IEEE-INNS Intern. Joint Conference on Neural Networks (IJCNN 2001),Washington, 2001, pp. 1261-1266, IEEE/Omnipress, 2001

Integration of a Fuzzy ART Approach in a Biologically Inspired Sensorimotor Architecture[1]

A. Heinze, H.-M. Gross, and D. Surmeli
Ilmenau Technical University, Department of Neuroinformatics, P.O.B. 100565,
D-98684 Ilmenau, Germany
{heinze}@informatik.tu-ilmenau.de

Abstract

We present a biologically motivated computational model that is able to anticipate and evaluate multiple hypothetical sensorimotor sequences. Our Model for Anticipation based on Cortical Representations (MACOR) allows a completely parallel search at the neocortical level using assemblies of rate coded neurons for grouping, separation, and selection of sensorimotor sequences. For a vision-controlled local navigation of a mobile robot Khepera, we can demonstrate that our anticipative approach outperforms a reactive one. We also compare our explicitely planning approach with the implicitely planning Q-learning. Further we show the advantages of ART [2] as a realization of MACOR for real-world tasks.

1 Introduction

Based on findings for the sensorimotor character of perception [1, 5], we developed an alternative approach to perception that avoids the common separation of perception and generation of behavior and fuses both aspects into one consistent neural process. In this approach, perception of space and shape in the environment is regarded to be an active process which anticipates the sensory consequences of alternative hypothetical interactions with the environment, that could be performed by a sensorimotor system, starting from the current sensory situation. This approach is supported by biological findings. For example, it was shown that such planning of motor actions takes place in the secondary motor areas [9]. Thach (1996) found that the premotor parts of the brain are active both in planning movements to be executed as well as in thinking about movements that shall not be executed.

Based on these findings, we developed our Model for Anticipation based on Cortical Representations MACOR, presented in sec. 2. It is intended as a general scheme for sensorimotor anticipation in a neural architecture. The model does not attempt to provide a detailed description of a specific cortical or subcortical structure, but we try to capture some general properties that are relevant to our "perception as anticipation"-approach in brain-like systems (for details see [3]). The objective of this paper is to demonstrate the efficiency of our anticipatory approach for a real-world sensorimotor control problem, the local navigation and obstacle avoidance of a vision-controlled mobile robot showing non-holonom movement characteristics. We compare the achieved navigation behavior with other non-anticipative approaches like reactive control and Q-learning-based control (sec. 3.2). Moreover, in section 3.1 we show that a fast generation of a stable representation of situations in real sensorimotor tasks may be realized more efficiently by systems derived from the Adaptive Resonance Theory (ART) than by statistical clusterers.

2 Model for generating visuomotor sequences

For real sensorimotor tasks, the development of MACOR requires compliance with a number of demands, such as a fast generation of stable representations of situations, the recording of rare, but behaviorally relevant situations and a simultaneous (without sensory pretraining) adaptation of sensorimotor transitions. In sec. 2.1, we demonstrate both the agreement of ART-systems with those requirements, and the feasibility of an integration of an ART-system into MACOR.
In sec. 2.2, we introduce the anticipation-based model MACOR in a realization based on Fuzzy ART.

2.1 Fuzzy ART for the generation of sensorimotor representations

In the following, the term statistical clusterer shall refer to systems that approximate the probability density of the input data (e.g. Self-Organizing Feature Maps or a Neural Gas [4]). Only statistically often presented input data will be represented with high accuracy by those clusterers. In contrast, ART-systems, and es-

[1] This work was partially supported by DFG Graduate college GRK 164.

pecially Fuzzy ART in fast learning mode, are able to capture input data after a single presentation and thus, independent of their probability density (plasticity). Its weights form subspaces in the input space, whose size depend on a single set parameter, the vigilance $\rho \in [0, 1]$. The higher the vigilance, the smaller the subspaces and therefore, the more accurate the representation of the inputs. When an input is presented to ART, and it can be represented by a neuron, the system will move into a resonant state. Otherwise, a new neuron must be inserted and its weights calculated from the input sample. Another essential feature is the stability of subspaces in Fuzzy ART. Since subspaces are only allowed to grow and are bounded by the vigilance, no learned data will ever be lost.

Stability and plasticity: In MACOR, the generation of a visuomotor representation and the adaptation of sensorimotor transitions (action evaluations) may be realized either simultaneously or sequentially. Although the effects of the sensorimotor representation can only be investigated separately following sequential learning, we are much more interested in simultaneous learning due to its fewer learning iterations and better biological plausability.

During simultaneous learning with a statistical clusterer, action evaluations will be associated with neurons that still undergo drastic changes of their representations. Thus, the evaluations must be relearned to reflect the new representations. Such changes are unavoidable and will occur frequently during the beginning of learning. They produce a slow convergence to the resulting behavior and may cause failures regarding the realization of system goals.

With Fuzzy ART, the representations of neurons are stable, such that the associated action evaluations do not need to be adapted to changing representations and simultaneous learning becomes possible.

In modification to the original algorithm for Fuzzy ART, which assumes an unlimited number of free neurons, MACOR restricts the number of available neurons to a maximum due to biological and technical resource limits. If a presented input cannot be represented by a neuron (no resonance) and there are no more available neurons, our modifications use the neuron with minimal distance as the best matching neuron, but without any subsequent weight adaptation.

2.2 Model for anticipation

In the framework of a visually-guided navigation task (sec.2.5), our architecture (see fig.1) processes as sensorimotor information a visual input yielded from an omnidirectional camera, supplemented by the last motor command executed.Within the architecture, those inputs are represented by a Fuzzy ART architecture [2]. The visual part is represented by the F2-nodes of a Fuzzy ART network. Each of these nodes contains several neurons, which represent the motor commands available to bring the system into the considered visual situation (visuomotor column). Between the resulting visuomotor representation of the current situation and that of the preceeding one, connections incorporating the certainty of that transition and its evaluation are adapted (sec.2.3).

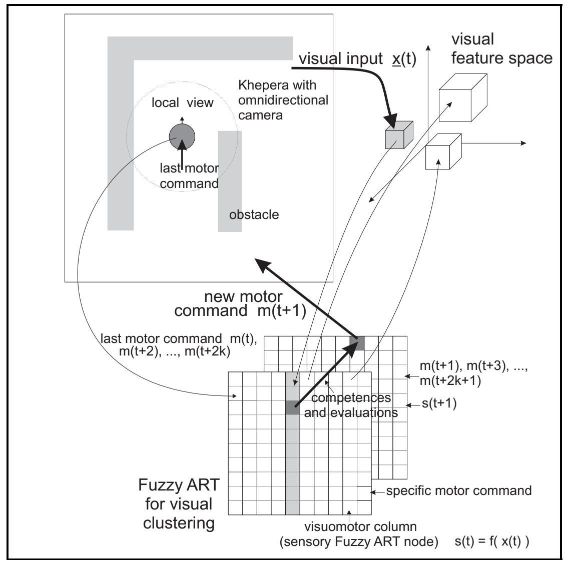

Figure 1: Overview of the architecture. To represent visuomotor inputs, each Fuzzy ART node representing a specific visual input contains several neurons to represent motor commands available to bring the system in this situation. Connections between the present and the preceeding visuomotor situations are adapted to reflect both the reinforcement incurred during that particular transition and the competence gained. By activation of a specific neuron and the subsequent spread of activity through its weights to activate other neurons, alternative sensorimotor sequences may be generated.

This architecture embeds the functionality required for a *parallel* generation of sequences of sensorimotor hypotheses (sec.2.4). In the context of the biological foundations of the architecture, this parallel generation is realized by a spread of activity through connections between the visuomotor columns. In the following investigations we use a restricted set of discrete motor commands to investigate the effect of anticipative behavior (see sec.2.5, too).

2.3 Learning within the map

The learning of sensorimotor connections between the present and the preceeding visuomotor situations takes place after each executed motor command. To evaluate the certainty of the existance of sensorimotor connections, we investigated several approaches, for example actual transition probabilities. Because of the clustering of the input space, high competence weights between neurons in the same sensorimotor assembly are established. The correspondingly lower competence weights between different situations yield lower sequence evaluations. These sequences are subsequentely not selected, but important for the consideration of movements resulting in a collision and thus for achieving the system goals.

Using the simple learning equ.1 realized as a sigmoidal function for **competence weights**, a specific weight can quickly reach a maximum value $w^{\mathrm{C}}_{\max}$, determined by the parameter μ and σ. Thus all weights hold the same value after a sufficiently high number of adaptation steps, such that the action selection (sec.2.4) depends only on the values of the evaluation weights. Therefore, the competence weights are only meaningful for unknown transitions, which are devalued by the competence. In equ.1, z_{ij} is the number of adaptations of the respective weight.

$$w^{\mathrm{C}}_{ij}(t) = \frac{1}{1+e^{\frac{z-\mu}{\sigma}}} \quad \mu = 2.8, \sigma = 0.4 \qquad (1)$$

For learning of **evaluation weights**, a simple form of reinforcement learning is used. Thus, the evaluation weights hold only the expected evaluation for the next transition, yielding a purely reactive behavior. Only by means of internal simulation, an anticipative behavior can be generated. The reinforcements r yield the evaluation of the generated behavior. In our experiments, they were chosen to reward a collision free, straight navigation behavior and to punish turns and collisions. For the adaptation of evaluation connections w^{r}_{ij}, the number z of adaptations of the respective connection w_{ij} is used to determine an adaptive learning rate $\frac{1}{z^r_{ij}}$.

$$w^{\mathrm{r}}_{ij}(\mathrm{t{+}1}) = w^{\mathrm{r}}_{ij}(\mathrm{t}) + \frac{1}{z^r_{ij}} \cdot (\mathrm{r} - w^{\mathrm{r}}_{ij}(\mathrm{t})) \qquad (2)$$

For comparision a special form of reinforcement learning, Q-learning [10, 8] was also investigated, because it is a model free, but implicitely planning approach.

$$w^{\mathrm{r}}_{ij}(\mathrm{t{+}1}) = w^{\mathrm{r}}_{ij}(\mathrm{t}) + \frac{1}{z^r_{ij}} \cdot (\mathrm{r} + \gamma \cdot w^{\mathrm{r}}_{ki}(t) - w^{\mathrm{r}}_{ij}(\mathrm{t})) \qquad (3)$$

$$w^{\mathrm{r}}_{ki}(t) = \max_{l} w^{\mathrm{r}}_{li}(t) \qquad (4)$$

By choosing the parameter γ to control the planning horizon $\gamma > 0$, the prediction of evaluations of preceeding transitions can be stored in one connection (strong Q-learning, equ.3 and 4). Q-learning becomes a planning free approach by choosing $\gamma = 0$ (weak Q-learning, equ.2).

2.4 Generation and evaluation of sensorimotor sequences

For generation of sensorimotor sequences, a specific neuron in a Fuzzy ART assembly is activated by the current sensory situation and the last executed motor command. This neuron propagates its activity y^{α}_j to all other interconnected neurons $i \in S = [0, n \cdot m - 1]$ using its competence connections w^{C}_{ij}, where n is the number of sensory assemblies and m the number of motor neurons within each assembly. The activated neurons may in turn activate further neurons, resulting in a mechanism of internal simulation and thus the generation of whole sequences of sensorimotor hypothesis (equ.5).

$$y^{\alpha}_i(t+1) = \max_{j \in S} w^{\mathrm{C}}_{ij}(t) \cdot y^{\alpha}_j(t) \qquad (5)$$

Since the maximal value of the competence connection is less than 1.0, a subsequent neuron will always be less activated than its predecessor. Also, a neuron can only be activated by its maximum input activity. This supplies a simple stopping criterion for the propagation of interconnecting sequences. Simultaneous to the parallel generation of sequences of sensorimotor hypotheses, the model realizes the selection of the best evaluated sequence by a backpropagation of local sequence evaluations, as shown in equ.6. This backpropagation starts as the start neuron activates further neurons and each evaluation is backpropagated to the respective sequence predeccesor.

$$y^{\beta}_i(t+1) = w^{\mathrm{r}}_{ij}(t) \cdot y^{\alpha}_i(t+1) + \max_{r \in S} y^{\beta}_r(t) \qquad (6)$$

The activity backpropagated to the start neuron represents the highest sequence evaluation in each time step. Thus in each time step, an action selection is possible, which improves as internal simulation goes on. This mechanism yields high cumulated sequence evaluations for well known and highly evaluated transitions. To realize a reactive behavior, the process of internal simulation runs only for one time step. This means the starting neuron activates further neurons which then propagate their evaluations directly back onto the starting neuron.

2.5 Experimental scenario

Because of their embodiment and situatedness, robots are ideal systems to demonstrate the advantages of an anticipation based sensorimotor control compared to a reactive one. To navigate successfully, for example, to avoid obstacles or to go through narrow passages, they have to consider their physical and mechanical properties and constraints (e.g. inertia, holonom or nonholonom kinematics). A mobile system, a robot or an animal, that is not able to learn and consider its constraints and their sensorimotor consequences, will not be able to evolve successful navigation behaviors.

To demonstrate the advantages of an anticipation based sensorimotor control, we used the mobile robot Khepera as a non-holonom system. Because of the used restricted action space, our system has to consider its constraints and is forced to start an early obstacle avoidance by internal simulation to realize a successful navigation behavior.

In our investigations, the visual sensory inputs were provided by an omnidirectional camera. After a transformation into a physiological color space [6] specially tuned receptive fields extract the color distribution around the robot, which gives an implicit description of the obstacle arrangement.

3 Results

We first show that Fuzzy ART is more efficient than a statistical clusterer for the generation of visuomotor representations as the basis for an adaptation of visoumotor transitions. To demonstrate the advantages of explicit planning systems, in sec.3.2 we provide a comparison between the following approaches:

- the presented system which explicitely plans by internal simulation,
- a purely reactive system which operates exclusively on the current sensorimotor situation, and
- an implicitly planning approach, Q-learning.

3.1 Local Navigation based on a Neural Gas or Fuzzy ART

We intend to show the advantages of Fuzzy ART compared to a statistical clusterer, here exemplarily a Neural Gas (NG), based on fulfillment of requirements on real sensorimotor systems, namely an *online*-enabled generation of sensorimotor representations and the generation of *actions appropriate for the system goals*. We measure the *necessary number of training steps* as well as the *generated behavior* for both systems.

Sensory representation: To generate sensorimotor representations, both systems were given 30 sensory assemblies with three motor commands each: move straight and left ($\phi = 12°$, $s = 2.1\ cm$), move straight and right ($\phi = -12°$, $s = 2.1\ cm$) and move straight ($\phi = 0°$, $s = 2.1\ cm$). In accordance with the results of preliminary investigations, the vigilance of Fuzzy ART was set to 0.8. For the NG, the parameters learning radius and learning rate were determined offline on a data set using a breadth search for the combination yielding the smallest clusterer error.

Figure 2 shows the expected lower clusterer error for the NG versus that for Fuzzy ART. Also evident is, that this error was achieved after the fifth presentation for the NG, whereas Fuzzy ART converged after the first presentation. The higher error for Fuzzy ART stems probably from the calculation of the smallest Euclidean distance from the middle of the subspace of the neuron. As those subspaces often do not exhibit similar ranges in the different input dimensions, but typically, span much greater ranges in one or a few dimensions, an input may cause a fairly high clusterer error despite being represented by the neuron.

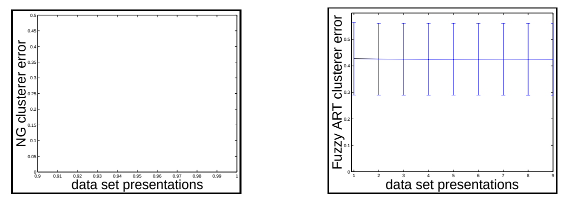

Figure 2: Clusterer error for Fuzzy ART and NG. Following each complete presentation (13100 visual inputs), we computed the average Euclidean clusterer error (for Fuzzy ART, from the middle of the corresonding subspace). Although the learning equations for Fuzzy ART utilize Fuzzy AND to determine the distance between input and weight vector, we use the Euclidean distance for comparability of both systems The clusterer error of Fuzzy ART is much higher, but it is achieved after the first presentation. The variance of the clusterer error for Fuzzy ART is also much higher than for the NG.

The following investigations will show that even with the high clusterer error, Fuzzy ART is still able to provide a foundation for the adaptation of appropriate sensorimotor transitions. The advantages of a fast learning and stability outweigh the drawbacks in accuracy. In a first investigation, we compared a NG fully trained according to fig.2 with a system based on Fuzzy ART. The NG was allowed five presentations of the data set for sensorimotor adaptation, whereas the Fuzzy ART used only the 9000 starting patterns in the data set. A first indication, that only Fuzzy ART may be used for an online-enabled generation of sensorimotor representations, is given by the necessary training times: ca. 4.5 hours for Fuzzy ART and 32.7 hours for 65500 steps of the NG. After that, the advantage of fast learning can be shown by the much lower number of necessary training steps.

Adaptation of sensorimotor sequences: After the generation of sensorimotor representations, competence- and evaluation weights were trained according to eqns. 1 and 2. For this training, ca. 11000 randomly selected and about 2000 reactively selected actions were executed (equals about 6.5 hours of training). After the adaptation of the weights, the generated reactive behavior of both systems was evaluated and depicted, together with corresponding reinforcement signals, in fig. 3.

The results in figure 3, show the generated behavior, a collision avoidance, for both systems. The system based on Fuzzy ART is characterized by early turning movements, which result from the sensory representations generated.

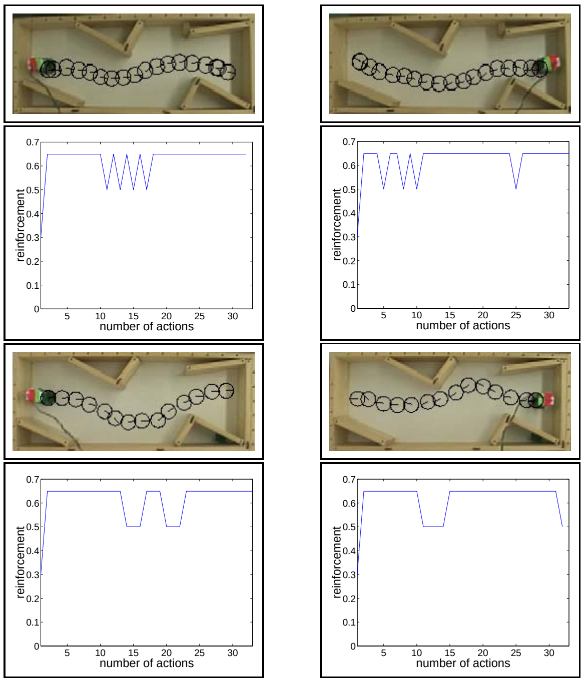

Figure 3: Local navigation using Fuzzy ART (**top**) and NG (**bottom**) in a system with 30 neurons and three motor commands per neuron. The circles mark the positions of the Khepera during the experiments. The reinforcement signals received are provided underneath the corresponding experiment. The system received less reinforcement for a turning movement than for a straight motion, whilst collisions were always punished by a reinforcement of 0.0. The trial ended in a collision or at the other end of the scenario. In the figure on the left top, a drift of the robot to the right is visible, in spite of an execution of a straight motion (cmp. reinforcement signals). This is caused by inaccuracies in the drives of the robot. However, the sucessful navigation is proof for an integration of this drift during the adaptation of sensorimotor sequences.

The advantage of fast learning with Fuzzy ART and its effect with respect to the behavior generated is shown in fig. 4. For these investigations, both systems were allowed 1000 steps for sensory pretraining. For the NG, the parameters learning radius and learning rate were again determined using a breadth search for the combination yielding the smallest clusterer error. All other experimental configuration is the same as before.

The results show that Fuzzy ART is more efficient for a generation of sensorimotor representations and based on it, an adaptation of transitions for an action evaluation, although there are very few such applications of ART-systems for sensorimotor tasks in the literature. In contrast to the NG with its high number of training steps for sensory adaptation, Fuzzy ART realizes an online-enabled learning with much fewer training steps.

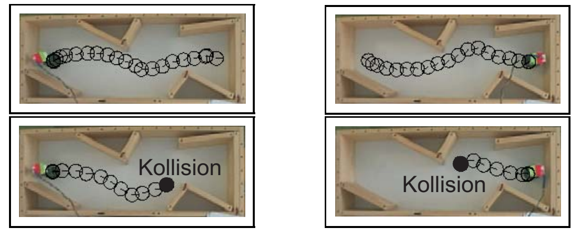

Figure 4: In contrast to the investigations in fig. 3, here both systems were allowed 1000 steps of training. Only Fuzzy ART (**top**) realizes a collision avoidance. The NG (**bottom**) is unable to generate an appropriate sensory representation as a basis for successful collision avoidance.

3.2 Comparison between reactive and planning approaches

For the investigation of anticipation-based behavior, we used Fuzzy-ART modules with 50 nodes (vigilance $\rho = 0.9$) for sensory clustering and also three executable actions per sensory assembly. After the structuring of the sensory clusters the competence and evaluation weights were adapted involving more than 10.000 typical sensorimotor situations. In Q-learning we used a planning horizon of $\gamma = 0.9$. Systems except Q-learning, which does not use competence weights, utilized equ.1 to adapt their competence weights. The behavior generated by the different systems was recorded over 10 trials each and is exemplarily depicted in fig.5. The generated behavior shows that only the explicitly planning system (fig.5 **middle**) was able to realize a successful obstacle avoidance considering the physical limitations of the non-holonom robot with a very limited action space. After the very first detection of a curve or obstacle, this system began to move to the left and pursued that movement until the turn was negotiated. In contrast, the reactive system (**left**) chose to drive straight forward until almost hitting the wall. Only directly in front of the wall it chose to turn right, thus unable to avoid a collision. The Q-learning system (**right**) initiated also an early turn to the left and continued until the turn was negotiated. Subsequently, the action 'go straight' was repeatedly chosen, because the visual situation in the middle of the floor is represented by the same Fuzzy ART node as the floor situation near the wall. Only in the situation directly in front of the wall, another Fuzzy ART node was activated and a turn initiated. Like Q-learning, our reactive approach also learned the action 'go straight' for the floor situation, but by means of internal simulation, our anticipatory approach chose mostly turns.

Figure 5: In the experiments, a trial ended either after a collision or a maximum of 40 steps. For explanation see text.

Further, in each trial the achievable total reinforcement was determined as the sum of the individual reinforcements received after each motor command. The respective mean values and variances together with typical trial lengths and the mean reinforcements for a single motor command are shown in table 1.

criteria		react. system	anticip. system	Q-learn.
r	mean	2.8	15.6	4.2
per trial	std	0.1	8.6	2.4
trial	mean	7.2	29.2	11.7
length	std	0.4	14.2	4.2
r per step	mean	0.61	0.55	0.57

Table 1: Mean reinforcement (r) and standard deviation (std) for 10 trials of the different systems. Although the explicitly planning system achieves the highest mean reinforcement per trial, it also yielded the highest std. This was caused by 4 of those 10 trials (maximal trial length was 40) that ended due to collision, most probably as a matter of the hardware fixation of the omnidirectional camera. The trials using the reactive and the Q-learning system all ended by collision and therefore lower trial lengths. Further, the mean reinforcements per step (not considering the step that resulted in a collision) show that the explicitely planning system got the lowest reinforcement per step as a result of more turns executed, which yield lower reinforcements than the movement 'go straight'.

These results give a strong indication that only the explicitly planning system is able to realize an obstacle avoidance for the inert, non-holonom roboter system. In comparison, the reactive system did not produce any avoiding actions at all. Though the Q-Learning system chose mostly proper motor commands it could not treat the perhaps non optimal sensorimotor representation.

4 Conclusions and outlook

We presented MACOR realized by Fuzzy ART that is able to anticipate and evaluate multiple hypothetical sensorimotor sequences. Using a framework of local navigation of a non-holonome robot, the advantage of anticipation based control was demonstrated empirically. The investigations performed demonstrate the advantage of anticipative systems compared to reactive systems and especially, the advantage of explicitly planning systems compared to implicitly planning ones.

This comparison is the basis of further investigations targeting a more complicated scenario, dealing with the aspect of changing system goals. These investigations are intended to show the advantages of reactively stored evaluation weights compared to weights whose values represent an implicit planning horizont, like Q-learning. These advantages become obvious for system goals that change over time, and which thus require readaptations of the sensorimotor transitions. With reactive evaluation weights in combination with internal simulation, only the actually different sensorimotor transitions must be relearned, which are only a few, to reflect the new system goal. So we want to show that an explicitly planning system may realize a flexible behavior according to a new task much faster than an implicitly planning system like Q-learning.

References

[1] Arbib, M.A. et.al., Neural Organization: Structure, Function and Dynamics, MIT Press, 1998

[2] G. Carpenter et.al. Fuzzy ARTMAP: A Neural Netw. Architure for Incr. Supervised Learning of Analog Multidimensional Maps, NN, pp.698-719, 1992.

[3] Gross, H.-M. et.al., Generative Character of Percept.: A Neural Architecture for Sensorimotor Anticipation, NN, 12:1101-1129,1999.

[4] Martinetz, T. M. et.al., Neural Gas Network for Vector Quantization and its Application to Time-Series Prediction, IEEE Trans.on NN, 4:558-568, 1993.

[5] Pfeifer, R. and Scheier, C., From Percept. to Action:The Right Direction?, Proc.of PerAc'94, IEEE Comp.Soc.Press, pp.1-11, 1994.

[6] Pomierski, T. and Gross, H.-M., Biological Neural Architecture for Chromatic Adaptation Resulting in Constant Color Sensations, Proc.IEEE Int. Conf. on NN, 734-739, 1996.

[7] Rabiner, L. R., A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition, Proc. of IEEE, 77(2):257-286,1989.

[8] Sutton, R. S. and Barto, A. G., Reinforcement Learning: An introduction, MIT Press: A Bradford Book, "Cambridge, MA, USA, 1998.

[9] Thach, W. T., On the specific role of the cerebellum in motor learning and cognition: Clues from PET activation and lesion studies in man, BBS, 19(3):411-431,1996.

[10] C. Watkins, Learn. from delayed rewards, King's Coll., Cambridge, UK, 1989.