Perceptual Homing by an Autonomous Mobile Robot using Sparse Self-Organizing Sensory-Motor Maps*

Rajesh P.N. Rao and Olac Fuentes
Department of Computer Science
University of Rochester
Rochester, NY 14627

Abstract

We present a biologically-motivated approach to the problem of perception-based homing by an autonomous mobile robot. A three-layered self-organizing network is used to autonomously learn the desired mapping from perceptions to actions. The network, which bears some similarities to the structure of the mammalian cerebellum, is initially trained by teleoperating the robot on a small number of homing paths within a given domain of interest. During training, the connections between input sensory layer and the hidden layer as well as those between the hidden layer and the motor output layer are modified according to the well-known competitive Hebbian learning rule. By employing a population averaging scheme for computing output motor vectors, the robot can subsequently home from arbitrary locations within the domain based solely on current perceptions. We describe preliminary results based on simulation for an actual mobile robot, equipped with simple photoreceptors and infrared receivers, navigating within an enclosed obstacle-ridden arena.

1 Introduction

A central problem in mobile robotics is that of autonomous goal-directed navigation in unstructured environments. Traditional methods for designing navigational controllers involve prewiring a fixed set of strategies based on heuristics and domain knowledge. Such systems however suffer from the inherent inflexibility of utilizing predefined behaviors and as such are unable to cope with the variations that are characteristic of unstructured environments.

Recent work in behavioral robotics has shown that in many instances, the uncertainties posed by unstructured environments can be circumvented to a large extent by endowing the robot with the ability to *autonomously learn* navigational behaviors (for example, [3] and [6]). In the spirit of this recent trend, we present a biologically-motivated framework for the autonomous acquisition of perception-based homing behaviors in mobile robots. Homing can be defined as the ability of an autonomous agent to navigate to a particular "home" location from arbitrary locations within a specific environment. Homing behaviors are almost universal in animals [10]. Assuming that complex animal behaviors emerged from pre-existing simpler ones, it is reasonable to assume that the acquisition of the homing behavior represents a significant step towards acquiring more complex navigational behaviors. Indeed, the general problem of learning to navigate between a given number of arbitrary locations can be decomposed into the simpler components of navigating between one-one, many-one, and one-many locations as shown in Figure 1 (a).

In this paper, we propose a three-layer network architecture that allows a robot to autonomously learn to home based only on its current perceptions. The network, which bears some similarities to the structure of the cerebellum, employs *competitive Hebbian learning* to modify the connections between the input sensory layer and the hidden layer as well as the connections between the hidden layer and the motor output layer during an initial training period which involves teleoperation of the robot (Figure 1 (b)) in an enclosed arena (Figure 1 (c)). The robot is equipped with four orthogonally-placed infrared detectors (for measuring the strength of the modulated infrared light being emitted by a beacon placed at an arbitrary location near the arena), six horizontally-positioned photoreceptors (for measuring the amount of light from a light source located near the arena), and six tilted photoreceptors (for measuring light intensity value due to the color of the floor and surrounding obstacles).

*This work is supported by NSF research grant no. CDA-8822724, ARPA grant no. MDA972-92-J-1012 and ONR no. N00014-93-1-0221.

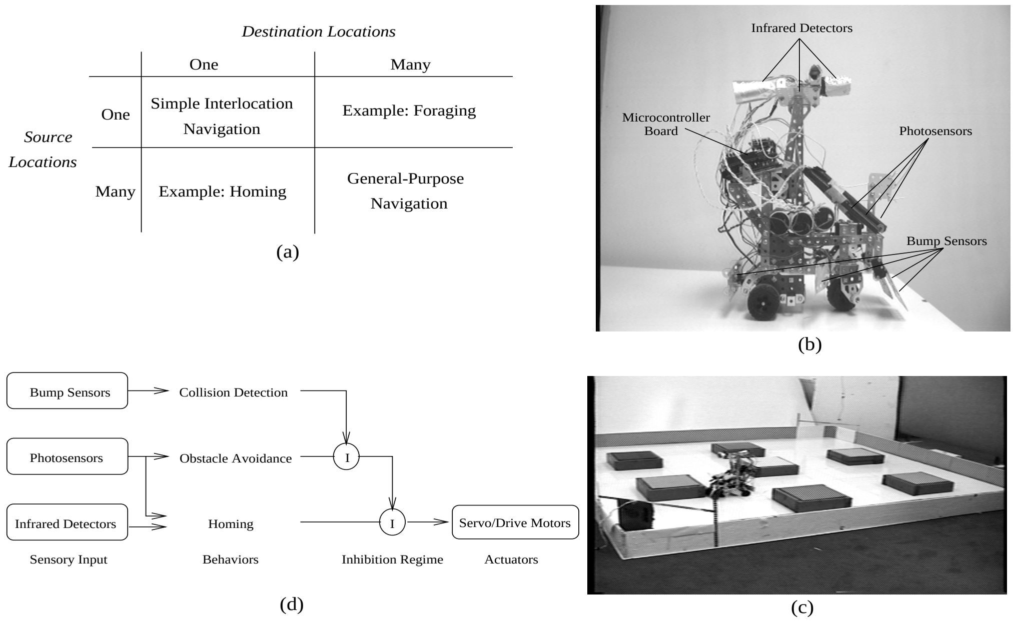

Figure 1: (a) A hierarchical decomposition of the general navigation problem. (b) The mobile robot (c) The robot in its environment. (c) Robot control architecture: the collision detection and obstacle avoidance routines were autonomously learned as described in [2].

2 The Approach

One relatively straightforward method for perceptual homing is to use a look-up table or associative memory to store a large set of perception vectors from different locations in conjunction with corresponding set of actions required to reach home [9]. This approach however has at least three drawbacks: (a) the address space formed by the perception vectors is usually quite large (in our case, 256^{16}) and therefore, storing fixed reference vectors for every possible scenario becomes infeasible; (b) the training time increases drastically as the look-up table size increases, and (c) simple look-up table strategies usually fail to generate the appropriate responses to novel situations. Our approach addresses problems (a) and (b) by using only a *sparse* subset of the perceptual address space. Generalization in novel scenarios is achieved by *self-organizing* both the prototype perception vectors and motor vectors in response to incoming training vectors, and using population averaging to generate motor responses during the homing phase.

2.1 Sparse Sensory-Motor Maps

Figure 2 shows a diagram of the three-layer network architecture used. The first layer consists of n units representing the input sensory vector $\vec{s}$. The hidden layer consists of m units while the motor output is represented by k units. Let $\vec{w}_i$ ($1 \leq i \leq m$) represent the vector of weights between hidden unit i and the input layer, and let $\vec{m}_i$ represent the vector of weights between hidden unit i and the output layer.

During training, the robot is guided home from a small number of arbitrary locations within the arena with the help of the teleoperation joystick. The weight vectors $\vec{w}_i$ are modified in accordance with the competitive Hebbian learning rule:

1. Calculate the distances $d_i = ||\vec{w}_i - \vec{s}||$ between $\vec{w}_i$ and the current sensory vector $\vec{s}$.

2. Select the hidden units $N(i^*)$ representing the $f(t)$ least distances among all hidden units where t represents the number of iterations.

3. For $i \in N(i^*)$, modify weight vectors according to:

$$\vec{w}_i := \vec{w}_i + \gamma(d_i, t)(\vec{s} - \vec{w}_i) \tag{1}$$

   where the gain function γ is given by the Gaussian $\gamma(x,t) = e^{-x^2/2\sigma(t)}$

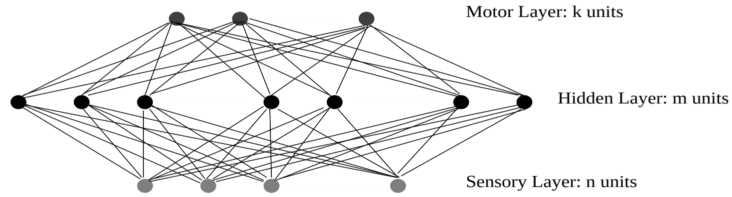

Figure 2: The Network Architecture.

4. Repeat the above steps until home is reached or teleoperation is terminated.

The corresponding motor input vectors $\vec{m}$ (for example, steering angle and speed) are used to modify the weights $\vec{m}_i$ for $i \in N(i^*)$:

$$\vec{m}_i := \vec{m}_i + \beta(t)(\vec{m} - \vec{m}_i) \tag{2}$$

In the above, $f(t)$ and $\sigma(t)$ as well as the gain term $\beta(t)$ $(0 < \beta(t) < 1)$ are gradually decreased with increasing t.

2.2 Population Averaging of Motor Output Vectors

After the initial training phase, the network generates motor responses on the basis of current perceptions using a *population averaging scheme*:

1. Calculate the distances d_i between $\vec{w}_i$ and the current sensory vector $\vec{s}$ and select the hidden units $N(i^*)$ as above with f, for instance, defined to be a constant function.

2. Let m_{ij} $(1 \leq j \leq k)$ denote the weight from motor output unit j to hidden unit i. Then, the output of motor unit j is given by:

$$o_j = \sum_{i \in N(i^*)} e^{-d_i^2/2\sigma} m_{ij} \tag{3}$$

   where σ defines the spread of the weighted averaging.

By averaging the output, considerable generalization to novel situations is achieved due to the inherent interpolation evident in such a step. The above averaging method is a form of interpolation using *radial basis functions* [11] and is inspired by recent neurophysiological evidence [8] that the *superior colliculus*, a multilayered neuronal structure in the brain stem that is known to play a crucial role in the generation of saccadic eye movements, in fact employs a population averaging scheme to compute saccadic motor vectors.[1]

2.3 Relation to the Structure of the Cerebellum

While the network architecture presented above was not designed with the explicit goal of modeling any particular neurobiological structure, we briefly note in this section some striking similarities to the organization of the cerebellum (and therefore to the cerebellar model of Marr [7], the CMAC of Albus [1] and Kanerva's sparse distributed memory [5]). The cerebellum is a relatively large neuronal complex found in the mammalian brain and responsible for motor control and the coordination of muscular movements in response to sensory stimuli [4]. Sensory inputs to the cerebellum are provided by the mossy fibers (from the cerebral cortex) which synapse on the granule cells. Thus, the mossy fibers represent the input vector $\vec{s}$ while the granule cells correspond to the hidden layer of Figure 2. The granule cell axons, called parallel fibers, synapse with Purkinje cells which form the output layer of the cerebellum. A second set of inputs is provided by the climbing fibers (originating in the brain stem) which are paired one-for-one with the Purkinje cells. The climbing fibers can thus provide the initial motor training input $\vec{m}$ which are used to modify the synapses between the hidden layer (granule cells/parallel fibers) and the output layer (Purkinje cells). Three additional types of cells found in the cerebellum include the basket cells, the stellate cells, and the Golgi cells. The basket cells receive input from the climbing fibers and synapse inhibitorily on the Purkinje cells; their function could thus be to prevent modification of synapses when training has been suspended. The Golgi cells receive input from the parallel fibers and synapse on the granule cell bodies; they are thus in a position to select a given number

[1] We refer the interested reader to an earlier paper [12] for a related method for visuomotor learning of saccades for a robot head.

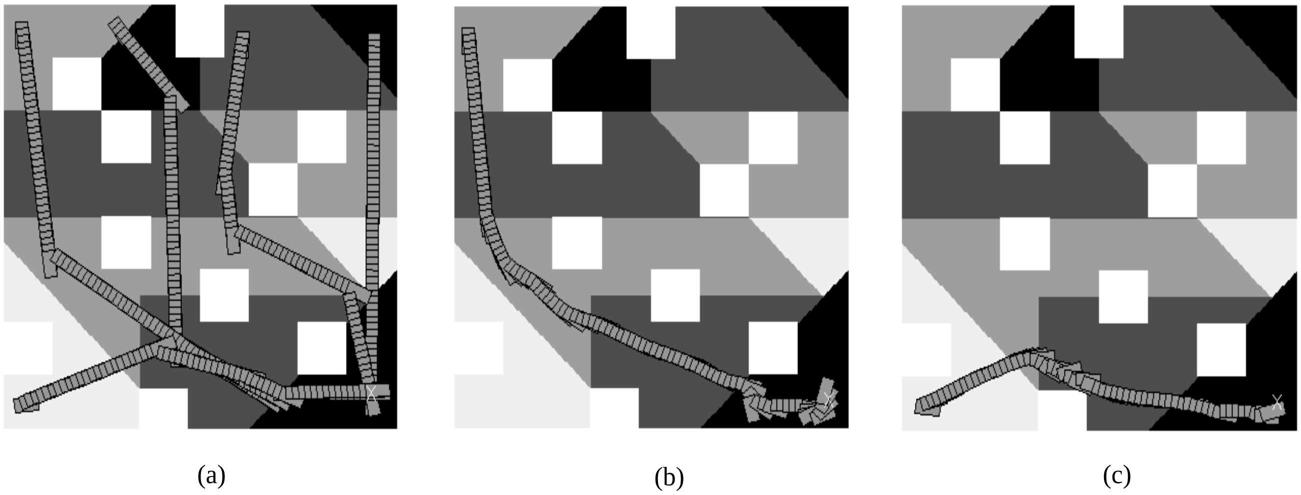

Figure 3: (a) The paths on which the network was trained by teleoperating the robot. Floor color is depicted by shades of grey and obstacles are colored white. The home position is marked by an 'X'. (b) and (c) depict the paths chosen by the robot when placed at two different positions within the arena. Note the slight deviations from the training paths caused by mild perceptual aliasing; the obstacle avoidance behavior is automatically invoked when the deviations cause encounters with the wall or the square obstacles.

of winning hidden units (similar to the number $f(t)$) amongst all hidden units. Finally, the stellate cells are believed to play a role in adjusting the Purkinje cell thresholds; they can therefore mediate the weighted averaging described above.

3 Experimental Results and Conclusions

The network was tested by simulating the robot in its environment. Robot perceptions were computed using the simplifying assumption that light/infrared intensity at a particular orientation θ and at a distance r from the source is proportional to the solid angle subtended by the source (i.e. $cos(\theta)/r^2$), assuming unit area for the robot receptors. Figure 3 illustrates some typical results obtained.

Current work involves making suitable modifications to the method (such as sampling perceptions only between specified intervals) in order to allow close to real-time implementations on our robot. A particularly difficult problem that can be expected to be more severe in a real environment is that of *perceptual aliasing*: perceptions from two different locations may sometimes appear indistinguishable to the robot. Perceptual aliasing can be reduced by augmenting the sensory vector with history information, but this increases the dimensionality of the sensory space. One attractive method to retain history information while keeping the dimensionality manageable is to use *eigenperceptions* rather than the raw perceptions themselves. Incoming perceptions are first pre-processed by projecting them along the principal eigenvector directions and the coefficients thus obtained are used for modifying the network weights. Such a scheme both reduces the dimensionality of the data, thereby maintaining feasibility, while at the same time preserving the major variations in sensory data. Ongoing work involves the integration of eigenperceptions with the method proposed in the paper.

Acknowledgments

We would like to thank Professor Chris Brown for his encouragement and support for building the mobile robot, and Mike Van Wie for his help in the construction of the robot.

References

[1] James S. Albus. A theory of cerebellar functions. *Math. Biosci.*, 10:25–61, 1971.

[2] Olac Fuentes, Rajesh P. N. Rao, and Michael Van Wie. Hierarchical learning of reactive behaviors in an autonomous mobile robot using stochastic hillclimbing. Submitted, January 1995.

[3] D. Gachet, M.A. Salchis, L. Moreno, and J. R. Pimentel. Learning emergent tasks for an autonomous mobile robot. In *IEEE/RSJ International Conference on Intelligent Robots an Systems*, pages 290–297, 1994.

[4] M. Ito. *The Cerebellum and Neuronal Control.* New York: Raven Press, 1984.

[5] Pentti Kanerva. Sparse distributed memory and related models. In Mohamad H. Hassoun, editor, *Associative Neural Memories*, pages 50–76. New York: Oxford University Press, 1993.

[6] Ben J. A. Krose and Marc Eecen. A self-organizing representation of sensor space for mobile robot navigation. In *IEEE/RSJ International Conference on Intelligent Robots an Systems*, pages 9–14, 1994.

[7] David Marr. A theory of cerebellar cortex. *J. Physiol. (London)*, 202:437–470, 1969.

[8] James T. McIlwain. Distributed spatial coding in the superior colliculus: A review. *Visual Neuroscience*, 6:3–13, 1991.

[9] Randal. C. Nelson. Visual homing using an associative memory. *Biological Cybernetics*, 65:281–291, 1991.

[10] F. Papi. *Animal homing.* New York : Chapman and Hall, 1992.

[11] T. Poggio and F. Girosi. Networks for approximation and learning. *Proc. IEEE*, 78:1481–1497, 1990.

[12] Rajesh P.N. Rao and Dana H. Ballard. Learning saccadic eye movements using multiscale spatial filters. In G. Tesauro, D.S. Touretzky, and T.K. Leen, editors, *Advances in Neural Information Processing Systems 7.* Cambridge, MA: MIT Press, 1995. (To appear).