Building Complete Autonomous Agents: A Case Study on Categorization

D. Lambrinos

AILab
Department of Computer Science
University of Zurich
Winterthurerstrasse 190
CH-8057 Zurich, Switzerland
email: lambri@ifi.unizh.ch

C. Scheier

AILab
Department of Computer Science
University of Zurich
Winterthurerstrasse 190
CH-8057 Zurich, Switzerland
email: scheier@ifi.unizh.ch

Abstract

With the advent of New Artificial Intelligence it has become clear that in order to understand intelligence it is important to build complete agents. A complete agent is capable of behaving autonomously in an environment without human intermediary. It has to incorporate, among other things, capabilities for categorization, for navigation, and for "deciding" what to do. In this paper we illustrate how such an integration can be implemented on a mobile robot. We introduce a new control architecture, the so-called Extended Braitenberg Architecture (EBA). It consists of loosely coupled processes that run in parallel. They implement various behaviors such as exploration or recharging. The task of the robot is to collect certain types of objects in the environment. We show how the category learning can be embedded in the overall architecture, i.e. we illustrate how categorization can be viewed from a complete agent's perspective. Results on the behavioral performance as well as the underlying internal dynamics are presented.

1 Introduction

In 1961 the Japanese psychologist Masanao Toda proposed to study so-called "Fungus Eaters" as an alternative to the traditional ways of academic psychology ([13]). Rather than performing ever more restricted and well-controlled experiments on isolated faculties (e.g. categorization, learning or memory) and narrow tasks (e.g memorizing non-sense syllables) one should study *complete* systems. "Complete" in this context means that the agent has to be capable of behaving autonomously in an environment without human intermediary. It has to incorporate, among other things, capabilities for categorization, for navigation, and for "deciding" what to do. Toda's hope was that such a system which is capable of interacting with it's environment autonomously will lead to better insights into the nature of intelligence than looking at isolated fragments of the very complex human mind (for a more detailed discussion of "Fungus Eaters", see [10]).

A very similar view is adopted by "New Artificial Intelligence" where intelligence is studied using autonomous agents (i.e. mobile robots) that interact autonomously with their environment (see e.g. [10] for an overview). One of the basic competences these agents have to be equipped with is the ability to make distinctions, i.e. the ability to *categorize*. Traditionally, categorization has been treated as an information processing module: the sensors receive a particular input which is processed and mapped onto an internal representation (e.g. a category node). This view dominates most psychological models of categorization (see e.g.[6]). Similarly, in computer vision systems categorization is seen as a problem of matching the visual input to a stored representation or model of objects (see e.g. [5] for an overview).

In our previous work we have developed an alternative approach to categorization adopting the New AI framework ([8],[11],[12]). The main idea is to view categorization as a *sensory-motor coordination* rather than an isolated perceptual (sub-)system. This is achieved by including the robot's own actions into the classification process. In this paper we considerably extend this framework. First, we introduce a new control architecture, the so-called *Extended Braitenberg Architecture* (EBA), in which the categorization can be easily embedded. Second, the sensory-motor complexity of the robot is significantly increased. While the sensory-motor system in the previous experiments consisted of IR sensors and two wheels the robot used in this paper is additionally equipped with a CCD camera and an arm with a gripper mounted at the end. A third improvement concerns the categorization mechanisms used. Previously, the basic categorization mechanism was a conditioned association between learned sensory-motor mappings and some behaviors. In the experiments presented in this paper categorization is achieved by a learned *reentrant mapping* between visual and haptic feature maps. The term *reentry* refers to the fact that there are reciprocal connections between the feature maps (e.g. [3]). Reentry is necessary to correlate perceptual informa-

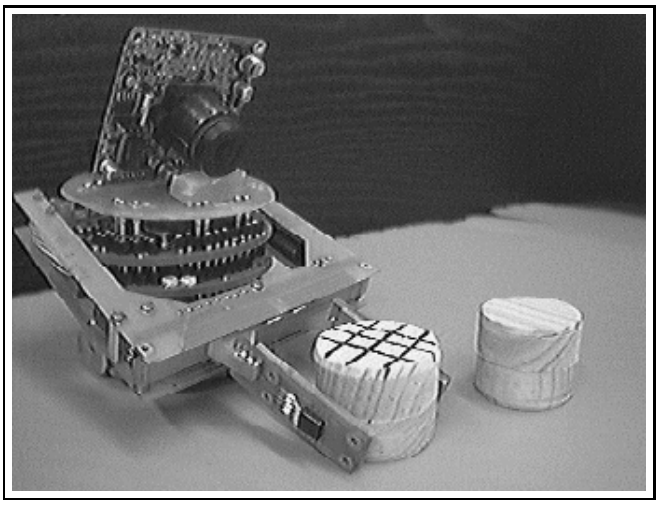

Figure 1: The robot and its environment. Explanations see text.

tions from different modalities on the basis of their temporal contiguity. A final extension of our previous work is the concept of *attentional sensory-motor loops* which are modulated by the category-specific responses of the reentrantly connected feature maps. In essense, categorization consists of breaking or enhancing the attentional sensory-motor loops depending on the type of object encountered and the resulting activity in the feature maps. The *task* of the robot is to collect certain types of objects and bring them to a home base. At the same time it has to sustain itself by regularly visiting the charging station.

2 Experimental Set-up

The mobile robot used in the experiments is is a $Khepera^{TM}$, 55mm in diameter and 32mm high (weight 70g) (see figure 1). The effector system consists of two wheels which are individually driven by DC motors, and an arm with a gripper installed at the end. The maximum object size that can be grasped with the gripper is about 40mm. There are two types of objects in the environment (see figure 1). Wooden objects with texture on the surface and others without texture. Textured objects have been made conductive by rapping a metalic wire around them. The sensory system consists of a visual and a haptic system. Input to the *visual system* is provided by a miniature (1/3", 18g) B/W CCD camera (ces VPC-465). The camera is equipped with a built-in lense with 90 degrees viewing angle. In order to reduce computational load all visual processing was done on a visual server based on a Pentium 133Mhz PC. The video image is sampled at a maximum of 30 fps using a video framegrabber (PMS from Media Vision). Images were grabbed at resolution of 640 x 480 pixels and were spatially averaged to provide an image of 160 x 120 pixels. The main part of the control architecture was run on a workstation (SUN SPARC 10). Communication between the visual server and the workstation was done using the TCP/IP communication protocol. The input to the *haptic system* is as follows. Two wheel encoders provide position information with a resolution of 0.08mm. Arm and gripper position are sensed by position sensors coupled with the respective motors. The arm position sensor takes values from 0 (bottom back) and 255 (bottom forward), the gripper positions sensor takes values from 0 (open) to 255 (closed). Conductivity of objects can be read by a conductivity sensor which takes values from 0 (nonconductive, e.g. plastic, wood) to 255 (conductive, e.g. metal). Objects inside the gripper can be detected by an optical barrier that is mounted on the gripper. The optical barrier takes values from 0 (no object) to 255 (object presence). Finally, there are eight IR sensors, six in the front and two in the back with an angular resolution of 60 degrees each. The maximal distance which the IR sensors can detect is around 40mm. Because of the small distance that they can sense and their arrangement around to body surface of the robot we use them as skin (pressure) sensors.

3 Architecture

3.1 The Extended Braitenberg Architecture

In the real world, agents always have to do several things and at least some of them will be incompatible. To decide what the agent should do in a particular situation is one of the important functions of a control architecture. In the literature this problem has been called "action selection". The term is inappropriate because it introduces a particular bias on how the issue should be tackled (see below). A comprehensive review of "action selection" is given in [14]. There is a fundamental problem with most approaches. They explicitly or implicitly rely on the assumption that what is expressed behaviorally has an internal correspondence. If we see an agent following an object we suspect an object-following module (which is sometimes called a "behavioral layer"). But there is a frame-of-reference issue here. Behavior is by definition the result of a system-environment interaction and the internal mechanisms cannot be directly inferred from the behavior alone. Of course, there must be something within the organism which leads to this behavior. But it would be a mistake to infer that, if we want an agent to follow objects, we have to define a special module for object-following. But this is precisely what is often done. [1] proposed an alternative scheme which does not suffer from this problem. Instead of having modules (or behavioral layers) there are a number of simple processes which all run in parallel and continuously influence the agent's internal state. An extension of this approach is the "Extended Braitenberg Architecture" (EBA). Let us illustrate the point directly with our case study. The basic architecture is shown in Figure 2. There are a number of processes functioning in parallel. The processes are implemented as neural networks. Each process receives weighted input from sensors and effectors (proprioception) and from other processes. They all write simultaneously onto the effector variables where they are summed by a particular

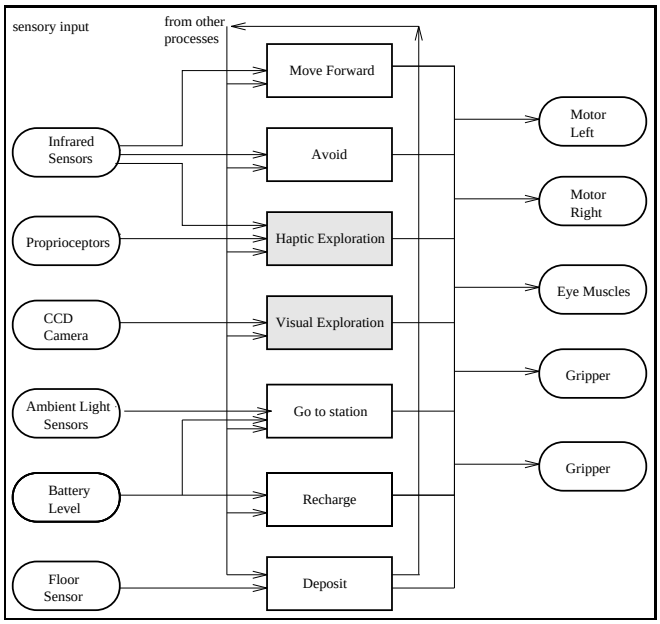

Figure 2: The control architecture of the robot: The shaded processes (haptic and visual exploration) are explained in more detail below.

summation scheme. The resulting values of these effector variables determine the behavior of the agent. For example, the motor speeds are calculated as:

$$\mathbf{s}(t) = (s_l(t), s_r(t)) = \left(\sum_{i=1}^{N} o_i^l(t), \sum_{i=1}^{N} o_i^r(t) \right) \quad (1)$$

where s_l, s_r is the speed of the left and the right motor and o_i^l, o_i^r is the output of the i-th process to the speed quantity of the left and the right motor, respectively and N is the number of processes. There are the following networks: *move forward*, *avoid*, *haptic exploration*, *visual exploration*, *go-to-station*, *recharge*, and *deposit*. This sounds very much like the traditional approaches. The main difference is the following. All the networks run all the time. The influence they exert on the behavior of the agent varies depending on the circumstances. So, under certain conditions they will have no influence and, in others, they will constitute the major influence; but they are not on or off. We give only a short description of each process in Figure 2 (for technical details see [7]. Because the *haptic exploration* and the *visual exploration* network are core aspects of the current architecture we treat them separately below.

move forward: The activity of this process is inversely proportional to the activation of the two front IR sensors.

avoid: The activity of this process is a weighted sum of the activity of all IR sensors such that obstacles on the right will make the robot to turn left and thus avoid the obstacle (and *vice versa*).

go to station: Whenever the battery level is low the robot should go to the charging station and recharge. The motivation to go to the charging station, $M^{(GC)}(t)$, is a convolution of the motivation function associated with the battery level, $M_E(t)$, the distance from the charging station, $M_D(t)$, and the presence of an object in the gripper, $M_O(t)$, respectively (see below). The robot finds the charging station by using the light sensors, i.e. by performing phototaxis behavior :

$$\overline{GC(t)} = M^{(GC)}(GC_l(t), GC_r(t)) \quad (2)$$

$$M^{(GC)} = (M_E, M_D, M_O) \quad (3)$$

$$GC_l(t) = \sum_{i=1}^{6} w_{il}^{(GC)} \left(\frac{1}{\frac{|i - i_{max}(t)|}{5} + C} - C \right) \quad (4)$$

$$GC_r(t) = \sum_{i=1}^{6} w_{ir}^{(GC)} \left(\frac{1}{\frac{|i - i_{max}(t)|}{5} + C} - C \right) \quad (5)$$

where l,r is the contribution of the process to the left and right motor, respectively, i_{max} is again the index of the maximally active ambient light sensor and C is a constant. The weights $w_{ij}^{(GC)}$ are chosen such that this process will actually lead to phototaxis behavior.

recharge: Once the robot is in the charging station the *recharge* process causes it to slow down or stop (depending on the activity of the other processes). The activity of this process is a function of the energy inflow, dE. The energy inflow is large when the battery level is low and *vice versa* (see below). Thus, when the robot enters the charging station with a low battery level dE will be large and as a consequence the *recharge* process will try to stop the robot :

$$\overline{RC(t)} = (RC_l(t), RC_r(t)) \quad (6)$$

$$RC_l(t) = w_l^{(RC)} \frac{1}{1 + exp(\alpha - \beta dE(t))} \quad (7)$$

$$RC_r(t) = w_r^{(RC)} \frac{1}{1 + exp(\alpha - \beta dE(t))} \quad (8)$$

where $w_i^{(RC)} < 0$ so that large activity in $RC_i(t)$ will decrease the motor speeds significantly or even halt the robot. In the experiments presented in this paper, there is (a) an energy decrease proportional to the distance traveled, (b) a constant decrease of energy due to a metabolic rate and (c) a decrease of energy when the agent is carrying objects. Thus, the energy dynamics are as follows:

$$E(t+1) = E(t) - MR - D(t) - (IR_7 + IR_8) * D(t) \quad (9)$$

where $E(t)$ is the energy or battery level, MR is a constant metabolic rate, and $D(t)$ is the distance traveled between two time steps. The metabolic rate causes energy loss even when the robot is not moving. The distance traveled, $D(t)$, is approximated by taking the summed and averaged difference of the wheel encoder values WE at two subsequent steps. Since we are

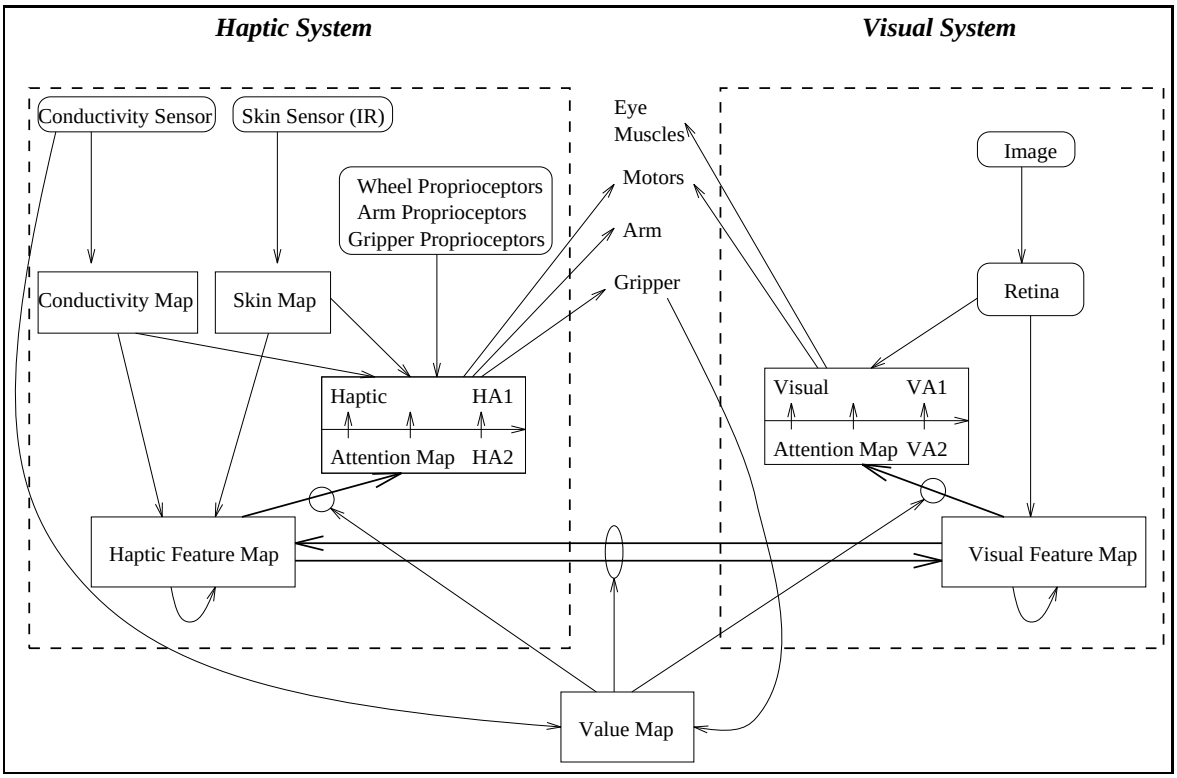

Figure 3: Overview of the visual and the haptic exploration systems (the shaded processes in figure 2). Each system consists of sensory sheets, an attentional map and a feature map. There is a learned crossmodal interaction via reentrant connections between the haptic and visual feature map.

interested in having a complete robot one important task of the robot is the keep its battery level in a certain range, i.e. it should sustain itself over an extended period of time. There are two aspects of this: First, the agent should go to the charging station whenever its battery level is low. Second, the distance from the charging station has to be taken into account: traveling a large distance costs a lot of energy. The motivation function of the *go to station*, M^{GC}, has three components M_E,M_D and M_O. M_O is a linear function of the object presence and arm position sensor. The M_D component implements the dependence of the *go to station* process activation on the distance from the charging station. If the robot is far away from the charging station the influence will be greater than when it is close to it. We use the activation of the light sensors to estimate this distance:

$$M_D = 1 - IR(i_{max}) \quad (10)$$

M_E is a function of energy. We used the following function:

$$M_E = 1 - \frac{1}{1 + exp(\alpha - \beta E)} \quad (11)$$

Typical values for α and β are 15 and 45 respectively. In essence, the motivation is largest, when energy is low and the distance is large (i.e. the agent is far away from the charging station). Other functions can be used such as a quadratic [9].

3.2 Sensory-motor coordination

As pointed out earlier, categorization can be viewed as a process which involves action as well. Categorization then becomes a matter of *sensory-motor coordination* rather than information processing alone. Instead of looking at a particular (fixed) sensory pattern we let the agent explore the object. In the present architecture there are two types of exploration behaviors: visual and haptic exploration. This part of the architecture is shown in greater detail in figure 3. The *visual exploration system* consists of a retina, a visual attention map and a visual feature map. We use an image of 160x120 pixels as an input to the *retina*. An artifical fovea (20x20 pixels) is extracted from the center of the retina. The spatially averaged periphery of the retina gives input to area VA1 of the *visual attention map*. VA1 responds to bright spots in the image. The visual attention map is connected to the wheel motor map. These connections implement a continuous mapping between location of activity in the attention map and the resulting translational and rotational movements. As a result, the robot always orients its body towards areas of interest (bright spots in the present model). In addition, the attention map

is connected to artificial eye muscles that move the fovea to bright spots. Thus, together with the motor map and the eye muscles the attention map forms a complete (visual) *attentional sensory-motor loop*: it brings the robot to relevant places in the environment while at the same time keeping the eye focussed on the spot where the robot is heading to.
The fovea in turn is connected to the *visual feature map*. The feature map responds to relevant features in the fovea. In the present model these are defined to be horizontal and vertical edges, i.e. texture. The visual feature map is connected back to area VA2 of the visual attention map via modifiable weights. The main idea behind these weights is that the categorical responses of the classification couple (i.e. the two feature maps) should modulate the attentional sensory-motor loops by either enhancing or breaking it. Area VA2 consists of a population of inhibitory and excitatory units that are connected via prewired weights to area VA1. The activity of these units increases as the weights from the feature map to the attention map evolve. These weights are topographic in the sense that areas that code for strongly textured input are mapped onto the excitatory population, while areas that code for bright or non-textured objects are mapped onto the inhibitory population. High activity of the inhibitory population will lead to a suppression of activity in VA1 and in turn to a breaking of the visual attentional loop. The *value map* implements a general bias in the system. The value map receives input from the resistivity sensor and the arm proprioceptors. The basic motivation behind these connections is that the robot should learn only when it explores an object. Activity in this map acts like a gating function and is used as a reinforcement signal for the synaptic modifications between the feature and attentional map as well as for the reentrant connections between the two feature maps. Finally, the visual feature map is connected to the haptic feature map via reentrant connections. Again, the synaptic growth of these connections is modulated by activity of the value map. The value map is strongly active when the robot has something in the gripper. In this way the robot only learns about objects when it is exploring them with its arm-gripper system.
The overall functionality of the *haptic exploration system* is similar to the one just described for the visual system. The sensory maps consist of a *conductivity map*, and a *skin map*. They get input from a conductivity and a skin sensor, respectively. Both maps are connected to area HA1 of the *haptic attention map*. Similar to the visual equivalent, the haptic attention map and the wheel, arm and gripper motor maps form an attentional sensory-motor loop. Again there is a continuous mapping of location of activity in the attentional map and the resulting translational and rotational movements. The main result is that as soon as the robot has made body contact with an object (which is sensed by the skin sensors) it will bring the object to the front of its body. This "haptic focussing" leads to an increase of activity in area HA1 of the attention map and in turn to large activity in the arm motor map, causing the robot to lower it's arm. Lowering the arm leads to increased activity in the arm proprioceptors and in turn to an even more increased activity in HA1. As a result the robot will start exploring the object by closing the gripper. Area HA2 consists of inhibitory and excitatory populations of units. They receive inputs from the haptic feature map map and they are fully connected with the units of area HA1. The projections between the haptic feature map and HA2 are topographic in the sense that areas that code for non-conductivity are mapped onto the inhibitory population, while areas that code for conductivity are mapped onto the excitatory population. HA2 and HA1 are fully interconnected. High activity of the inhibitory population will lead to a suppression of activity in HA1. As a result the gripper will be opened and the arm will be lifted. At the beginning of a trial the main source of activity in the haptic attention map stems from the skin map. This is a kind of "haptic reflex" that makes the robot haptically track and explore objects. Over time connections between the haptic feature map and the haptic attention map (area HA2) evolve. This leads to an amplification of attention for relevant objects and to a breaking of the haptic attentional sensory-motor loop when the robot encounters irrelevant objects.

3.3 Activation and learning rules of neuronal fields

The generic equation that describes the activation rule for the attention maps is:

$$a_i^{AM}(t) = \sum_{i=1}^{N^{FM}} a_j^{FM}(t) w_{ij}^{FM}(t) + \sum_{i=1}^{N^s} a_j^s(t) w_{ij}^s(t) \quad (12)$$

where $a_i^{AM}(t)$ is the i-th unit of the attention map, $a_j^{FM}(t)$ is the activation of unit j of the feature map, $w_{ij}^{FM}(t)$ is the weight from unit j in the feature map, $a_j^s(t)$ is the activation of unit j of the sensory map and $w_{ij}^s(t)$ is the weight connecting unit j to unit i. The activity of the units in the feature maps, $a_i^{FM}(t)$, is computed using a "leaky integrator" activation function [2]:

$$a_i^{FM}(t) = \beta a_i^{FM}(t-1) + \alpha [\sum_{i=1}^{N^r} a_j^r(t) w_{ij}^r(t) + \sum_{i=1}^{N^s} a_j^s(t) w_{ij}^s(t)] \quad (13)$$

where $0 < \beta < 1$ relates to the time constant of the unit, $0 < \alpha < 1$ is the attack parameter, $a_j^r(t)$ is the activation of unit j of the other reentrantly connected feature map, $w_{ij}^r(t)$ denotes the reentrant connection strength from unit j to unit i, $a_j^s(t)$ is the activation of unit j of the sensory map,$w_{ij}^s(t)$ is the weight connecting unit j to unit i. The connections between the sensory map and the feature map are chosen such that they result in a population coding of the average sensory activity (conductivity/skin map activations for the haptic and texture for the visual feature map, respectively). Equation 13 makes the activation of the units not only depend on the weights and the current

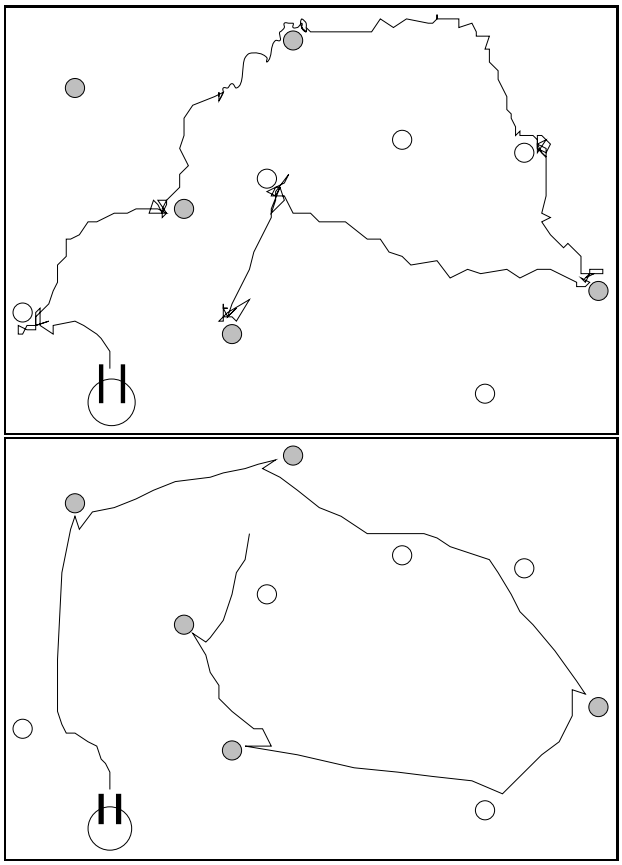

Figure 4: A typical trajectory of the robot before (top) and after (bottom) learning had occured. For visualisation purposes objects were removed from the gripper once the robot had grasped them.

inputs but also on the activation of the previous time step. Besides being a biologically plausible activation function, the leaky integrator function has the important property that it leads to stability of responses and a kind of "low pass" filtering on the input. The connection strenghts between the haptic and the visual feature map, and between the feature maps and the attention maps are computed using a Hebbian learning scheme:

$$\Delta w_{ij} = v(t)(\eta a_i(t) a_j(t) - \epsilon a_i(t) w_{ij}(t)) \qquad (14)$$

where w_{ij} represents the strength of the connection between the presynaptic unit j and the postsynaptic unit i, $v(t)$ is the value signal (activity of the value map), $a_j(t)$ is the activation of the pre-synaptic unit, and η, ϵ are the learning rate and decay parameters respectively.

4 Results

Experiments were conducted on a flat arena (100 cm x 100 cm) with walls (8 cm height) on each side. Objects were of 1.5 cm in diameter and 2 cm high, the shape of the objects was cylindrical (see figure 1). There were conductive and non-conductive objects in the environment. The conductive objects had a strongly textured surface while the non-conductive had a white or only slightly textured surface. The robot's task was to collect the conductive objects. We present results on the behavioral level as well as the underlying internal dynamics.

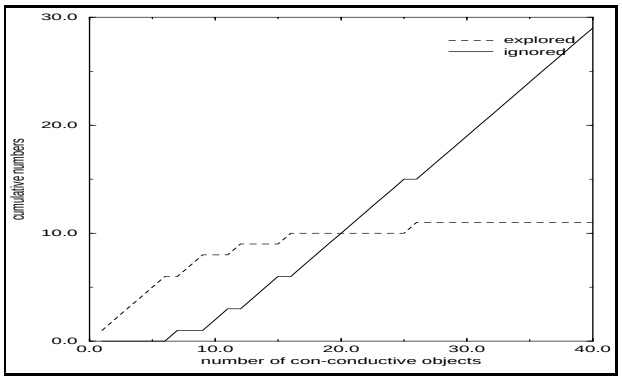

Figure 5: Cumulative number of exploration and ignoration steps for 40 non-conductive objects. The data are means over 20 trials.

4.1 Behavior

The behavior of the robot as it moves around in the environment and explores objects is shown in figure 4. The trajectories were recorded with a video camera and then hand traced. Figure 4 (*top*) shows a typical trajectory at the beginning of a trial. White and shaded circles indicate non-resistive/non-textured and resistive/textured objects, respectively. It can be seen there is no distinct behavior for the two types of objects. Rather the robot approaches all objects and explores them. Figure 4 (*bottom*) shows a typical trajectory after the robot has encountered 10 objects of each type. Two main results can be taken from the traces in figure 4. First, the robot has stopped exploring both types of objects. Rather the behavior is now governed by the dynamics of the classification couple. Second, the robot "ignores" non-conductive objects while it grasps the conductive ones (without first exploring them). We use the term "ignoration" instead of "avoidance" to indicate that there is no separate avoidance module. Rather the avoiding is achieved by breaking the attentional sensory-motor loop. In order to quantify this learning process the number of non-conductive objects explored and ignored was recorded for 20 trials. Each trial ended when the robot had encountered 40 non-conductive objects. Figure 5 shows the averaged cumulative number of non-conductive objects the robot explores and ignores. At the beginning of the trials the agent always explored non-conductive objects. After it had explored around 6 objects it started to ignore them. Because the weights had not been sufficiently evolved it still explored some of the objects. After having encountered around 12 objects the robot only explored because of errors in the sensory readings. One major problem when working with mobile robots is that they can get stuck. In

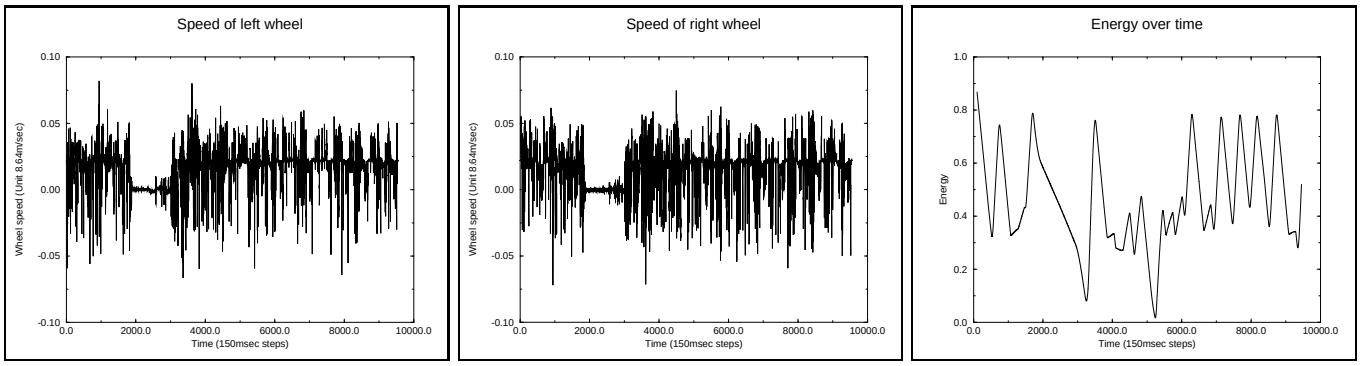

Figure 6: The speed of the left and right wheels and the energy over time for a typical trial (see text for explanation).

our set-up the agent can, for example, get trapped in a corner when there are some objects located nearby the corner. There are several ways to approach this problem. One is to introduce special processes which survey the success or failure of others ([4]). In our approach, however, this is not necessary for the following reason. Energy decreases even when the robot is stuck because of the metabolic rate. After a certain amount of time this will lead to an increased activity in the *go to station* process. As a result, there will be an increase of activity in the speed quantities and the robot will start moving again. This is shown in figure 6. As can be seen in the figure, the agent gets trapped after about 1800 steps: The speed of the left and the right wheel become zero. At the same time, the energy level starts to decrease because of the metabolic rate and the energy level gets very low at around 2400 steps. Because of the low energy level the activity in the *go to station* process starts to increase and contribute to the motor speeds. As a result, the robot starts moving again. Because it cannot really get out of the impasse initially, there is an increased wiggling behavior (the bursts around 3100 steps in the figure). Because of this the agent finally manages to get out of the impasse and go to the charging station. Another result that can be seen in figure 6 is that the robot is in fact self-sufficient in the sense that it keeps it's battery level in a safe range.

4.2 Internal dynamics

In the following we present results on the categorization dynamics. One of the main consequences of using reentry as the mechanism for categorization is that temporal correlations or synchronizations in the activity between neuronal areas emerge. In order to quantify these correlations we estimated the level of synchronization between the visual and the haptic feature map. We have quantified synchronization as the product of the average activities of the areas V1, V9 and H1, H9, in the visual and haptic feature maps, respectively. Units in area V1 selectively respond to non-textured regions in the fovea while units in V9 respond to textured ones. Similarly, H1 and H9 of the haptic feature map selectively respond to non-conductive and conductive stimuli, respectively.

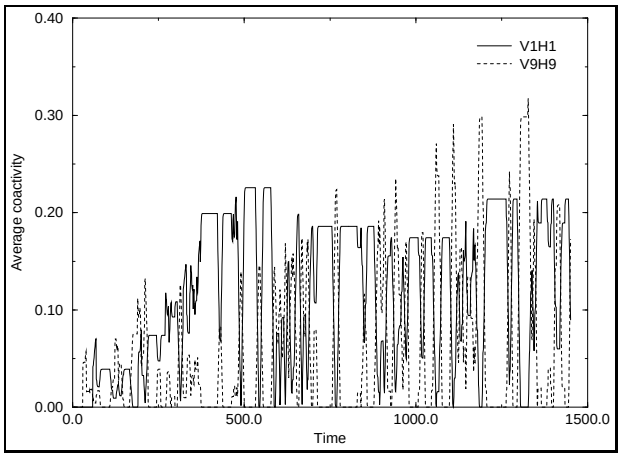

Figure 7: Level of synchronization between areas H1, V1 and areas H9, V9. See text for explanation.

Figure 7 shows the evolution of synchronization between the classification couple as the robot explores objects. Three main results can be taken from this figure. First, the *level of synchronization* between areas V1, V9 and areas H1, H9 (denoted as V1H1 and V9H9 in figure 7, *top*) at the beginning of the trial is significantly lower than after learning has occured. For example, until around 300 steps the mean synchronization level is around 0.05 while later in the trial it increases up to 0.3 implying levels of average activity around 0.55. This increase is due to the evolving reentrant weights that couple the activities in the two feature maps (see below). Second, the *difference* in the synchronized activities between the two areas (V1H1 and V9H9) increases significantly. Third, the *coherence* of synchronous activity between the maps increases. As can be seen in the figure, the dynamics of the coactivities become strongly coupled over time. This is also reflected in the correlations V1H1-V9H9 that reach an average level of around -0.9 (data not shown). Finally, figure 8 shows the evolution of the average weights from the feature maps to the attentional maps. These weights couple the categori-

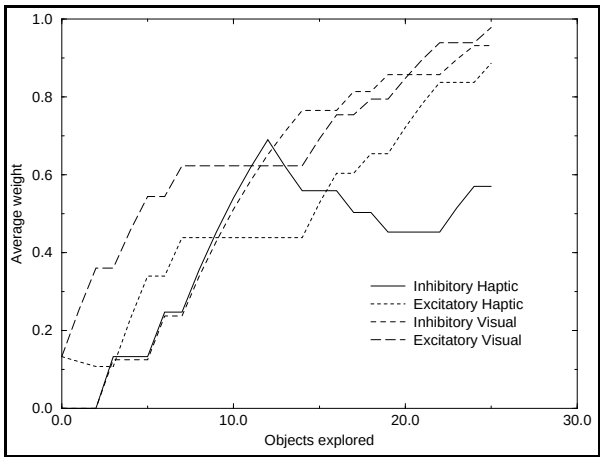

Figure 8: Average weights from the feature maps to the attention maps. See text for explanation.

cal responses to the attentional sensory-motor loops. As can bee seen in the figure, weights to both excitatory and inhibitory populations in the attention maps increase significantly with the number of objects the agent has explored. More specifically, area V1 of the visual feature map (responding to non-textured stimuli) becomes associated with the inhibitory population of the attention map (area VA2, see figure 3). Strong weights also evolve between area V9 (responding to textured stimuli) and the excitatory population of the visual attention map. Together these connections lead to a suppression of activity in the visual attention map for non-textured objects while the encounter of textured objects lead to an enhanced activity. Similar weights evolve between the haptic feature map and the haptic attention map. Area H1 (coding non-resistive input) becomes associated with the inhibitory population of the haptic attention map while area H9 which selectively responds to conductive objects becomes associated with the excitatory population of the attention map. Again, these connections have the consequence that the activity in the haptic attention map gets suppressed for non-resistive objects and enhanced for resistive ones.

5 Summary and Conclusions

In this paper we have addressed the problem of integrating the various competences that are required for a complete autonomous agent. We have proposed a new architecture, the EBA, which is based on the principle of parallel control with loosely coupled processes. EBA provides a natural solution to the notorious problem of when the agent should do what, i.e. the action selection problem. We have demonstrated how this could work even when a "difficult" competence such as category learning is embedded in the architecture. In our case study categorization is based on the principle of sensory-motor coordination rather than the one of information-processing. We have shown that based on this architecture the agent is able to learn to categorize the objects in its environment. The categories that evolve are expressed in the synchronization of activity in the classification couple.

References

[1] V. Braitenberg. *Vehicles.* Kluwer Academic Publishers, 1984.

[2] G. J. Chappell and J. G. Taylor. The temporal kohonen map. *Neural Networks*, 6:441–445, 1993.

[3] G. M. Edelman. *Neural Darwinism.* Basic Books, New York, 1987.

[4] E. Gat. On the role of stored internal state in the control of autonomous robots. *AI Magazine*, Spring 1993.

[5] J. E. Hummel. Object recognition. In M. A. Arbib, editor, *The hand book of brain theory and neural networks*, pages 658–660. MIT Press, Cambridge, MA., 1995.

[6] J. K. Kruschke and M. A. Erickson. Five principles for models of category learning. In Z. Dienes, editor, *Connectionism and Human Learning*, pages 365–413. Oxford University Press, Oxford, 1995.

[7] D. Lambrinos and C. Scheier. Extended braitenberg architectures. Technical Report AI.95.10, AI Lab, Computer Science Department, University of Zurich, 1995.

[8] D. Lambrinos, C. Scheier, and R. Pfeifer. Unsupervised classification of sensory-motor states in a real world artifact using a temporal kohonen map. In *Proceedings of the International Conference in Artificial Neural Networks ICANN'95*, pages 467–472, Paris, October 1995.

[9] D. McFarland and T. Boesser. *Intelligent Behaviour in Animals and Robots.* MIT Press, Cambridge, Massachusetts, 1993.

[10] R. Pfeifer and C. Scheier. *Introduction to "New Artificial Intelligence".* AI Lab, Computer Science Department, University of Zurich, November 1995. Lecture Notes.

[11] R. Pfeifer and C. Scheier. Sensory-motor coordination: the metaphor and beyond. *Journal of Robotics and Autonomous Systems*, in press.

[12] C. Scheier and R. Pfeifer. Classification as sensory-motor coordination: a case study on autonomous agents. In *Proceedings of the Third European Conference on Artificial Life ECAL95*, pages 656–667, Granada, Spain, June 1995.

[13] Masanao Toda. *Man, robot and society.* Nijhoff, The Hague, 1982.

[14] T. Tyrell. *Computational Mechanisms for Action Selection.* PhD thesis, University of Edingburgh, 1990.