Complex Neural Architectures for Emerging Cognitive Abilities in an Autonomous System

Philippe Gaussier (*) (**), Stéphane Zrehen (**)

(*) ENSEA ETIS, 6 av du Ponceau, 95014 Cergy Pontoise Cedex, France
(**) Laboratoire de Microinformatique, EPFL-DI, CH-1015 Lausanne
E-mail: gaussier@ensea.fr or zrehen@di.epfl.ch

Abstract

In this paper, we propose a novel neural architecture named PerAc which is a systemetic way to decompose the control of an autonomous robot in perception and action flows. We first present an application of the PerAc architecture to the simulation of a vision system with a moving eye. Then we propose a second application where the robot learns to return from any starting place to a previously discovered and learned position without any a priori symbolic representation.

Keywords
SENSORY-MOTOR LOOP - VISION - NAVIGATION - NEURAL BUILDING BLOCK

1 Introduction

Realization of neural control architectures that allow an autonomous robot to behave as a rat or simpler as an ant is a great challenge in Artificial Intelligence. Trying to create animals-like robots is based on old cybernetics concepts of self stability and homeostasy. But this kind of animal robots can only be vegetative. To be really autonomous, they have to be able to choose between different behaviors and to access to a certain form of freedom [McFarland94]. Designing such machines requires to be interested in cognitive science and more precisely in neurobiology, ethology and psychology. Indeed, neurobiology provides information about what the atomic elements of an intelligent system can be, i.e., a model of a formal neuron [Rumelhart86] or of a cortical column [Burnod89]. It also puts forward the need of having physiological data about the brain architecture [Van Essen83]. On another level, psychology and ethology are useful because of their description of brain functionalities [Treisman88] and their measures of quantitative brain performances. Such a frame is very important to point the directions in which engineers must advance in order to build machines that overstep their actual limitations. As a matter of fact, most of today's robots compute their actions from their perceived input by using models of their environment and are not able to imagine other models when they get in an unforeseen situation. They give pretty good results when their environment is adapted to their work but they are almost blind in a natural world: Too much data to analyze saturate the analysis capabilities of their " logical " brain. Moreover, in the industrial domain too, each new model of arm manipulator need to be modelized before planning to use it to manipulate objects. Would it not be more interesting to have them only learn their task?

   In this paper, we will present an autonomous mobile robot named Prometheus that can learn to return to an interesting place (its goal) in an unknown environment. The N.N. structure of Prometheus " brain " is based on the idea that Prometheus is intended to be in interaction with its environment like in

the enactivism paradigm [Maturana87], [Stewart91]. Contrary to an expert agent that knows how to reply to an arbitrary question about navigation problems or objects manipulation for instance, Prometheus is just something that learns to agree with its environment and its internal motivations. It has no global or complete representation of its world. What is stored in its " memory " is only what needs to be learned to act correctly in a particular situation. Should the universe collapse, the robot memory would have no more meaning.

In the first part, we will expose vision problems that appear in the task of recognizing marks in a scene, and the kind of information that Prometheus extracts. Next the visual information will be used to guide the robot in the direction of its goal. These application constraints lead us to search for as regular and simple as possible N.N. structures. They are brought together in the PerAc (Perception-Action) architecture and concept (fig. 1) which is inspired by works by Albus[Albus91], Burnod [Burnod89], Brooks [Brooks 86], Carpenter and Grossberg [Carpenter87] and Edelman [Edelman87].
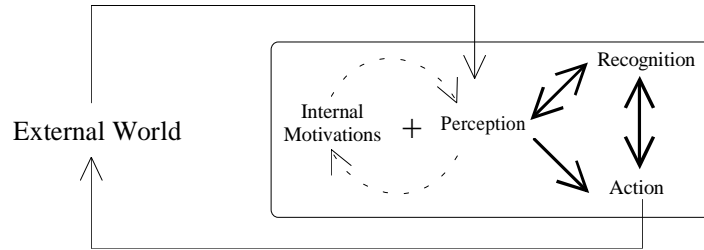

Figure 1: The concept of PerAc architecture to control autonomous robots and one PerAc unit bloc in bold.

PerAc is a systematic neural structure that allows on-line learning. It involves two data streams associated respectively to perception and actions in each part of the robot controller. From each perceived input, we suppose we can extract reflex information to control directly the robot action like in the behaviorist paradigm. But there is also a mechanism to recognize the input pattern that can take control of the robot action and avoid the reflex pathway. Both mechanisms can be controlled by a small number of internal motivations that influence the neural activity and the weight modification laws. For instance, a pain signal can provoke an increase of the random activity of the neurons, which allows the robot to quickly escape reflex solutions and to explore the whole action possibilities. In such a phase, the robot seems to be really stressed like a rat in a skinner box when electric shocks are used to force it to discover and to learn a particular behavior. In the same way, pleasure increases the robot vigilance and allows it to learn what seems to have been the cause of the pleasure signal [Gaussier94a], [Gaussier94b].

Prometheus' " brain " is made of two PerAc unit blocks which function exactly in the same manner. The first one performs the visual scene analysis whereas the second one is concerned with target retrieval. In each PerAc block, motor and perceptive flows are processed and recombined through four groups of neurons representing input and output of each flow. We will show that the problem of learning to recognize objects or scenes and next to return to a previously discovered interesting location can be achieved with only two PerAc blocks push-pully connected to each other. We will emphasize the importance of the choice of a coherent neural code that can be applied to code the Prometheus' eye saccades and the direction of Prometheus movements. At last, we will conclude by showing how this architecture can be generalized to other tasks.

2 A Neural code to control autonomous robots

Well known studies in psychology suggests that animals are able to use objects in their environment to locate themselves. These objects are named landmarks. For instance, Morris [Morris75] proposed an experiment in which a rat is trained to swim in a tank toward an invisible platform. Fixed marks on the

walls of the tank are visible from any point in the tank, and they constitute the only information available to the rat for its localization. Other experiments by [O'Keefe 89] show that a rat can find a goal in a maze by using landmarks such as usual objects (a light, a marker pen, a towel ...).
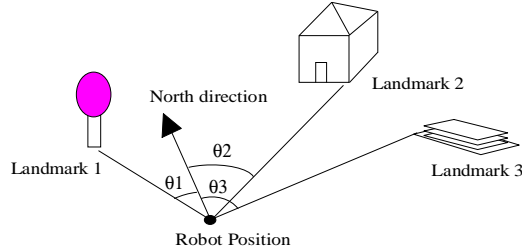

Figure 2: example of a landmarks configuration that the robot can use in a localization task

Moreover, both these experiments show that hippocampus plays a major role in this work. They have found that particular cells in the rat hippocampus respond maximally when the rat is at a particular position and that their activity decreases as the rat is displaced. It seems also that this response does not depend of the rat orientation in its environment. That means, the rat must be able to translate all its visual information in order to present them all the time in the same orientation. It must use something like a switching mechanism that can be modelized by sigma-PI units [Rumelhart86], [Koch85]. The same mechanism could also explain the capability to recognize an object whatever its orientation is. It has been shown that the visual recognition time depends of the angular variation between the learned object and the present object [Farrah88]. So, we can imagine such a switching mechanism that would rotate objects to simplify their recognition and another one that would be usefull to build a scene representation that does not depend on the eye or head position. We have chosen to code the movement command of the robot and its eye in polar coordinates. Each neuron corresponds to a particular movement orientation. For instance, the ocular saccades are represented as vectors associated to a grid of neurons that represents 32 orientations and 32 intensity of possible movements. The direction of the robot movements is also expressed in the same coordinates. That simplifies the connectic problems of linking several neurons groups. Indeed, the retinal image directly provides information to activate an eye saccadic movement in the retinal coordinate and to make possible the goal tracking by the robot itself. The quantization precision is not really important because the use of probabilistic neurons allows to make movements with a precision that depends only on the sampling time. For example, if the robot can only move forward or 90° left or right, it can move in the direction of 25° if the left-neuron is activated randomly twice more than the straight-ahead-neuron. The precision of such a probabilist control can be very efficient and seems to explain human and animal manipulation precision [Georgopoulos89]. Simple reflexes can be easily constructed to control the ocular movement in the direction of " something " in the retinal image and in the same manner to force the robot to move in that direction (fig. 3).
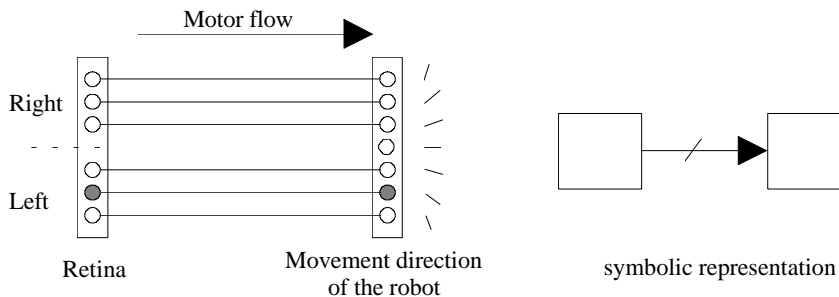

Figure 3: representation of a reflex link in the motor flow of the PerAc architecture

Now, we have to find how to recognize landmarks and to extract angular position from a landmark to another one to build the first part of the robot control.

3 Visual scene recognition

Prometheus' visual system tries to solve this problem by emulating a moving eye. Its own task is to learn several objects and to recognize them in a scene, where they can be scaled, rotated, deformed, occluded or noisy. The first important feature of Prometheus that allows it to solve this problem is that it has a limited vision of the scene. It cannot see all the objects at once. It needs to move it eye from one object to the other. This limitation requires to have a sequential functioning which simplifies learning and recognition. Indeed, the task to locate " where " is the object is now independant of those concerned with deciding " what " the object is. Such a mechanism is consistent with neurobiological data about temporal and parietal areas in the brain which are involved in those specific tasks [Burnod89], [Gilbert83].
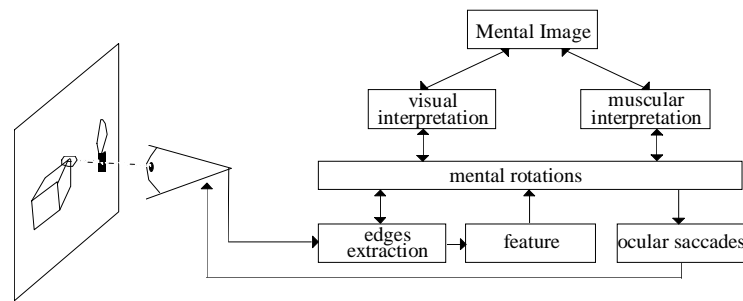

Figure 4 : General architecture of the vision system

Prometheus' " brain " has several cortical areas associated to the primary visual areas, to the visual recognition, to the motor control of the ocular saccades and to the association of visual information and motor information (fig. 4) [Gaussier92a], [Gaussier92b]. We distinguish two connected levels of processing. The first one is involved in low level processing. It is massively parallel. It extracts contours of the image [Grossberg85] which are diffused to obtain local maxima that correspond to the characteristic points that draw the robot attention [Seibert89], [Watt83]. The second one processes a state space transformation of the input picture, i.e., a log-polar transformation [Schwartz80] which is tolerant to rotations and changes of scale but very sensitive to shifts in positions. It is the reason why we prefer the robot's eye to focus its attention on a corner than on the gravity center of objects. For instance, with the second solution, if the object is occluded, the position variation of the object gravity center is terrible and makes recognition impossible. With the first solution, if the robot focuses on angular zone, it risks only to lose a few focus points on all the focus points used to recognize the object. The sequential object exploration is then a good trick to provide redundancy and movement information to help the recognition. To sum up, two data streams are present in Prometheus vision system: one perceptive which identifies the contour image around the focus point, and one motor which guides ocular saccades. Both interact on each other. The scheduling memory of the local recognitions and actions can explain attentional processes that lead us to first explore one possibility before " thinking " to the next one.

The local visual interpretation is performed through a mechanism that simulates the mental rotations. A motor map is used to control the eye movements or the focus of the attention in association with the visual recognition. Both visual and motor data are joined in a kind of " frontal area " where temporal integration is used to recognize sequences. They define a *sub-symbolic* mental representation of the studied object. Recognition of an object consists in using the same scan path during learning and during the utilization phase. It somewhere imitates human behavior in front of an " unknown " object with the simulation of ocular saccades [Norton71] and with a similar recognition time to recognize objects which have been rotated a lot [Farrah88].

Due to those considerations, Prometheus' vision system does not need any complex hierarchical structure to recognize objects. Moreover, the object concept in Prometheus is not linked to the need to analyze a closed region in the image. An object can be composed of several isolated pieces. So a scene with all or a part of its most relevant objects can be considered as a single object. Its recognition will depend on the robot's capability to recall the scan path used during learning to go from one focus point belonging to one piece of object to the next one. We have chosen to use angles between edges as focus points. From the contour image, we use a sort of OFF-Center cell (fig. 5) that provides a maximum response when there is a sharp corner in the neighborhood. A competition mechanism identical to those used to extract edges is used to find feature points at a particular resolution.
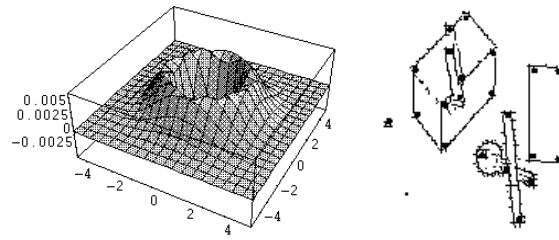

Figure 5: a) The filter used to find corners: it is the difference of gaussian mask: a OFF-Center cell. b) an example of feature point extraction on contour image. Big black dots represent features points

In addition, Prometheus perceives the optical illusion depicted on fig. 6 if we consider that the perception of segment lengths and line discontinuity is due to eye movements from one segment extremity to the other.
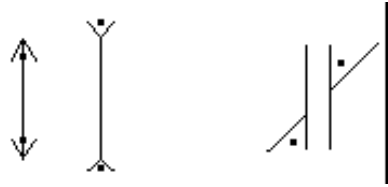

Figure 6: examples of optical illusions explained by boundary diffusion (Muller Lyer's and Poggendorff's illusions)

During training, the robot extracts the characteristic points in the scene and it performs an invariant transformation from each of these points. During the interpretation, the robot focuses its eye on a characteristic point (a corner), it performs an invariant transformation (i.e., a polar logarithmic transformation) and then a mental rotation to match the present target with the learned representation. To complete its interpretation or to remove any ambiguity, the robot focuses on the other characteristic points used during learning according to learned saccadic movements (fig. 7). Objects can thus be recognized in a real scene even if they are partially occluded or rotated or if there is noise. At last, a mechanism of time integration is introduced to simulate a short term memory. Thanks to it, Prometheus will be able to interpret a particular area according to the previous interpretation.
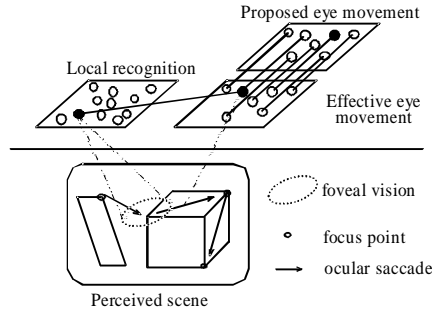
Figure 7: Functioning of Prometheus.Four neural groups are involved (see fig.8). Prometheus focuses its eye on one of the cube's vertices. The ocular saccade it will perform is due to the combined activation of one neuron in the local recognition group and a neuron in the proposed eye movement group. The performed saccade thus corresponds to the one learned when exploring this cube's vertex for the first time.

When a characteristic point has been chosen, an inhibition mechanism prevents the robot from choosing it all the time. However, a problem remains. The points to inhibit are in $\{\log(\rho),\theta\}$ space and when the robot changes its focus points, it loses the origin of the transformation. So, there is a new mapping in the state space. If we perform a simple feedback, it is not the neuron corresponding to the previous mapping which will be inhibited. Consequently, we assume that the brain has a mapping of the picture expressed in rectangular coordinates. This space must be like an *internal universe* and we also need an inverse polar transformation. For details, see [Gaussier92b]. The complete architecture of the vision part is shown on fig. 8. Each arrow represents a link between two groups of neurons. The arrows crossed by one short line represent one-to-one neuron links whereas the arrows crossed by two short lines represent one-to-all neurons links. Commonly, the one-to-one links are reflex pathways (fig. 3) and are considered as unmodifiables as in classical Pavlovian conditionning.
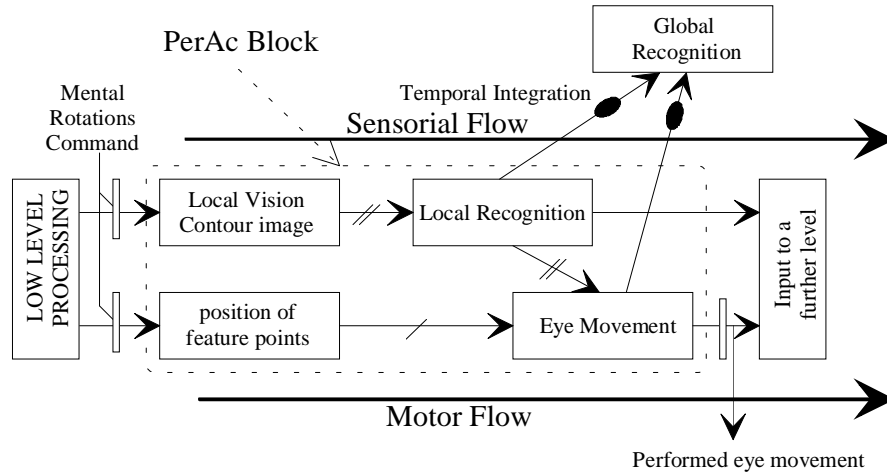

Figure 8: The PerAc architecture for visual scene interpretation. Each block is a group of neuron. There is topology preservation in Local Vision, feature points and Eye Movement groups. Local recognition and Eye movement are WTA.

The eye movement group is a topological map with a WTA (Winner Takes All), with input in the perceptive and motor flow: the position-of-feature-points group proposes a movement, and the local recognition group is associated to a given movement. This latter group has been inplemented with a WTA but a Probabilistic Topological Map (PTM) [Gaussier94a] could be used advantageously for local recognition. There is a global recognition group, which learns with the help of a teacher, and which works according to a counterpropagation algorithm [Hecht-Nielsen87]. However, it does not belong to the studied unit block and it is not necessary to solve the complete robot task.

a)   b)
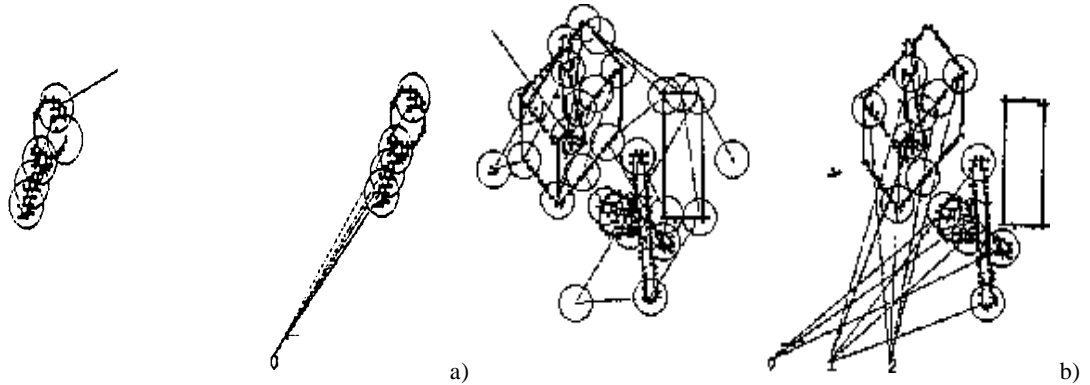
Figure 9: a) Learning of a key labeled as object 0. b) Observation of a complex scene: scan path of the eye (ocular saccades) and interpretation of each zone pointed out

4 Target retrieval using landmarks

The local visual recognition and the information about the ocular movements can be joined to provide information about " what " the landmarks are and " where " they are from each other. Simple product or logical AND neurons can be used to merge those different information type in a map of neurons that reacts only if a particular landmark is recognized at a particular place (fig. 10). Moreover, this model seems to be biologically plausible and to agree one part of the hippocampus architecture [McNauton89], [Zipser85]. A short term memory represented by recurrent positive feedback links is used to obtain a spatial image of the position of the different landmarks in the observed environment from the sequence of input activation. Next, a simple diffusion mechanism allows to use the topological information about ocular movement to match a learned panoramic pattern with the present one even if the angles are not exactly the same.
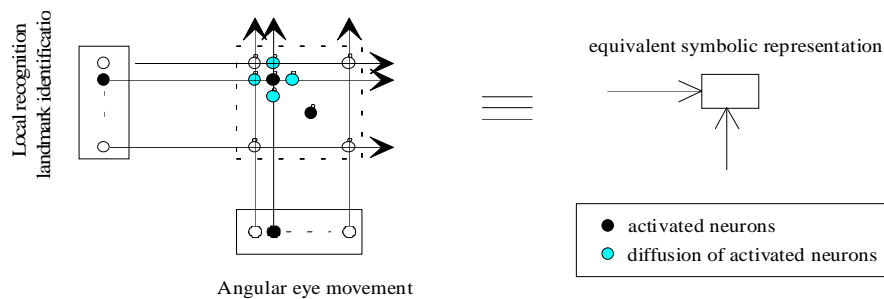

Figure 10: Recombination of visual and motor flow as an input to the Place Fields Cells

At the beginning, we suppose Prometheus moves randomly looking for something interesting. When it finds " food ", it first eats one part of it then it moves around it to learn the food direction from particular locations what is the movement needed to go in the food direction. Later, when the robot wants to find " food ", it considers the information of the place fields associated to the food and goes in the direction associated to the most activated place fields (competitive mechanism). Thus at each time step, the distance to the target is reduced (fig. 11) and it is bounded to return to the learned position of the food.
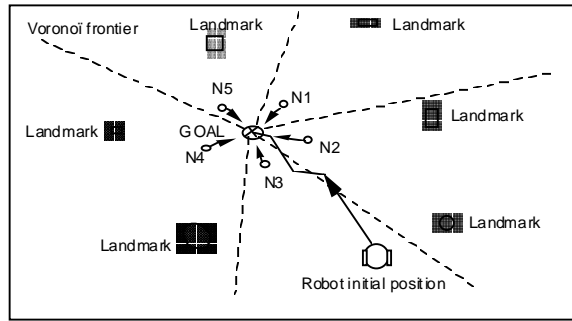
Figure 11: Local exploration around the target represented by the large black circle. The agent records at certain points (represented by small circles) their relative position to the landmarks (represented by squares) and the direction to the target.

The learning phase is more complex because it is an unsupervised and an on-line process. When Prometheus eats " food ", it triggers a reflex which allows to somehow circle around the food at a certain distance, in order to visit evenly placed locations around it. At each of these well chosen locations, a place field learns the relative position of the robot to the landmarks, and the direction heading towards the target (fig. 12). We propose here to detail the neural network used for landmark-based navigation (fig.13). We focus mainly on designing the simplest architecture for the desired behavior. As in the vision part, four neuron groups are involved, all are Topological Maps. Two one-dimensional maps are used to represent movement directions and one two-dimensional PTM for localization, i.e. place cells. There is an internal representation of the world expressed in a referential independent to the robot's orientation with respect to its surroundings. As a consequence, two groups are used for movements, because one must correspond to movement directions with respect to an absolute direction and be associated to localization, while the other corresponds to the movement actually performed by the robot, which means that it takes the robot's orientation into account.
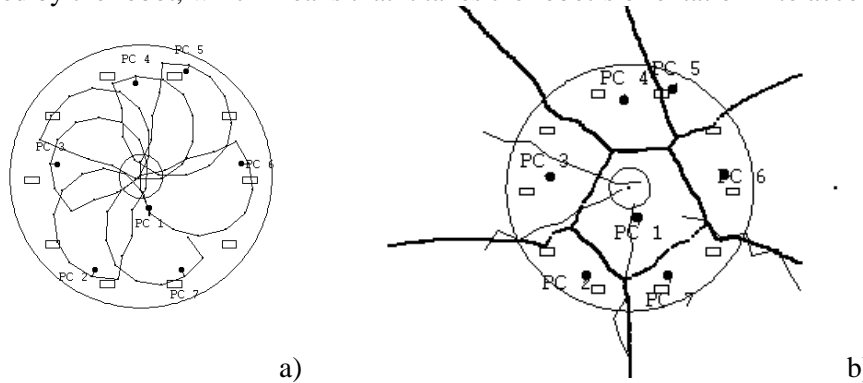

a)   b)

Figure 12: a) Local exploration around the target represented by the large black circle. The robot records at certain points (represented by small circles) their relative position to the landmarks (represented by squares) and the direction to the target. The numbers correspond to the place-field number in its neuron group. b) Different trajectories. The Place-cells (PC) are indexed by their order during exploration. The Voronoi tessellation is represented by the thick lines, the landmarks by the rectangles and the target by the inner circle. The large circle represents the limit beyond which the target is not perceived. Thin lines represent trajectories from various starting-points.

When a movement direction is selected, the robot makes one step of given length in that direction. The input to this network are the north direction, and the food and landmarks positions in the robot's visual space. We assume that a compass is available. It could be replaced by a vestibular system or a gyroscopic mechanism that would produce low precision information about the body orientation (a local landmark could also be used but it would reduce the generalization capabilities of the robot to very distant situations).
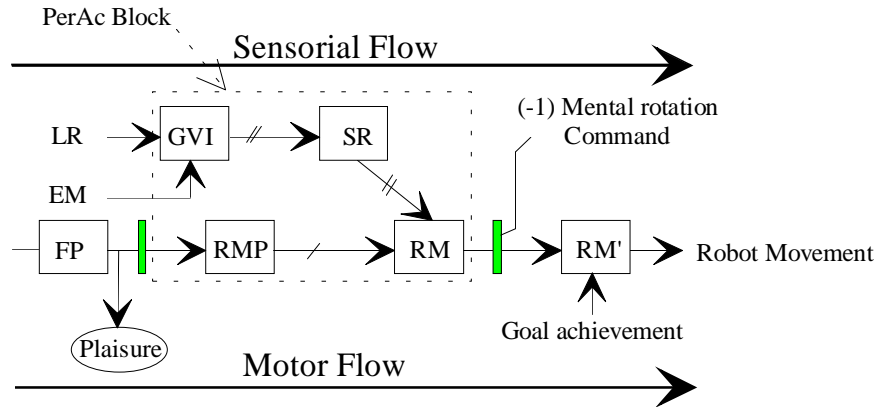
Figure 13: The navigation neural network. SR is the Scene Recognition group. Its input is the Global Visual Input group which corresponds to the Landmarks Recognition associated to the Eye Movement. The Robot Movement group is a WTA. When the food is visible (Food Proposal group), the chosen direction in RM corresponds to the food position, because of high-valued one-to-one links between RMP and RM groups. The RM' group is also a WTA and it corresponds to the Robot Movement in the environment. When Goal achievement is activated. It activates throw a high intensity reflex a particular neuron in RM', causing the robot to turn in a given direction, thus giving rise to ellipsoid trajectories. The black rectangles represent a shifting mechanism used either to provide an invariant representation of the input, or to transform invariant representations into extracorporal ones. Pleasure is emitted when the food is in sight. It works as a chemical substance emitter, and increases learning throughout the whole network.

Just as for humans and most mammals, we assume that the immediate visual angle is limited. Therefore, food is perceived only when it is located in a given orientation ahead of the robot. The same goes for the landmarks, but we assume that when a position must be recorded, Prometheus rotates in order to see in all directions. This supposes that when exploring a scene, it can make ocular saccades and move its head as well, thus spanning the whole surrounding space. The functioning of the neural network is easier to understand by starting from the end, that is the one-dimensional neural map corresponding to the movements. We used two different maps, because the " exploration " reflex must activate a " turn left by a certain angle " from the actual angular position of Prometheus. On the other hand, learning of place cells and of associations must be learned in a fashion which is independent on that position, thus producing an internal representation of the robot's world. Indeed a place cell represents a position in space, and not an orientation of the robot. Therefore, its activation should depend only on Prometheus position, which is obtained by using a shifting mechanism similar to the one described about vision. In the same manner, the first WTA should record directions of movement independent of the robot's direction.

When food is in sight, a neuron corresponding to its angular position relative to the robot's facing position is activated in the Food Proposal Map. The shifting mechanism activates a neuron in the Robot Movement Proposal by adding an angle corresponding to the angle between the robot and the north. If there is pleasure at that moment, a place cell learns the invariant landmarks position, and the association with the movement in RM due to the reflex link from RMP. The inverse shifting mechanism is applied to the output of that group, by substracting the same angle. This activates the neuron in the effective RM' map which corresponds to the actual movement to be performed by Prometheus.

The achievement of the robot goal (to eat food) triggers a movement reflex that remains active for a certain amount of time. The provoked trajectories after reaching food thus take an ellipsoidal shape, which ends after a while. As soon as food is in sight (given a limited visual angle) the position of the landmarks is recorded. This supposes that when pleasure is active, the robot moves its " head " in order to see landmarks in all possible directions.

We have implemented the neural network described above on a Khepera robot. Due to the tremendous computing time required, we simulate part of the input. We assume that the position of landmarks is known, and we compute with wheel-odometry the position of the robot and where it should see the landmarks and food. When this learning phase is over, it becomes possible to launch the robot

from a place where it is not supposed to see the food, and it appears from fig. 11 that it always takes the right direction, whatever its starting point (fig. 12). The distance from the place fields' recorded positions from which the robot can be launched grows with the angular resolution and with the width of the diffusion applied to the input. More realistic trajectories can be obtain if the movement is performed according to a probabilistic vote rather than a determinist WTA mechanism.

5 Conclusion

All these examples lead naturally to the question of the definition of emergence, central to the constructivist paradigm [Maturana87] in which it represents an alternative to the classical cognitivist paradigm [Lakoff87]. At first sight it corresponds to the application of the holistic principle -the whole possesses features that cannot be found in any of its subcomponents-, but there is no clear undiscussed definition of emergence. Nevertheless, we have given several examples of phenomena which exhibit some kind of emergence, as the features of those systems cannot be explained by any of their components.

In learning with neural networks, we have two examples of holistic phenomena. In the vision application, the system in which emergence appears is the vision system plus the image itself. The approach is based on Gestalt Theory, according to which the image as a whole contains more information than its parts. This includes all the possible ambiguities and optical illusions which are not present in subcomponents of the image. Moreover, the optical illusions are due to the vision system performing its operations on the image. This is a good example of a structural coupling between the system and its " environment " cherished by [Maturana87]. All these examples, which by no means pretend to propose a clear definition of emergence, at least show that it is necessary to set up global solutions to cognitive problems. One cannot be content with studying only a function to be approximated, or the behavior of a single processing element, since it cannot be known a priori which role or how important this element should be in the whole system. Moreover, by studying only subcomponents, one loses the opportunity to use the important dynamics of the system [Gaussier94b].

Throughout this paper, we have insisted on the importance of active perception. We have shown that using action simplifies the interpretation of perception: each action is a choice and conditions entirely the future of the robot. The greatest advantage of this type of approach is that it makes cognition sequential, thereby avoiding the possible large duplications and relaxation mechanisms needed by massively parallel systems such as the connexionist systems proposed by Feldman [Feldman85] or by the PDP group [Rumelhart85].

Prometheus proves that a complete autonomous navigation system has no need of an explicit symbolic representation: high level capacities use all the emergent phenomena due to the lower levels. The robot learns to categorize its external world according to what is relevant to " him " and not to us. The information it stores depends only on its action capabilities and on its perception of the world's complexity. Prometheus' " brain " architecture is summarized on fig. 14. The PerAc blocks of which it is made appear to be a kind of basic building block and a systematic tool to combine motor and perceptive information. In addition, the PerAc architecture takes into account the dynamical aspect of the robot's behavior and solves robot control problems in which " autonomy " is needed. Indeed, Perac architecture relies on the postulate that the recognition of any cue can be simplified if the system can act on it. This justifies to cut any perceived cue into two parts: a) a motor part which is the result of a hardwired conventional processing and b) a cognitive one which proposes to learn/record important situationa to allow a quicker adaptation of the system's response.
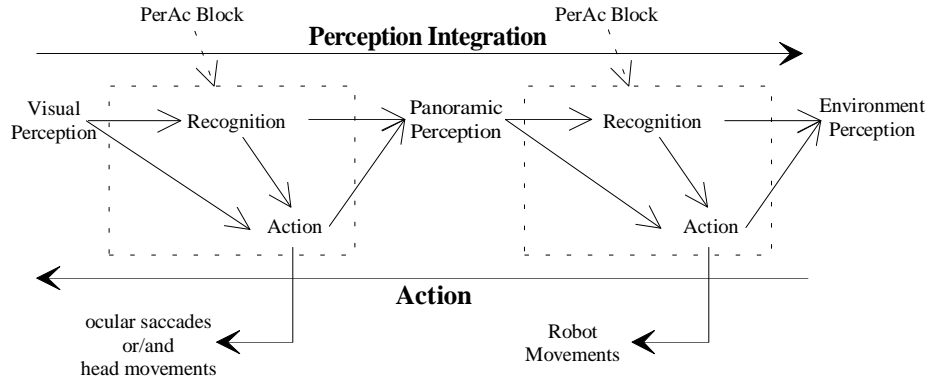
Figure 14: The global neural network for visual scene recognition and navigation. *It is made of two PerAc blocks, the first one for vision and the second for navigation.*

Next, our model proposes an alternative to the classical scheme of hierachical classification because we propose to integrate not only static recognition information but also motor information provided by the input cue or/and the local recognition. For instance, in the recognition problems or in the classification of high dimention data, a commonly accepted method for avoiding to lose topology information consists in classifying local features before taking the results as inputs to higher levels. That constitutes a bottom up architecture with a pyramidal shape: the higher the level is, the less there are nodes to code the more abstract information [Linsker86], [Fukushima82] (fig. 15). From this point of view, the PerAc concept allows to reduce significantly the number of levels from the real world to " sufficiently " abstract levels. For instance, in the vision system, we have only two levels: 1) low-level processing and 2) global object recognition.
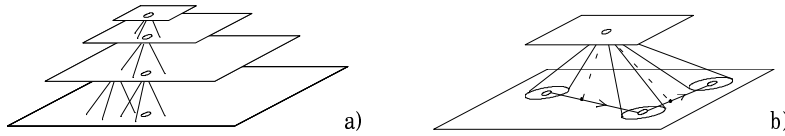

figure 15: a) classical pyramidal structure for hierarchical classification. b) PerAc structure (less levels)

To realize the same kind of task, Fukushima [Fukushima82] needs a number of layers that directly depends on the invariance expected in its image analysis. In PerAc, the reduction of task complexity is due to the *a priori* knowledge we introduce about the nature of the input image and about the relevance of the focus point. But that *a priori* information has nothing to do with the one needed by methods of recognition by modelization. As a matter of fact, here the information can be explained by the ontogenesis of the system and by the fact that we suppose that input have their own topology and that simple competitive/cooperative mechanisms can always be used to locate important features in any perceived cue. Moreover, our model agrees the motor theory of speech recognition which postulates we recognize speech signals by trying to imitate the heard sound. The information used for the recognition is in that case the sequence of articulations needed to imitate the sound. It is obviously more variable than the original sound and it must take into account the mechanical limitations of our phonatory system as well as of our knowledge about the possible succession of actions that could produces intelligible words and sentences.

Clearly, the PerAc or any other neural architecture is nothing without a good model for all the neural groups involved. This is the reason why we believe it is really important to progress simultaneously in the design of interesting neural groups, in our case an on-line learning topological map, and in the architecture to use them. See [Zrehen94] in this same conference. This research process should lead to define an explicit parallel langage to " program " animal robots with adaptation and autonomy capabilities.

Future work will be concerned with finding ways to optimize the architecture parameters, and to extend that kind of networks to more complex tasks, always relying on the constructed level to obtain

the next. A special stress will be put on introducing goal generation and resolution [Burnod89] and to improve the cheap " limbic system " we use for modeling internal motivation. Indeed, it is the element in the robot's artificial " brain " that influence most the overall behavior.

Acknowledgments

References

[Albus91]   Albus J.S. , Outline for a Theory of Intelligence, *IEEE trans. on syst. man and cybern.*, vol 21, No 3, p 473-509, may/june 1991

[Brooks86]   Brooks R.A., A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation* **RA-2**, p 14-23, 1986.

[Burnod89]   Burnod Y. , An adaptive neural network : The cerebral cortex, *Masson,* Paris, 1989.

[Carpenter87]   Carpenter G. & Grossberg S., A Massively Parallel Architecture for a Self-Organizing Neural Pattern Recognition Machine. *CGVIP* **37**, p 54-115, 1987.

[Edelman87]   Edelman G., Neural Darwinism : The Theory of Neuronal Group Selection*, Basic Books,* New-York, 1987.

[Farrah88]   Farrah M., Hammond K. M. , Mental rotation and orientation-invariant object recognition: Dissociable processes, *Cognition*, No 29, 1988, p 29-46, 1988.

[Feldman85]   Feldman J.A, Connectionist models and parallelism in high level vision, *CVGIP*, 31, p 178-200, 1985

[Fukushima82]   Fukushima K., Miyake S., Neocognitron : A New Algorithm for Pattern Recognition Tolerant on Faults and Shifts in Position, *Pattern Recognition* **15**, 6, p 455-465, 1982.

[Gaussier92a]   Gaussier P. , Cocquerez, J.-P., Neural Networks For Complex Scene Recognition: Simulation Of A Visual System With Several Cortical Areas, *in Proceedings of IJCNN*, Baltimore, Vol. 3, p 233-259, 1992.

[Gaussier92b]   Gaussier P., Simulation d'un systeme visuel comprenant plusieurs aires corticales, *Doctoral Thesis*, Paris XI - Orsay, 1992.

[Gaussier94a]   Gaussier P. & Zrehen S., A Topological Neural Map for On-line Learning: Emergence of Obstacle Avoidance in a Mobile Robot, *SAB 94*, Brighton, not yet known, 1994

[Gaussier94b]   Gaussier P. Zrehen S., A Constructivist Approach for Autonomous Agents, *Wiley & Sons*, Ed N. Thalman, 1994

[Gaussier94c]   Gaussier P. Zrehen S., Why Topological Maps are Useful for the Design of Autonomous Agents, *proceedings of PerAc'94*, this volume, 1994

[Georgopoulos89] Georgopoulos A., Neural interpretation of Movement: Role of Motor Cortex in Reaching, *FASEB J.*, 13, p 2846-2857, 1989.

[Gilbert83]   Gilbert C.D., Microcircuitry of the visual Cortex, *Ann. Rev. Neurosci.*, 6, p 217-247, 1983

[Grossberg87]   Grossberg S., Mingolla E., Neural Dynamics of Surface Perception: Boundary Webs,Illuminants, and Shape-from-Shading, *CVGIP No 37*, p 116-165, 1987.

[Hecht-Nielsen87] Hecht-Nielsen R., Counterpropagation Networks. *Applied Optics* **26**, 23, p 4979-4984, 1987.

[Koch&Ullman85] C. Koch C., Ullman S., Shifts in selective visual attention: towards the underlying neural circuitry, *Human Neurobiol*, No 4, 1985, p 219-227

[Lakoff87]   Lakoff G., *Women, Fire and Dangerous Things: What Categories Reveal about the Mind*, The University of Chicago Press, Chicago, 1987.

[Linsker86]   Linsker R., From basic network principles to neural architectures : emergence of spatial-opponent cells. *Proc. Natl. Ac. Sci.* **83**, p 7508-7512, 1986.

[Maturana87]   Maturana & Varela, *The tree of knowledge*, Shambhala ed., 1987.

[McFarland94]   McFarland D., Animal Robotics - From Self-Sufficiency to Autonomy, *this volume/*

[McNauton89]   McNauton B.L., Neural Mechanisms for Spacial Computation and Information Storage. *In Neural Connections, Mental Computation, ed Nadel et al.,* MIT Press, p 285-350, 1989.

[Morris82]   Morris R.G.M et al, Place Navigation impaired in rats with hippocampal lesions, *Nature* 297, June 1982.

[Norton71]   Norton D., Stark L., Eye Movements and Visual Perception , *Scientific American*, vol 224(6), p 34-43, 1971.

[O'Keefe89]   O'Keefe J., Computations the Hippocampus Might Perform. *In Neural Connections, Mental Computation, ed Nadel et al.,* MIT Press, p 225-284, 1989.

[Rumelhart86]   Rumelhart D.E. et al., *Parallel Distributed Processing*, MIT-Press, Cambridge, 1986.

[Seibert89]   Seibert M. & Waxman A.M., Spreading Activation Layers, Visual Saccades, and Invariant Representations for Neural Pattern Recognition Systems, *Neural Networks* **2**, p 9-27, 1989.

[Stewart91]   Stewart, J., Life=Cognition: The epistemological and Ontological significance of Artificial Life, *proccedings of SAB 91 Paris,* Bourgine P. & Varela F. eds, MIT Press, pp 475-483, 1991.

[Schwartz80]   Schwartz L., Computational anatomy and functional architecture of striate cortex: a spacial mapping approach tp perceptual coding, *Vision Res*, Vol 20, 1980, p 645-669, 1980.

[Treisman88]   Treisman A., Features and Objects: The Fourteenth Bartlett Memorial Lecture, *The quarterly journal of experimental psychology*, 40A (2), p 201-237, 1988.

[Van Essen83]   D.C. Van Essen & J.H.R. Maunsell, Hierarchical organization and functional streams in the visual cortex, *TINS*, p 370-375, September 1983.

[Watt83]   Watt R.J. & Morgan M.J., The recognition and representation of edge blur: evidence for spacial primitives in human vision, *Vision Res*, Vol 23, No 12, p 1465-1477, 1983.

[Zipser85]   Zipser D., A Computational Model of Hippocampal Place Fields, *Behavioral Neuroscience* **99**, 5, p 1006-1018, 1985.