

Synthetic animals in synthetic worlds

Agnès Guillot and Jean-Arcady Meyer

Ecole Normale Supérieure
Groupe de BioInformatique - URA 686
75230 Paris cedex 05, France

E-mail: guillot@wotan.ens.fr, meyer@wotan.ens.fr

Abstract

This paper provides an overview of various research efforts that aim at designing adaptive animats, i.e. synthetic animals able to survive in synthetic worlds. The control architectures of these animats are either programmed by a human designer or, more or less, automatically determined by means of nature-mimicking processes such as learning, evolution and development. The paper concludes with a brief discussion of the directions in which future animat research should be oriented.

Keywords

ANIMATS – CONTROL ARCHITECTURE – DEVELOPMENT – EVOLUTION – LEARNING

1 Introduction

Synthetic worlds are inhabited by synthetic animals - or *animats* - which live, feed, reproduce and die in these worlds. Animats are equipped with sensors, with actuators and with control architectures that endow them with more or less sophisticated adaptive abilities. They are at the heart of a significant number of research programs [CLIF94, MEYE91a, MEYE93] in both fundamental and applied perspectives. On the one hand, indeed, the synthesis of animats can be expected to help in understanding how real animals manage to survive in real worlds. On the other hand, the corresponding mechanisms and working principles may prove to be effective in devising truly autonomous and adaptive robots.

The literature on animats falls into four broad categories [GUIL94, MEYE91b, MEYE94], according to whether it concerns animats whose control architectures have been *programmed* by a human designer, or animats whose behavioral plasticity is due to biologically-inspired automatic processes - like the processes of *learning*, *evolution* or *development*.

In order to illustrate the diversity and the originality of the animat approach, this paper will describe several research efforts drawn from the corresponding literature. It will conclude with a brief discussion of the directions in which future work should be oriented.

2 Preprogrammed Behaviors

Many animats exhibit adaptive behaviors because they have been purposely preprogrammed by a human designer. The work of Beer [BEER90], for instance, belongs to the field of "neuroethology", which aims at reproducing as faithfully as possible available knowledge about the

nervous systems of actual animals. Beer has elaborated a model enabling an artificial cockroach to display sequences of behaviors that ensure its survival in a simple simulated environment - a rectangular area with obstacles, walls and food patches. The whole control architecture of the insect consists of four interconnected neural networks that control locomotion, edge-following and both appetitive and consummatory behaviors involved in feeding (Figure 1).

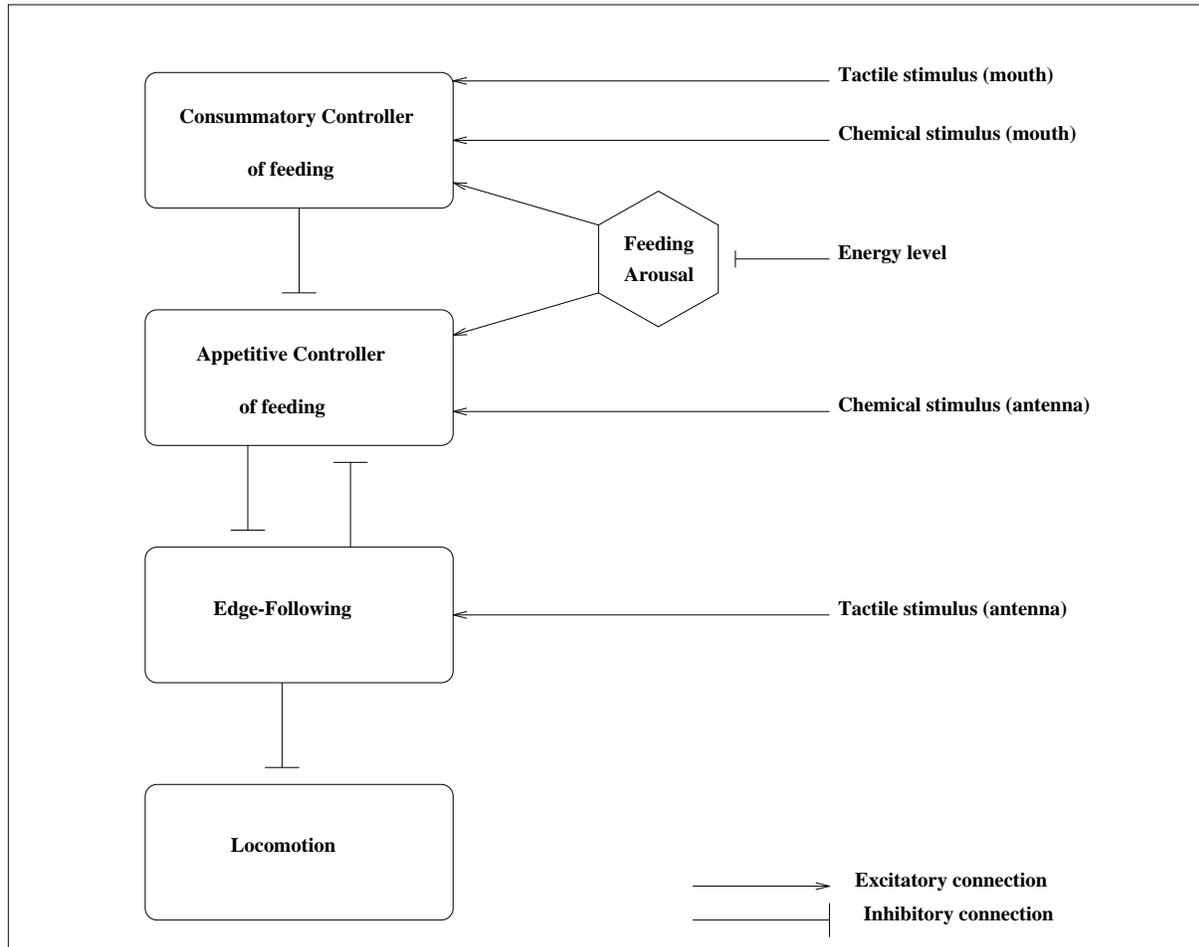


Figure 1. Global control architecture of Beer's artificial insect. After [BEER90]

Neurons correspond to sensory units, motor units, motivation units or interneurons. The network depicted on Figure 2a, for instance, ensures the insect's locomotion and governs the rhythmic motion of its legs. This network calls upon three motor neurons: the neurons involved with *stance* and *swing* determine how forcefully the leg is propelled forward or backward, while the *foot* motor neuron determines whether or not the foot is set down. Motion's periodicity is due to a *pacemaker* neuron P, and the force applied in each stance phase, together with the periodicity of P discharges, depends on a general level of excitation controlled by the *command* neuron C. The sensors essential to the operation of such a network are two neurons that emit a signal whenever a leg reaches an extreme angle. Lastly, a central connection between the pacemakers (Figure 2b) synchronizes the movements of the insect's legs, thus guaranteeing its

stability. In addition, other neurons not described here allow the insect to make use of its motor equipment for specific purposes, such as avoiding obstacles.

Another neural network enables the insect to reach food patches when it is hungry (Figure 3a). The odors detected by the chemical sensors located on each antenna (ACS) are compared by two neurons, LOS and ROS. The difference detected is used to generate a rotation towards the strongest odor caused by the excitation of an appropriate interneuron -LT or RT- governing the lateral extension of the front legs. When the energy rate of the insect decreases, the activity of an *energy sensor* neuron (ES) also diminishes, thus disinhibiting a *feeding arousal* neuron (FA) that otherwise would be spontaneously active. This neuron then excites a *search command* neuron (SC), that will decide whether or not to head for food.

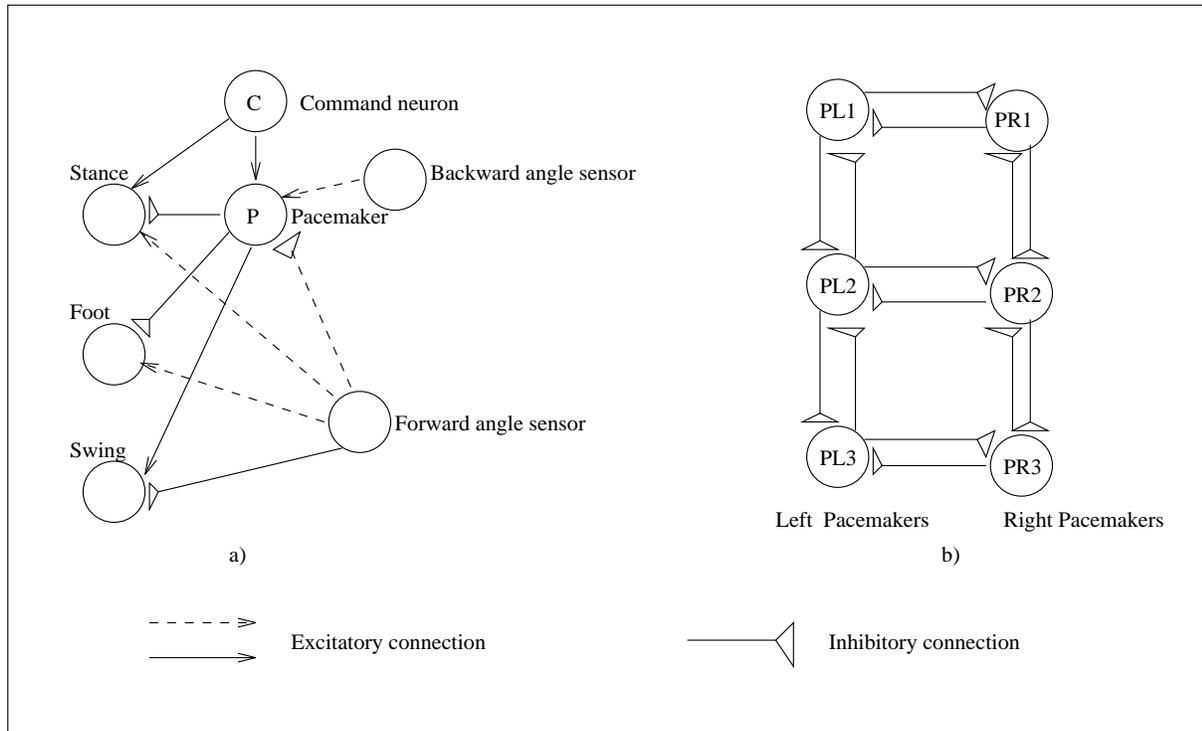


Figure 2. Beer's artificial insect circuits. a) leg controller. b) central coupling between pacemakers. After [BEER90]

A third network (Figure 3b) governs food ingestion. When the chemical (MCS) and tactile (MTS) sensors in the mouth indicate that food is present (FP), and when the insect is motivated enough to feed (FA), the *consummatory command* neuron (CC) is activated and forces the *pacemaker* neuron (BP) to produce rhythmic signals that make the *motor* neuron (MO) open and shut the mouth. When the energy rate increases, the activation level ES inhibits FA, which in turn ends up by suppressing the activity of BP and causes the feeding to stop. Moreover, a positive feedback loop between FA, PB and MO modulates realistically the frequency of chewing during a meal.

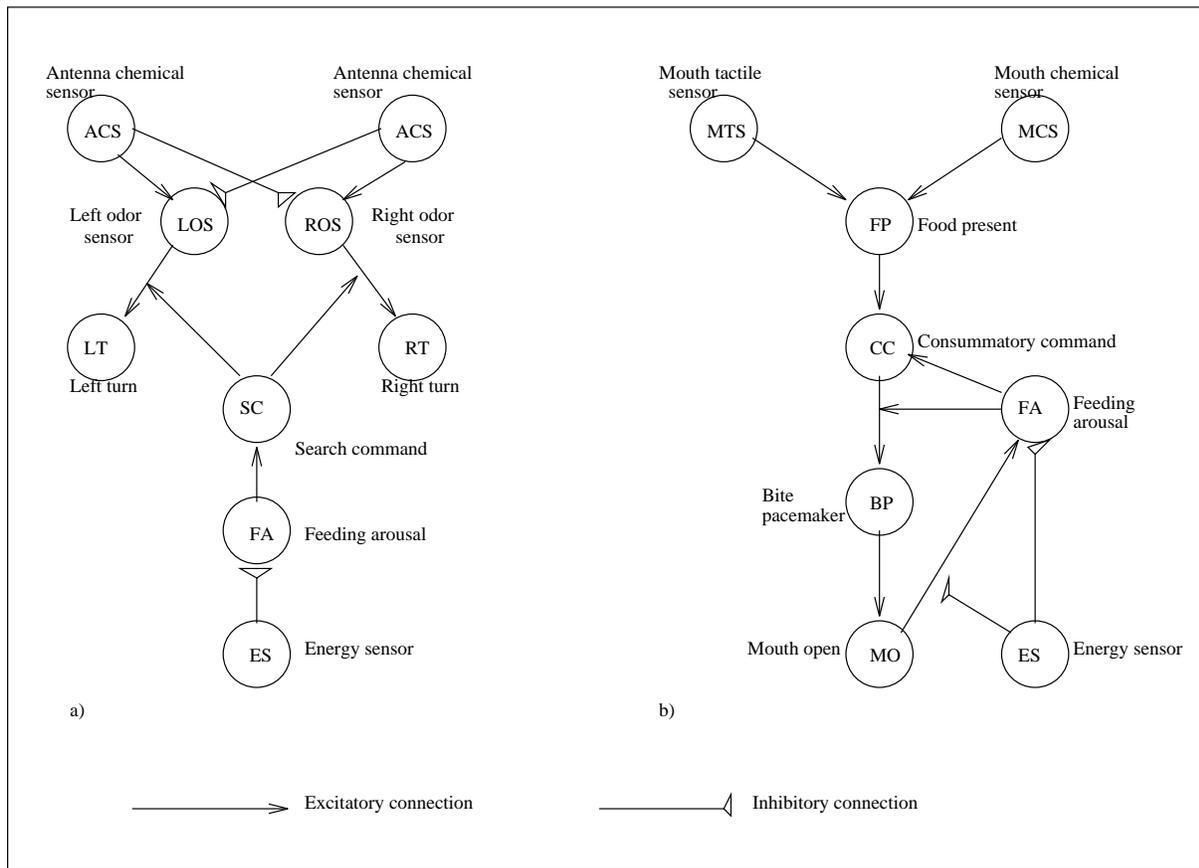


Figure 3. Beer's artificial insect circuits. a) appetitive controller of feeding. b) consummatory controller of feeding. After [BEER90]

Considering that the same neurons are involved in the initiation and control of these behaviors, a specific organization must preclude the simultaneous occurrence of incompatible acts - i.e., those calling upon the same motor units. The solution retained is a *hierarchical organization*, where the consummatory part of feeding takes precedence over the orientation towards food, which in turn is dominant with respect to the obstacle avoidance (Figure 1). According to such an organization, exploration is the behavior engaged in by default, while locomotion is activated in the course of every behavior entering into this hierarchy.

This architecture allows an artificial insect to perform realistic successions of behaviors such as those illustrated in Figure 4: at point A, the insect detects food, but cannot reach it; instead of staying there, it edge-follows the obstacle (point B); at point C, it loses the food odor, wanders and edge-follows the wall (point D); at point E, it detects food again and finally reaches it (point F).

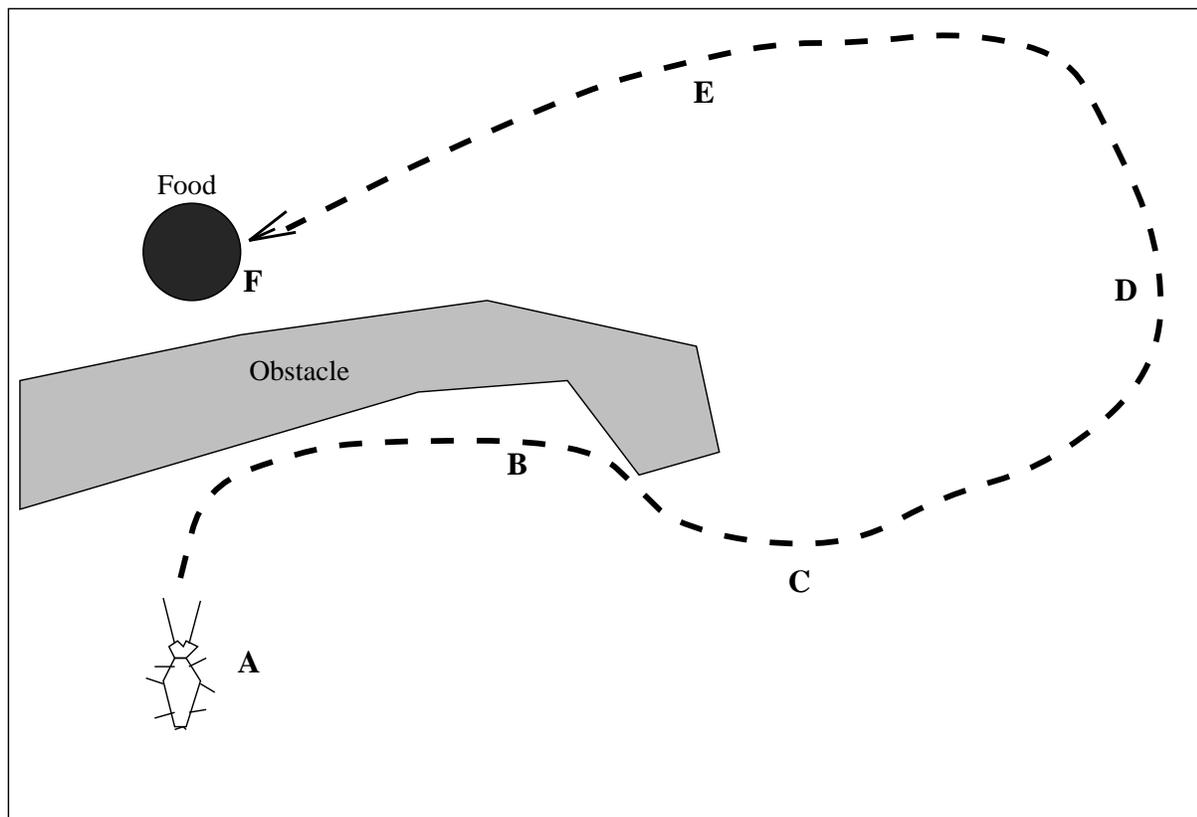


Figure 4. An adaptive behavioral sequence of Beer's artificial insect. After [BEER90]

This work demonstrates that adaptive behavioral sequences can emerge in an animat endowed with a very simple nervous system connected in a clever way by its designer. The insect architecture is totally distributed, but not uniformly: it must be structured so as to ensure an adapted alternation between behaviors in response to changes in the animat's internal state or in the external environment.

3 Learned Behaviors

To reduce the role of a human designer to a minimum, many studies address the way animats can autonomously improve the adaptiveness of their behaviors while experiencing new situations in their environment. In particular, in the situation of *reinforcement learning*, the animat has to discover, by trial and error, how to coordinate its actions in order to maximize a cumulative reward over time. Such a reward is either a positive or a negative signal coming from the environment.

The Barto and Sutton model [BART81] allows an animat to learn to orient itself in a two-dimensional environment, using odorous landmarks . This environment contains one central landmark (C) surrounded by four others (N,S,W,E). Each landmark emits an odor the gradient of which decreases with distance. The task of the animat is to learn how the four peripheral landmarks are associated to the central one, and to navigate toward this goal even if it ceases to emit its odorous signal. The architecture of the animat is a neural network with special neurons inspired from Klopff's "hedonistic" neuron concept [KLOP80]. As an input, this network receives a combination of four odorous signals (XN, XS, XW, XE) associated with the peripheral

landmarks (N, S, W, E), and furnishes, as an output, a combination of motor signals (YN, YS, YW, YE) associated with the four spatial directions. The reinforcement signal Z corresponds to the odour of the central landmark C, a signal that the animat seeks to maximize (Figure 5).

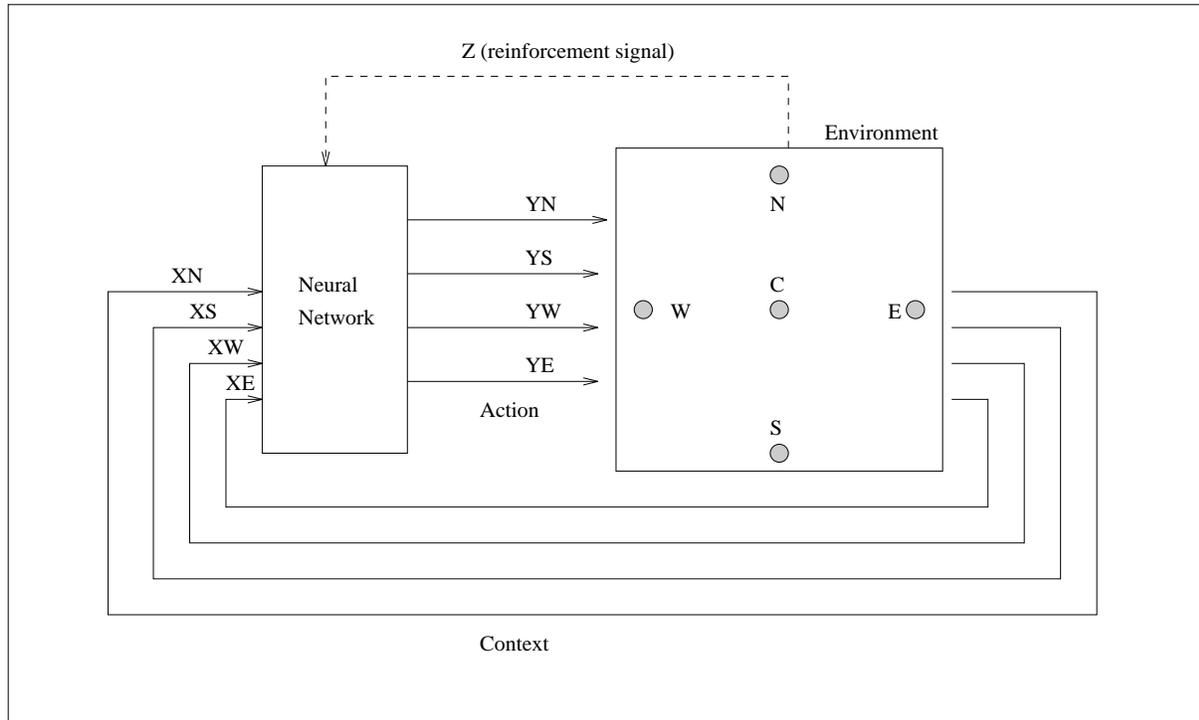


Figure 5. Neural architecture of Barto and Sutton’s animat interacting with its environment. After [BART80]

Learning takes place as the animat moves in the environment. The neurons’ synaptic strengths are updated so that, when the activation of a motor neuron at a given spot brings about a motion in a direction where Z increases, this neuron will have a better chance of being activated on the same spot in the future. Conversely, a motion in a direction where Z decreases will decrease the probability of activating the corresponding neuron.

After training, the animat gradually learns what direction it must move in from any given point in order to reach the goal. It thus can find the goal from any starting point by selecting the correct displacement at each intermediate point. This kind of adaptive ability is known as *route-navigation* [GALL90].

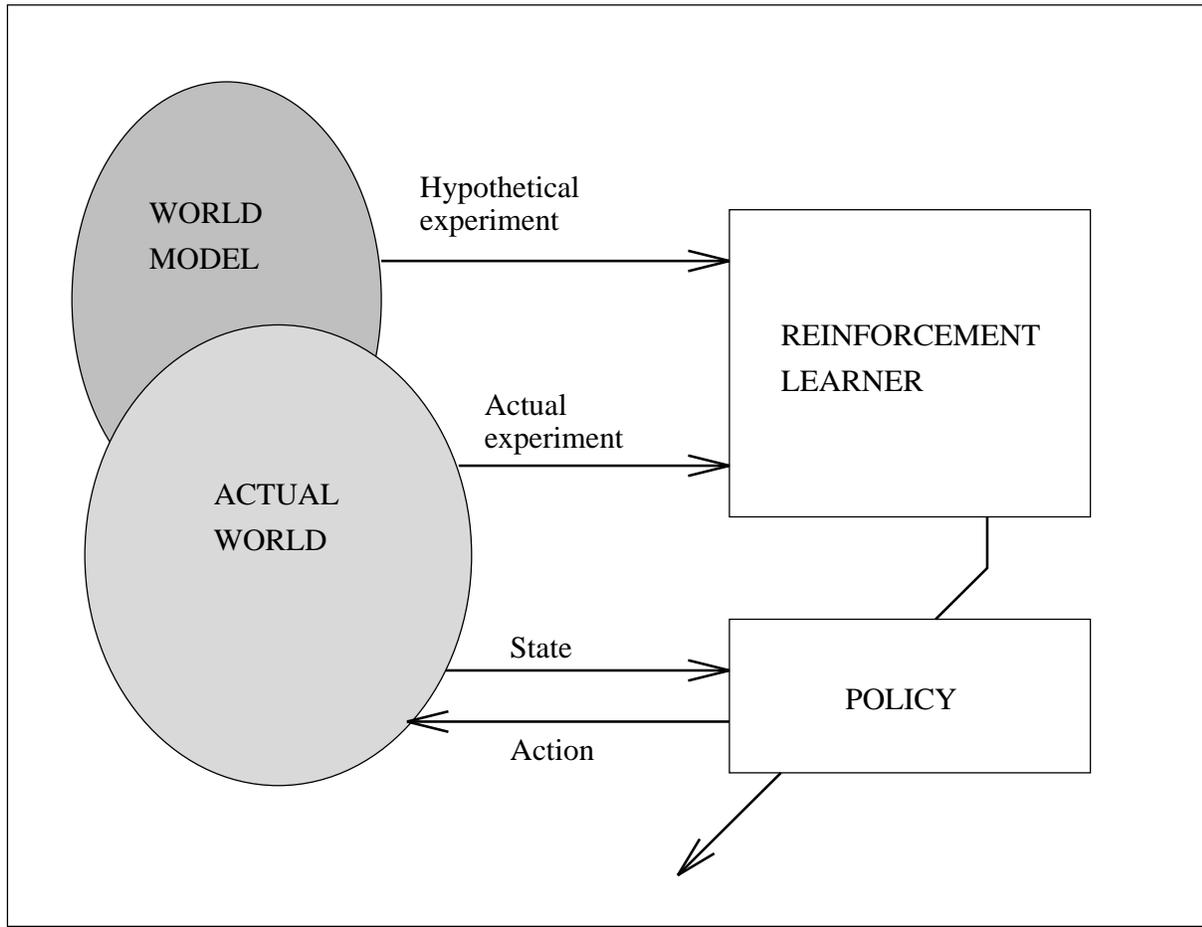


Figure 6. Dyna architecture. After [SUTT91]

Barto and Sutton’s animat illustrates a typical *behaviorist* architecture: it doesn’t need to build an internal representation of the environment to be able to navigate adaptively. On the contrary, the DYNA architecture proposed by Sutton [SUTT91] includes a world model and adds some planning abilities to a reinforcement learning process. DYNA is composed of four modules (Figure 6):

- the real environment, that changes state in relation with the animat’s movements and that distributes a reward signal,
- the internal world model, that the animat elaborates for itself and that is intended to represent the one-step input/output transitions of the real world,
- the policy function, relied on by the animat to determine what action to initiate in response to each possible state of the real environment,
- a primitive reinforcement learner that improves the policy function over time.

The world model and the policy function are progressively modified as the animat experiments with the operational laws of its world. These modifications depend on two types of experiments that the animat may alternate between: *actual* experiments, carried out on the real environment and *fictitious - or hypothetical-* experiments, that make use of the internal world model. Such hypothetical experiments endow the animat with planning abilities and make its behavior depend on its expected consequences.

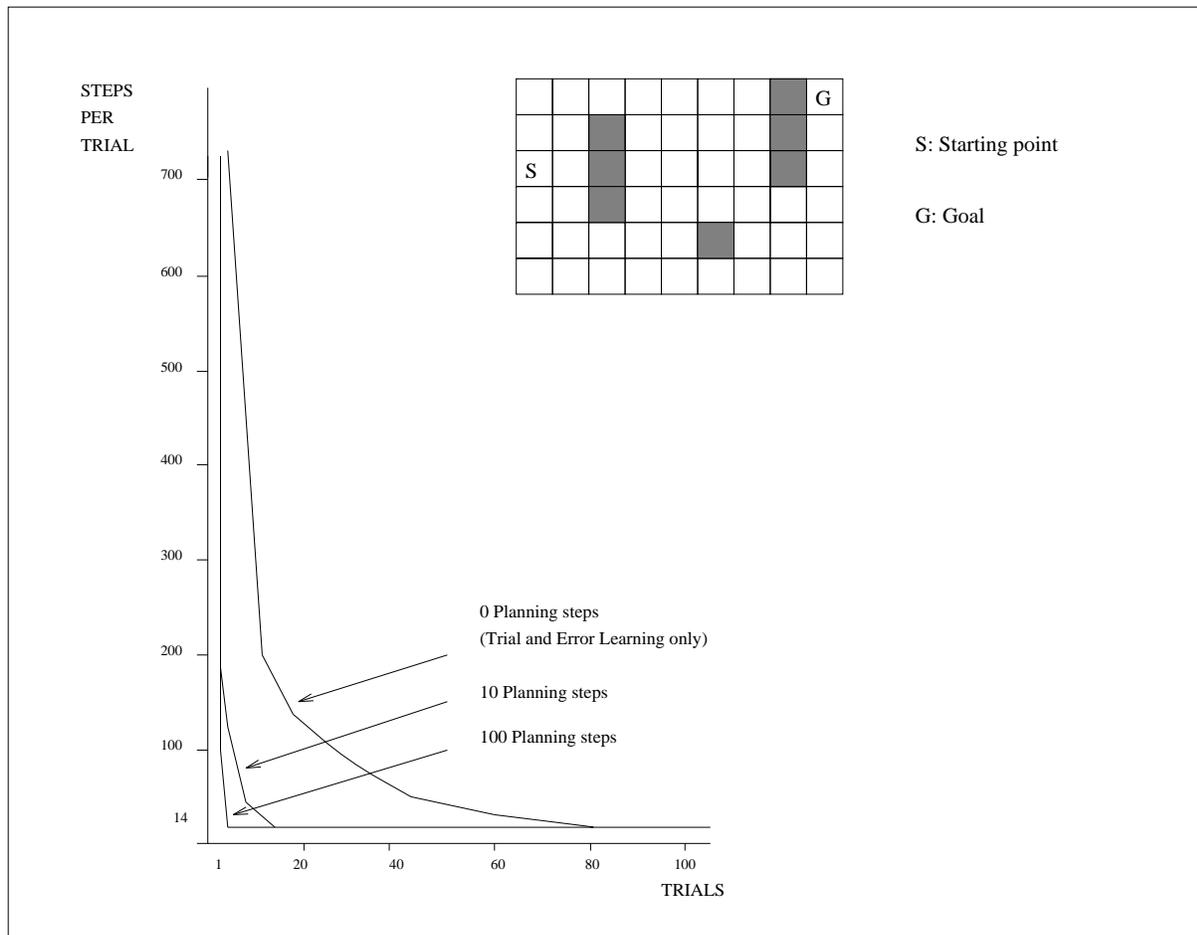


Figure 7. An example of the learning curves obtained in a specific navigation task with Dyna. The more hypothetical experiments (k planning steps) using the world model, the faster the optimal path is found. After [SUTT91]

Simulations effected on various DYNA architectures indicate that the corresponding animats are able to learn the shortest path leading from the starting state S to the goal state G , in the obstacle-encumbered environment on Figure 7. Moreover, such trial-and-error learning is expedited when animats avail themselves of the planning possibilities afforded by their internal world model and interleave one or more hypothetical experiments with actual experiments. What is more, some of these architectures exhibit generalization capabilities in changing environments.

4 Evolved Behaviors

Cliff et al. [CLIF93] have simulated an evolving process during which the neural architecture and the sensors' properties of each individual in a population of animats are improved over successive generations. During this process, *genotypes* of offspring are inherited from those of their parents and altered under the influence of mutation and crossing-over operators. At each

generation, the adaptive value of each *phenotype* -called its *fitness*- is evaluated by a test of the aptitude of each animat to generate the behavior sought by the experimenter, thereby allowing the genotypes that perform best to reproduce from one generation to the next, as well as eliminating those genotypes that perform most poorly.

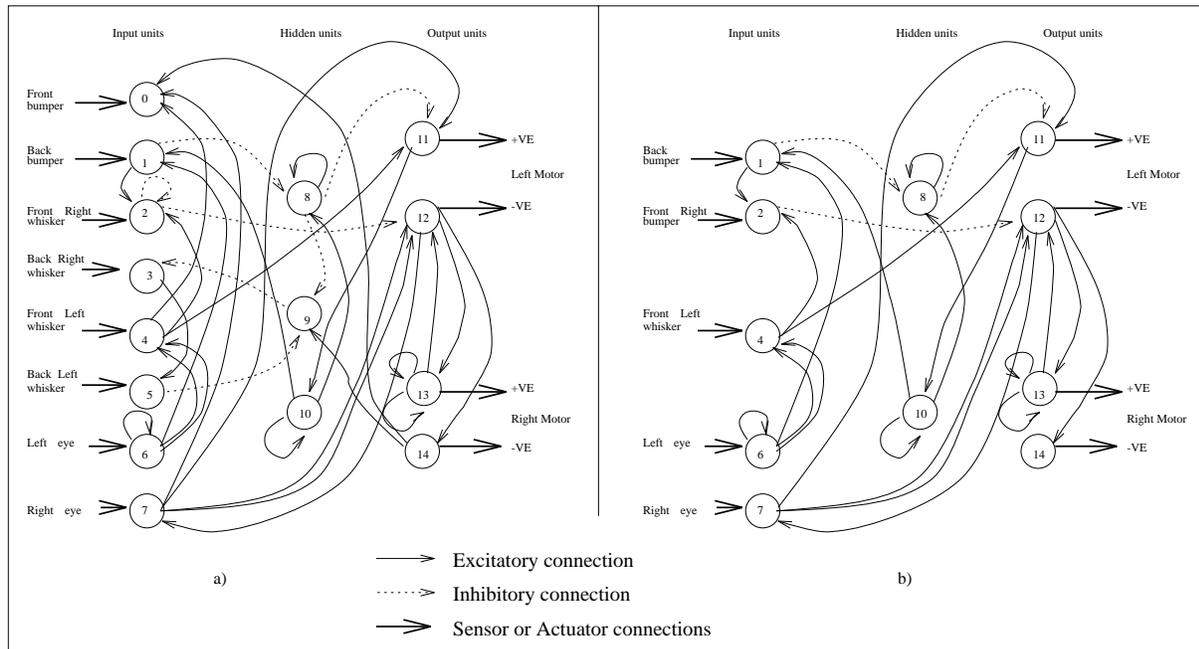


Figure 8. a) An evolved neural controller. b) The same network as on Figure a), with redundancies eliminated. After [CLIF93].

Instead of generating explicit control programs, Cliff et al. [CLIF93] make evolve the architecture of neural networks that directly link the sensors and actuators of an animat. The animat is equipped with various sensors: two forward and two backward whiskers, a front and a back bumper, and two photoreceptors. It is also equipped with actuators: two wheels and a trailing rear castor. The wheels have independent drives allowing turning on the spot and fairly unrestricted motion across a flat floor.

The architecture of the nervous system is general: the neurons are noisy linear threshold units, variable in number. If eight of them are input neurons -one neuron per sensor- and four of them output neurons -two neurons per motor- the number of interneurons is variable and is determined genetically. Likewise, if certain connections may interconnect two neurons, other connections result in temporarily preventing any transfer of information along specific direct connections. The number and nature of these various connections are genetically determined: the genotype of each animat consists of two chromosomes, one coding for the neural architecture, and the other coding for the properties of the visual sensors, i.e. the angle of acceptance and the eccentricity of the two photoreceptors.

The results show that it is possible to cause an animat's nervous system to evolve in such a way as to enable it to use its visual perception capabilities to avoid collisions with the wall of

an empty cylindrical room prior to making physical contact with the wall via one of its tactile sensors. Thus an animat can evolve in such a way as to be able to predict, from visual data alone, that a collision is likely in a near future, and to initiate appropriate evasive action. Examination of the evolved networks that generate such behavior reveals a complex connectivity with numerous redundancies. In earlier generations, the tactile sensors are widely used. Later, vision becomes progressively more dominant, and the tactile sensor input units are essentially used as internal neurons which process visual information (Figure 8).

5 Behavioral Development

If learning and evolution have already often been used for the automatic design of control architectures of animats, such is not the case with development- a point stressed by Meyer and Guillot [MEYE94]. However, a few applications combining development, evolution and learning have recently been published [KODJ94].

The work of Nolfi and Parisi [NOLF91], for instance, is concerned with the evolution of animats that consume food randomly distributed within a simple 2-D environment. Each animat is equipped with a sensory system that allows it to perceive the direction and the distance of the nearest food element and with a motor system that provides the possibility of turning any angle between 90 degrees left and 90 degrees right, and to move forward 0 to 5 steps. The nervous system of each animat is a bidimensional network with up to 40 neurons, whose development is coded in the animat's genotype. This genotype is a fixed-length string of 40 blocks, each block being made up of eight genes that describe the developmental fate of a given neuron. The first five blocks in the string correspond to sensory neurons, the last five blocks to motor neurons and the 30 intermediate blocks to internal neurons, which can be arranged in a maximum of 7 layers.

Within a given block, the first gene is the *temporal expression gene*, and specifies when, during development, the corresponding neuron will be expressed. Neurons scheduled to appear after the animat's death are non-expressed neurons. Two *physical-position genes* represent respectively the x and y spatial coordinates of the corresponding neuron. The *branching-angle gene* and the *segment-length gene* respectively control the angle of each branching of the neuron's axon and the length of each branching segment. The *synaptic-weight gene* determines the synaptic weight of each connection established by the corresponding neuron. In other words, in this model, all connections originating in a given neuron have the same weight. The *bias gene* represents the activation bias of the corresponding neuron. Lastly, the *neuron-type gene* specifies, in the case of a sensory neuron, whether this neuron reacts to the angle or the distance of food and, in the case of a motor neuron, whether this neuron determines the angle of turn or the length of a forward step.

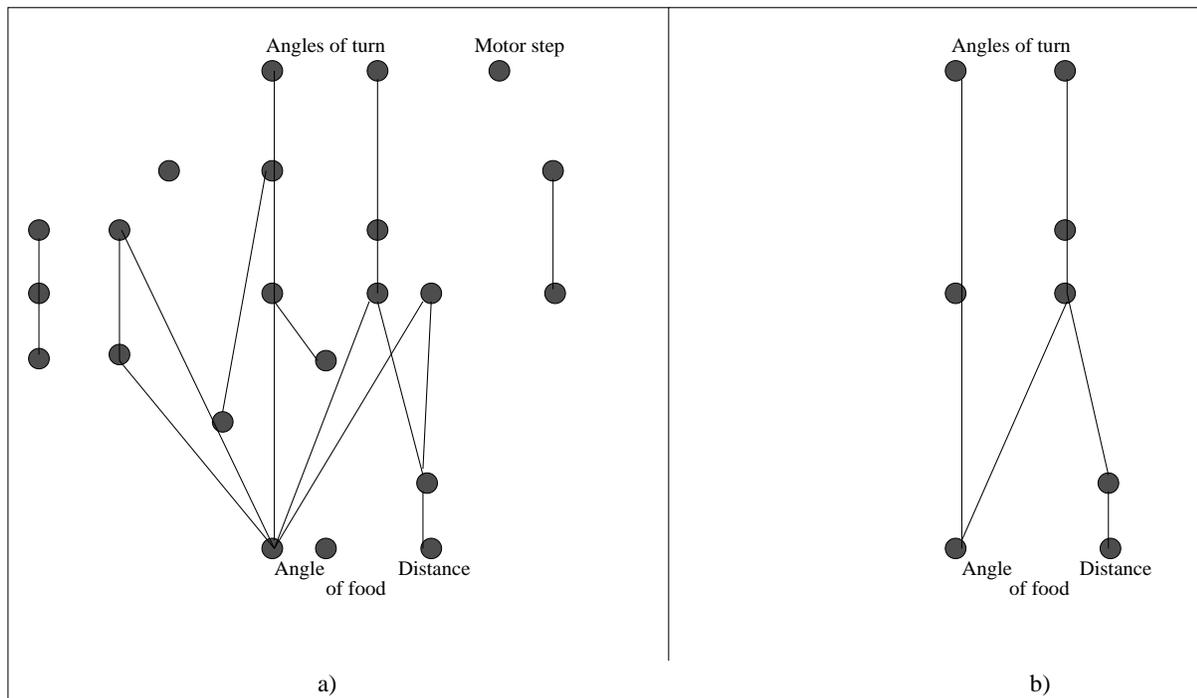


Figure 9. Growth process resulting from a randomly generated genotype. The lowest layer corresponds to sensory neurons, central layers to internal neurons and the upper layer to motor neurons. a) Connections established during the developmental process. b) Functional network mapping sensory input into motor output obtained after elimination of redundancies or unused neurons. After [NOLF91]

While, in approaches like that of Cliff et al. evocated above, there is a direct mapping between genotype and phenotype, in the work of Nolfi and Parisi, the nervous system of each animat changes during the animat's lifetime, according to the developmental instructions coded in the genotype. Thus, some neurons are created at birth, others appear later, and connections are established between two neurons when the growing axonal branch of a particular neuron reaches the soma of another neuron (Figure 9).

The results obtained by Nolfi and Parisi suggest that pairing an evolutionary with a developmental process is an efficient means for providing adaptive behavior in animats. The results also suggest that the architectures evolved tend to be structured into functional sub-networks.

6 Discussion and Conclusion

The various examples described above demonstrate that complex functionalities can arise out of simple models. In particular, it is now accepted that a few hundred neurons can be sufficient to control the motivational system and the behavioral sequences of a simple animat. These examples also demonstrate the fact that it is possible to analyze situations that are inaccessible to traditional observational and experimental methods, particularly when these implement an evolutionary process. They also suggest that the animat approach should help in assessing the

adaptive value of learning, evolution and development processes and in specifying how these processes interact with and complement each other. Finally, because research on animats seems to be an effective tool for studying how the highest intellectual abilities of man might derive from the simplest adaptive behaviors of animals, there is hope that this research will ultimately contribute to our understanding of the adaptive value and working principles of human cognition [MEYE95a, MEYE95b].

Nevertheless, the animat approach does present some limitations. For instance, the proof of principles that it yields are rarely used to best avail. Virtually never has research been conducted to ascertain whether a given simple solution to a problem is actually the simplest possible. Nor has it been demonstrated that any given adaptive capacity can be expressly ascribed to a specific global architecture rather than to a particular operational detail. Only a systematic comparison of several different versions of the same problem with as varied as possible a range of solutions can allow the respective advantages and the degree of originality of these solutions to be effectively evaluated.

More generally, it can be seen that, as of today, the animat approach is essentially empirical in nature and that it would gain from a broadening of its theoretical perspective. Fortunately, several research efforts have recently been initiated in such a direction [MEYE94]. For instance, the work of Wilson [WILS91] or Horswill [HORS92] on the characterization of environments and the adaptational problems to which they give rise, and that of Agre [AGRE91] or Chapman [CHAP92] on the theory of interactions between organisms and environments, constitute a valuable groundwork.

References

- [AGRE91] P.E. Agre. *The dynamic structure of everyday life*. Cambridge University Press. 1991.
- [BART81] A.G. Barto and R.S. Sutton. Landmark learning: an illustration of associative search. *Biological Cybernetics*. 42:1-8. 1981.
- [BEER90] R.D. Beer. *Intelligence as adaptive behavior: an experiment in computational neuroethology*. Academic Press. 1990.
- [CHAP92] D. Chapman. *Vision, Instruction and Action*. The MIT Press. 1992.
- [CLIF93] D. Cliff, I. Harvey and P. Husbands. Explorations in evolutionary robotics. *Adaptive Behavior*. 2,1:71-108. 1993.
- [CLIF94] D. Cliff, P. Husbands, J.A. Meyer, and S.W. Wilson (Eds). *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. MIT Press/Bradford Books. 1994.
- [GALL90] C.R. Gallistel. *The organization of learning*. The MIT Press. 1990.
- [GUIL94] A. Guillot and J.A. Meyer. Computer simulations of adaptive behavior in animats. *Proceedings Computer Animation'94*. IEEE Computer Society. 1994.

- [HORS92] I. Horswill. Characterizing adaptation by constraint. In F. Varela and P. Bourguine (Eds). *Toward a practice of autonomous systems. Proceedings of the First European Conference on Artificial Life*. The MIT Press. 1992.
- [KLOP80] A.H. Klopff. *The hedonistic neuron: a theory of memory, learning and intelligence*. Hemisphere. 1990.
- [KODJ94] J. Kodjabachian and J.A. Meyer. Development, learning and evolution in animats. *Proceedings of the PerAc'94 Conference*. IEEE Computer Society. 1994.
- [MEYE91a] J.A. Meyer and S.W. Wilson (Eds). *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*. MIT Press/Bradford Books. 1991.
- [MEYE91b] J.A. Meyer and A. Guillot. Simulation of adaptive behavior in animats: Review and prospect. In [MEYE91a].
- [MEYE93] J.A. Meyer, H.L. Roitblat, and S.W. Wilson (Eds). *From Animals to Animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior*. MIT Press/Bradford Books. 1993.
- [MEYE94] J.A. Meyer and A. Guillot. From SAB90 to SAB94 : Four Years of Animat Research. In [?].
- [MEYE95a] J.A. Meyer. The animat approach to cognitive science. In H. Roitblat and J.A. Meyer (Eds). *Comparative Approaches to Cognitive Science*. The MIT Press. In press. 1995.
- [MEYE95b] J.A. Meyer. Artificial life and the animat approach to artificial intelligence. In M. Boden (Eds). *Artificial Intelligence*. Academic Press. In press. 1995.
- [NOLF91] S. Nolfi and D. Parisi. Growing neural networks. *T.R. PCIA-91-15. Institute of Psychology. Rome*. 1991.
- [SUTT91] R.S. Sutton. Reinforcement learning architectures from animats. In [MEYE91a].
- [WILS91] S.W. Wilson. The animat path to AI. In [MEYE91a].

Authors Biographies

Agnes Guillot is Associate Professor of Psychophysiology at the University of Paris X and is pursuing her research within the Groupe de BioInformatique of the Ecole Normale Supérieure in Paris. She is interested in building causal and functional computer models of action selection in animals.

Jean-Arcady Meyer is currently Research Director at the CNRS and works in the Groupe de BioInformatique of the Ecole Normale Supérieure in Paris. His major research interests are the interactions of learning, development and evolution in adaptive systems, both natural and artificial. He initiated the SAB (Simulation of Adaptive Behavior: from animals to animats) series of International Conferences and is the Editor-in-Chief of the journal Adaptive Behavior.