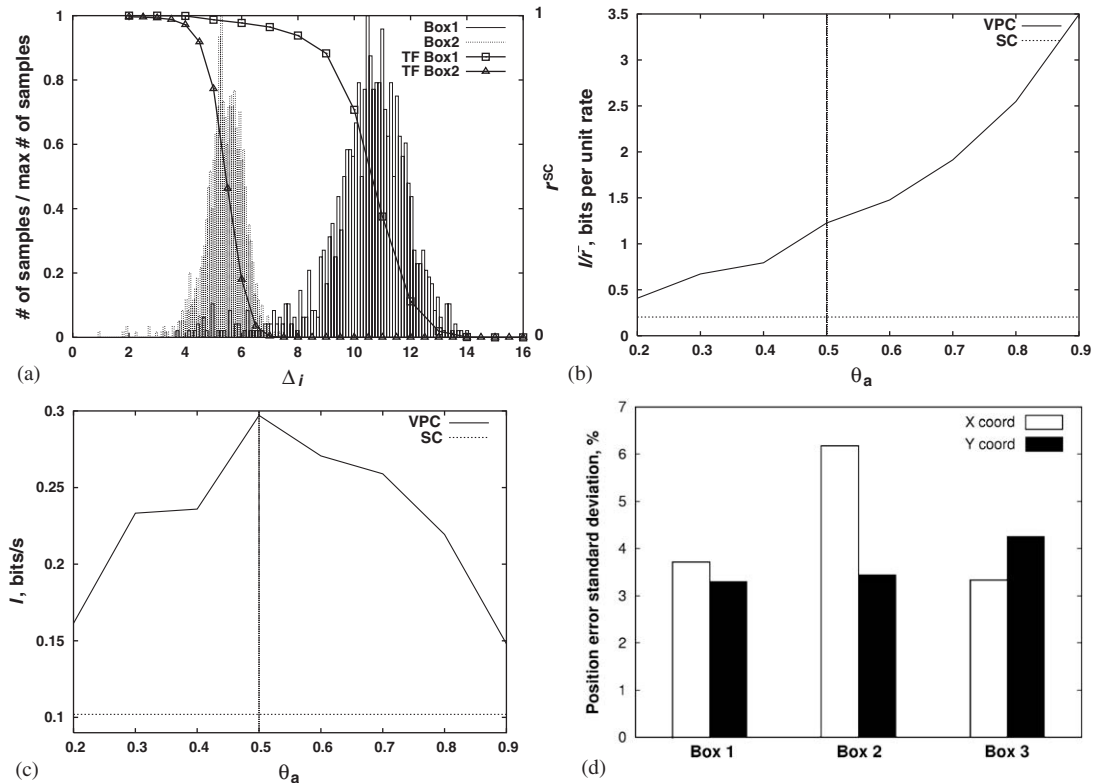


Fig. 1. (a) Normalised histograms of differences Δ_i with optimal transfer functions for Box 1 (squares) and Box 2 (triangles); (b,c) SC and VPC location specificity I/\bar{r} in bits per unit rate (b) and information rate I in bits/s (c) averaged over 50 cells versus θ_a . The vertical lines on both panels correspond to a θ_a that maximises I ; (d) The standard deviations of position estimation error along horizontal and vertical directions in percents of the length of the environment along that direction. For the cylinder (Box 3) only the locations close to the centre were analysed.

Fig. 1b where we plot the dependency of cell specificity I/\bar{r} on the values of θ_a . In contrast, the rate of position information I (Eq. (3), Fig. 1c) has a maximum at $\theta_a = 0.5$. Intuitively it means that if a cell's spatial receptive field is very small, the activity of this cell is not informative most of the time. On the other hand, if the specificity of the cell is too low, the cell is active in a large portion of the environment making the cell non-informative again.

To calculate the gain in the information rate of the place cells compared to the sensory cells, we calculated I and I/\bar{r} for the sensory cells using Eq. (3). The result ($I = 0.102$ and $I/\bar{r} = 0.204$) suggests that the combination of information from several sensory cells into one place cell indeed leads to an increase of both information rate and location specificity (at the optimal value of $\theta_a = 0.5$ in the VPC population $I = 0.297$ and $I/\bar{r} = 1.227$, Fig. 1b, c). As reported in [11,9] the information per spike of CA1 place cells (i.e. I/\bar{r}) is around 1.5–2.5 bits which is slightly more than the information per unit rate in the model. Higher information rate in biological place cells may be due to the combination of different sensory modalities in addition to vision [9].

In order to estimate the accuracy of self-localisation, we let the robot move in random directions for another 1000 time steps (with fixed population sizes). At each time step the error between actual position and its estimation encoded by the VPC population was calculated. Fig. 1d shows the standard deviations of the error for the three test

environments. In spite of the fact that the environments had such different statistics, as illustrated in Fig. 1a, the SD of localisation error is limited to 7% of the length of the environment. Exactly the same model was used in all the three environments.

4. Conclusions

We presented a simple biologically plausible model of self-localisation which is able to build an accurate spatial representation in statistically different environments. A slightly different approach to the problem of sensory adaptation could be to assume a particular type of the transfer function, e.g. a sigmoid, and use an online algorithm to learn the optimal parameters [3]. Results of simulations assuming such a parameterised transfer function were identical to those reported here, but required a careful choice of initial parameter values (data not shown). The method presented here is not exclusive to visual input and can be applied to other types of sensory processing.

References

- [1] A. Arleo, F. Smeraldi, W. Gerstner, Cognitive navigation based on nonuniform gabor space sampling, unsupervised growing networks, and reinforcement learning, *IEEE Trans. Neural Network* 15 (3) (2004) 639–652.

- [2] J.J. Atick, A.N. Redlich, Towards a theory of early visual processing, *Neural Comput.* 2 (3) (1990) 308–320.
- [3] A.J. Bell, T.J. Sejnowski, An information-maximization approach to blind separation and blind deconvolution, *Neural Comput.* 7 (1995) 1129–1159.
- [4] N. Brenner, W. Bialek, R. de Ruyter van Steveninck, Adaptive rescaling maximizes information transmission, *Neuron* 26 (2000) 695–702.
- [5] A. Guazzelli, M. Bota, M. Arbib, Competitive hebbian learning and the hippocampal place cell system, *Hippocampus* 11 (2001) 216–239.
- [6] J. O'Keefe, N. Burgess, Geometric determinants of the place fields of hippocampal neurons, *Nature* 381 (1996) 425–428.
- [7] J. O'Keefe, L. Nadel, *The Hippocampus as a Cognitive Map*, Clarendon Press, Oxford, 1978.
- [8] A.D. Redish, *Beyond the Cognitive Map, From Place Cells to Episodic Memory*, MIT Press-Bradford Books, London, 1999.
- [9] E. Save, L. Nerad, B. Poucet, Contribution of multiple sensory information to place field stability in hippocampal place cells, *Hippocampus* 10 (2000) 64–76.
- [10] D. Sheynikhovich, R. Chavarriaga, T. Strösslin, W. Gerstner, Biomimetic neural learning for intelligent robots, Chapter: Spatial Representation and Navigation in a Bio-inspired Robot, Springer, Berlin, 2005, pp. 245–264.
- [11] W.E. Skaggs, B.L. McNaughton, K.M. Gothard, E.J. Markus, An information-theoretic approach to deciphering the Hippocampal code, in: *Advances in Neural Information Processing Systems*, vol. 5, Morgan Kaufmann, San Francisco, CA, USA, 1992, pp. 1030–1037.
- [12] S.M. Smirnakis, M.J. Berry, D.K. Warland, W. Bialek, M. Meister, Adaptation of retinal processing to image contrast and spatial scale, *Nature* 386 (6620) (1997) 69–73.



Ricardo Chavarriaga has received his M.Sc. in Electronics engineering from the Pontificia Universidad Javeriana, Cali, Colombia, in 1998. Since 2001 he is a Ph.D. student at the Laboratory of Computational Neuroscience, Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland. His research is focused in the neurological basis of spatial learning and navigation.



Thomas Strösslin received the M.Sc. in Electrical Engineering from ETH Zürich, Switzerland, in 1999 and the Ph.D. in Computational Neuroscience from EPFL, Lausanne, Switzerland, in 2004. He joined the Okinawa Institute of Science and Technology (OIST), Japan, in 2005. He is interested in the neural bases of spatial memory and navigation, the role of neuromodulators in learning and in ways to incorporate these principles in artificial agents.



Wulfram Gerstner received his Ph.D. degree in theoretical physics from the TU Munich, Germany, in 1993, after studies in Tübingen, Berkeley, and Munich. He is a Professor and Head of the Laboratory of Computational Neuroscience, EPFL, Switzerland.



Denis Sheynikhovich has received his M.Sc. in Mathematics from the Institute of Fine Mechanics and Optics, Saint-Petersburg, Russia, in 1997. He is currently working on his Ph.D. at the Laboratory of Computational Neuroscience headed by Prof. Wulfram Gerstner, Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland. His research is focused on computational modelling of learning processes in hippocampal neural networks and their relation to behaviour.