Neural Processing Letters **15:** 69–76, 2002.

Emergence of Orienting Behavior in Ecological Neural Networks

PAOLO BARTOLOMEO,[1] LUIGI PAGLIARINI[2] and DOMENICO PARISI[3]
[1]*INSERM Unit 324, Centre Paul Broca, 2ter rue d'Alésia, 75014 Paris, France. e-mail: paolo@broca.inserm.fr;* [2]*Artificialia, Pescara, Italy;* [3]*Institute of Psychology, CNR, Rome, Italy*

Key words: attention, detection, genetic algorithm, orienting, perception

Abstract. We investigated the emergence of orienting behavior in artificial organisms that evolved following a genetic algorithm. These organisms live in a simulated environment containing food and danger elements and reproduce selectively based on the capacity of each individual to eat food while avoiding danger. When the amount of computational resources (number of hidden units) is adequate to the difficulty of the perceptual discrimination between food and danger, peripheral vision is sufficient to trigger stimulus identification. When the resources are scarce, the central portion of the sensory surface becomes a 'fovea', and the presence of a stimulus in peripheral vision triggers an orienting movement (foveation), before the organism can decide whether to eat or to avoid the object. Thus, orienting movements, as well as the segregation of processing resources into a high-definition fovea and a poor-definition periphery, may originate from a disproportion between complex perceptual tasks and (relatively) scarce computational resources.

1. Introduction

Seeing an object 'out of the corner of the eye' often induces movements of the eye and of the head to align the object with the retinal fovea, the region with the highest spatial definition for visual perception [1]. These movements reflect a cognitive process called orienting of attention [2], which is triggered by the sudden appearance of an object in the retinal periphery [3]. Orienting one's attention toward a visual object entails processing the object with increased speed and accuracy. This clearly represents an advantage when a quick decision has to be taken about which objects are to be approached (e.g., food) and which are to be avoided (e.g., danger).

Attention can also be oriented 'covertly', that is, without eye or head movements [2]. Over the last decades, psychological research on attentional orienting has predominantly explored covert orienting. However, covert orienting is mostly an artificial situation, chiefly motivated by the need to study manual reaction time to visual stimuli without the potential confound of the time taken by the eyes and the head to move. The more ecological approach of studying overt orienting has led to important findings in human neuropsychology. For example, a lesion in the right hemisphere of the brain which induces neglect for events occurring on the left side of visual space is also likely to produce a compulsive, 'automatic', orienting of the eyes towards the events presented on the right, non-neglected side [4]. This attentional attraction, coupled with a difficulty in re-orienting attention toward the left [5, 6], can explain why left-sided events are neglected.

Orienting movements of the body can be considered as a prototype of 'embodied' cognitive processes [7]. It is both more accurate and more parsimonious for the cognitive system to go and search the environment for immediately relevant information than to construct a detailed, all-purpose representation of it [8, 9], much in the same way as searching a telephone book spares us the effort of knowing by heart its contents.

The present study focuses on the conditions according to which orienting behavior can emerge in simple evolving neural networks. This modeling approach can shed light on the computational constraints that are at the basis of attentional phenomena. Cognitive neuroscience stems from the convergence of disciplines such as behavioral research, 'wet' neuroscience, and computational modeling [10]. Research on attention has particularly benefited from such an interdisciplinary approach [11].

In the present study, we used ecological neural networks [12], that is, neural networks that live in a physical (simulated) environment. Unlike 'classical' neural networks [13], ecological networks receive inputs that are not arbitrarily decided by the researcher but depend on the environment. By moving in the environment, an ecological network can influence what input it will receive from the environment which is the core of orienting behavior. In the simulations reported in this paper we used a 'genetic algorithm' [14, 15] to study the emergence of orienting behavior in artificial organisms as a function of the perceptual difficulty of the task and the amount of processing resources made available to the organisms. Our results suggested that that the emergence of a 'fovea' and, as a consequence, of orienting behavior, is due to the necessity of processing complex visual stimuli relative to the limited neural processing resources available.

2. Method

A simple artificial organism lives in an environment that contains randomly distributed food and danger elements. To reproduce itself, the organism must be able to reach and eat the food elements and to avoid eating the danger elements. Since only the best individuals reproduce, and offspring are similar but not identical to their parents because random mutations tend to modify the inherited genotypes, the selective reproduction of the best individuals and the random mutations result in the evolutionary emergence of efficient food-finding and danger-avoiding behavior.

The environment is a bidimensional grid of $100 \times 100 = 10\,000$ cells and it contains 200 randomly distributed elements, half of which are food and half dangers. Each individual organism lives in its own copy of the environment and both the organism and a single food or danger element are the size of a single cell. By moving in the environment the organisms must reach and eat the food elements and at the same time it must avoid eating the danger elements. At the end of their life each individual is assigned a ‘fitness’ score, which is the number of food elements minus the number of danger elements eaten by the individual during its lifetime. Lifetime is identical for all individuals (5000 cycles) and is composed of 100 epochs of 50 cycles each. At the beginning of each epoch, both the organism and the 200 food and danger elements are randomly positioned in the environment. A generation is made up of 100 individuals. Only the 20 individuals with the highest fitness reproduce by generating 5 offspring each. The $20 \times 5 = 100$ offspring are the next generation. We have run 5 replications of 3 different simulations (see below) by allowing the genetic algorithm to run for 2000 generations in each simulation.

The behavior of an organism is controlled by a simple neural network with 3 input units, 2 output units, and 5 or 15 internal units (according to the simulation). The 3 input units encode the presence and nature of food or danger elements near the organism. At any given time an organism has a heading. One of the 3 input units encodes the presence and nature of a food or danger element in the cell just in front of the organism. The other two input units encode the presence and nature of food or danger elements in the two cells respectively to the right and to the left of the organism’s body. The 2 output units encode 4 possible actions: move one cell forward (11), turn to the left (10), turn to the right (01), eat the food or danger element in the front cell, if the cell contains a food or danger element (00).

The organism’s behavior depends on the connection weights of its neural network. The connection weights of each individual organism are encoded in the genotype that the individual inherits from its single parent at birth. An initial population of 100 genotypes, for the 100 individuals of the first generation, is generated randomly with weight values ranging from -1 to $+1$. The genotypes of the 20 best individuals are inherited by the individuals constituting the second generation with the addition of random mutations. On the average, 10% of the weights inherited by each individual are modified by adding a quantity randomly selected in the same range between -1 to $+1$ to the weight’s current value.

Two variables are manipulated in the simulations: the perceptual complexity of the food and danger elements and the amount of computational resources available to the organisms. In the *low complexity* condition all food elements are identical and they are perceptually encoded by an activation value of 0.1 in the relevant input unit. The same is true for danger elements that are all perceptually encoded by an activation value of 0.9. In the *high complexity* condition, there are two types of food elements and two types of danger elements. Half of the food elements are encoded by an activation value of 0.1 and half by an activation value of 0.3. Similarly, half of the danger elements are encoded by an activation value of 0.7 and half by an activation value of 0.9. Hence, in the high complexity condition the neural network must classify two different activation values as belonging to the same class of either food or danger elements.

The variable 'amount of computational resources' is manipulated by varying the number of internal units in the organism's neural network. In the condition of *few computational resources* the internal units are 5 and in the condition of *many computational resources* the internal units are 15.

The two variables, perceptual complexity and computational resources, were combined in three different simulations:

(1) Low perceptual complexity and few computational resources
(2) High perceptual complexity and few computational resources
(3) High perceptual complexity and many computational resources.

3. Results

In all three simulations the genetic algorithm was able to find connection weights for the organisms' neural networks that resulted in the appropriate discriminative behavior of eating food elements and not eating danger elements. At the end of the simulations, the organisms tended to respond with an output of 00 (eating) to an input which encoded the presence of a food element in the front cell, and they responded with other outputs (move forward, turn either to the left or to the right) to an input encoding the presence of a danger element. Thus, the organisms could discriminate food from danger elements when a food or danger element was perceived frontally, i.e., with the central portion of their perceptual field.

More crucial to our study was the organisms' behavior when a food or a danger element was perceived in one of the two lateral cells, that is, in the periphery of the perceptual field. If the organism was able to distinguish food from danger elements when the stimulus was perceived peripherally, the organism would turn toward the stimulus if it was food (in order to eat it) whereas it would move forward or turn to the opposite side if the laterally perceived element was danger. In contrast, if the organism was able to distinguish food from danger elements centrally but not peripherally, the organism would turn toward a peripherally perceived element whether it was food or danger in order to 'foveate' it, and be able to recognize it as food or danger. In other words, in these circumstances the organism would exhibit orienting behavior.

To analyze the determinants of the occurrence of orienting behavior, we considered separately the results of the three different simulations.

SIMULATION 1 – LOW PERCEPTUAL COMPLEXITY AND FEW COMPUTATIONAL RESOURCES

At the end of the training period (after 2000 generations of the genetic algorithm), the organisms could effectively explore the environment and eat the food elements while avoiding the danger elements. The organisms, however, did not exhibit orienting behavior. They were able to determine whether a perceived element

Table I. Responses of the organisms to a danger presented in one lateral input cell.

	Number of times the organisms turns toward the stimulus	Number of times the organisms does not turn towards the stimulus
Simulation 1	7	43
Simulation 2	63	27
Simulation 3	20	73

was food or danger (and to react appropriately) both when the element was perceived frontally and when it was perceived laterally. Hence, they did not have to turn towards a laterally perceived element in order to ascertain if it was food or danger. This conclusion was supported by quantitative analysis of the organisms' behavior (see Table I).

We examined the behavior of the five best individuals of the last generation from each of the five replications of the simulation.

We tested these 25 individuals by providing each of them with 50 inputs and recording their output. The 50 inputs are encodings of laterally perceived dangers (2 dangers, one on the left and one on the right, $\times 5$ individuals $\times 5$ replications of the simulation). When these individuals perceived a danger element laterally they responded by turning to the danger only 14% of the time on average. The remaining 86% of the time they just moved ahead. In other words, in most of the cases they did not have to orient toward a laterally perceived danger element in order to recognize it as a danger because they could already recognize it as such with peripheral vision.

SIMULATION 2 – HIGH PERCEPTUAL COMPLEXITY AND FEW COMPUTATIONAL RESOURCES

In this simulation we were also able to evolve organisms that explored the environment efficiently and eat the food elements but not the danger elements. However, in this simulation we observed the emergence of orienting behavior. When the organisms perceived an element laterally (peripherally), they could not determine whether it was food or danger, and they had to turn towards the element in order to perceive it frontally (centrally). As soon as the element is 'foveated', it is recognized as food or danger and the organisms responded as appropriate, that is, by eating the element if it was food and by refraining from eating it if it was danger. The organisms of this simulation turned toward a laterally perceived danger 63% of the time on average (in contrast with 14% in the preceding simulation). Thus, in most cases they were not able to recognize a peripherally perceived element as food or danger, but they had to orient towards it in order to perceive it centrally and recognize it.

SIMULATION 3 – HIGH PERCEPTUAL COMPLEXITY AND MANY COMPUTATIONAL RESOURCES

The purpose of this third simulation was to examine whether the high level of perceptual complexity *per se* causes the emergence of orienting behavior or whether this depends also on the amount of computational resources available to the organism. When the best individuals of this simulation were tested in the same conditions of the two preceding simulations, they only had to orient toward a laterally perceived danger 27% of the cases in order to recognize it as a danger. Thus, although the elements present in the environment were perceptually complex, the organisms in Simulation 3 had enough computational resources in their nervous system (i.e., 15 hidden units) to permit the perceptual discrimination between food and danger already in the peripheral visual field.

The difference in the number of orienting responses observed in the three simulations was statistically significant [$\chi^2(2) = 40.48$; $P < 0.0001$]. Post-hoc tests indicated that this difference depended on the high number of orienting reactions observed in Simulation 2; Simulation 1 vs. Simulation 2, $\chi^2(1) = 30.22$; $P < 0.0001$; Simulation 2 *vs.* Simulation 3, $\chi^2(1) = 24.75$; $P < 0.0001$; Simulation 1 *vs.* Simulation 3, $\chi^2(2) = 2.52$; P n.s.

4. Discussion

We have been able to evolve a form of orienting behavior in simple artificial organisms living in a simple environment. In Simulation 2 (high perceptual complexity, few neural resources), when an element is perceived laterally (peripherally) the organism is unable to decide if the perceived element is food or danger and therefore it turns (orients) toward the element in order to perceive it centrally and recognize it as food or danger.

Orienting behavior seems related to a dishomogeneity in the perceptual field of organisms. Some portions of a perceptual field have more perceptual capacity; in general, very abstract, terms we can call these portions the ‘fovea’, while other portions can be called the ‘periphery’. When an input falls on the periphery of some perceptual field, the organism responds by movements that cause the input to fall on the fovea. In this way the organism is able to extract more information from the input than the information obtainable from the periphery. These movements constitute the orienting behavior.

Why do the perceptual fields of organisms tend to be dishomogeneous and to organize themselves into fovea and periphery? Our results suggest that these characteristics emerge when the organisms’ neural system must be able to process complex stimuli with limited computational resources. Our simulations show that both when stimuli are simple and computational resources (number of hidden units) are limited (Simulation 1), as well as when stimuli are complex and computational resources are more abundant (Simulation 3), the perceptual field of the organisms is homogeneous and there are no orienting responses. In both cases the internal computational resources needed to process the input are sufficient for the level of complexity of the stimuli and they can support a homogeneously powerful perceptual field. Whatever the portion of the perceptual field on which the stimulus is falling, the neural system can recognize a food or danger with no need for orienting movements. This is because the neural processing resources are proportional; either the stimuli are simple and the computational resources are few, or the stimuli are more complex and there are more computational resources.

On the other hand, when the stimuli are complex and the available computational resources are few, as in the present Simulation 2, dishomogeneity in perceptual fields becomes an emergent property. The resources are too scarce for the neural system to be able to extract all the needed information from stimuli falling on all portions of the perceptual field. Therefore, the perceptual field divides itself into two parts. One part, the fovea, uses most of the computational resources and thus is able to recognize a perceived element as either food or danger. The remaining portions of the perceptual field, the periphery, can only detect the presence of a stimulus, but it cannot recognize the element as food or danger. Hence, orienting movements become necessary. The organism responds to the peripheral input with movements that cause the stimulus to fall on the fovea. At this point the input can be classified as food and danger and the organism can respond appropriately.

Why is it that in our organisms, and apparently also in real organisms, the fovea is the central portion of the perceptual field, and not any other portion? A possible answer is that when an element is located in front of the organism, it is able to produce the adaptively critical behavior, i.e., the act of eating the element if it is food and of avoiding eating it if it is danger. Hence, it is evolutionarily advantageous for the frontal (central) portion of the perceptual field to become the fovea and to leave the role of periphery to the two lateral portions of the field. This hypothesis could be tested by dissociating direction of movements and direction of eating, i.e. by evolving organisms that can only eat food when the food is placed laterally rather than frontally while at the same time moving ahead, i.e., frontally.

Finally, one should evaluate whether the limitation of computational resources can be the only explanation for the emergence of the fovea and of orienting behavior. There might be other explanations such as the fact that dishomogeneous perceptual fields, with one portion much more perceptually able than the other portions, automatically translate into a mechanism of focal attention. A mechanism of focal attention can be advantageous because in most instances we can respond with our motor organs only to one among the many objects present in our perceptual field at any given time [16].

References

1. Sokolov, E. N.: Higher nervous functions: the orienting reflex, *Annu. Rev. Physiol.*, **25** (1963), 545–580.

2. Posner, M. I.: Orienting of attention, *Q. J. Exp. Psychol.*, **32** (1980), 3–25.
3. Yantis. S.: Attentional capture in vision, In: A. F. Kramer, G. H. Coles, G. D. Logan, (eds.) *Converging Operations in the Study of Visual Selective Attention*, pp. 45–76, American Psychological Association: Washington, DC, 1995.
4. Gainotti, G., D'Erme, P. and Bartolomeo, P.: Early orientation of attention toward the half space ipsilateral to the lesion in patients with unilateral brain damage, *J. Neurol. Neurosurg. Psychiatry*, **54** (1991), 1082–1089.
5. Bartolomeo, P.: The novelty effect in recovered hemineglect, *Cortex*, **33** (1997), 323–332.
6. Posner, M. I., Walker, J. A., Friedrich, F. J. and Rafal, R. D.: Effects of parietal injury on covert orienting of attention, *J. Neurosci.*, **4** (1984), 1863–1874.
7. Ballard, D. H., Hayhoe, M. M., Pook, P. K. and Rao, R. P. N.: Deictic codes for the embodiment of cognition, *Behav. Brain Sci.*, **20** (1997), 723–767..
8. Ballard, D. H.: Animate vision, *Artificial Intelligence*, **48** (1991), 57–86.
9. O'Regan, J. K.: Solving the 'real' mysteries of visual perception: the world as an outside memory, *Canadian Journal of Psychology*, **46** (1992), 461–488.
10. Kosslyn, S. M. and Koenig, O.: *Wet Mind: The New Cognitive Neuroscience*. The Free Press: New York, 1992.
11. Parasuraman, R.: The attentive brain: issues and prospects, In: R. Parasuraman, (ed.) *The Attentive Brain*, pp. 3–15, The MIT Press: Cambridge (MA), 1998.
12. Parisi, D., Cecconi, F. and Nolfi, S.: Econets: Neural networks that learn in an environment, *Network*, **1** (1990), 149–168.
13. Rumelhart, D. E. and McClelland, J. L.: *Parallel Distributed Processing*. Explorations in the Microstructure of Cognition. 1986, 000–000.
14. Holland, J. J.: *Adaptation in Natural and Artificial Systems*. University of Michigan Press: Ann Arbor, Michigan, 1975.
15. Mitchell, M.: *An Introduction to Genetic Algorithms*. The MIT Press: Cambridge (MA), 1997.
16. Allport, D. A.: Visual attention, In: M. I. Posner, (ed.) *Foundations of Cognitive Science*, MIT Press: Cambridge, MA, 1989, pp. 631–687.