

# From Natural to Artificial Life: Biomimetic Mechanisms in Animat Designs

Jean-Arcady Meyer

AnimatLab  
Ecole Normale Supérieure  
46 rue d'Ulm  
75230 Paris cedex 05, France  
E-mail: meyer@wotan.ens.fr

## Abstract

*This paper describes several models that incorporate some biomimetic mechanisms into the control architecture of an animat that has to survive in a changing environment. The adaptive capacities of these mechanisms are compared to engineering solutions, in application domains that involve navigation, action selection, and evolutionary design. It is argued that, although impressive adaptive capacities are already afforded by such biomimetic mechanisms, these capacities are nevertheless limited by the shortcomings of current biological knowledge. It is also argued that these shortcomings could be remedied were biologists willing to take advantage of the bottom-up and synthetic point of view with which artificial life might complement their traditional top-down and analytic approach.*

## Keywords

ANIMATS – NAVIGATION – ACTION SELECTION – EVOLUTION – DEVELOPMENT

## 1 Introduction

In its endeavor to unravel the mechanisms and processes that sustain natural life, traditional Biology usually seeks to decompose a system into its constituent sub-systems and then to study these sub-systems in isolation from one another, according to a *top-down*, *analytical* and *reductionist* approach. In contrast, scientists involved in artificial life research attempt to *synthesize* behaviors characteristic of natural living systems with man-made artifacts like computers or robots. Ideally, their approach is bottom-up and starts with a collection of entities or modules exhibiting behaviors that are simple and well understood, organizing them into more complex systems in which internal interactions generate *emergent* life-like properties (Langton, 1989; Taylor, 1992).

Research in artificial life (Langton, 1989; Langton et al., 1992; Varela and Bourgine, 1992; Langton et al., 1994; Brooks and Maes, 1995; Moran et al., 1995) has given rise to a number of applications, notably regarding computer viruses, autocatalytic nets, cellular automata, artificial nucleotides, immune systems, market simulations, etc. In particular, it has been concerned with the conception and construction of *animats* - that is, of simulated animals or autonomous robots - whose structures and functionalities are as much inspired from current biological knowledge as possible, in the hope that these animats will exhibit at least some of the adaptive capacities of real animals and prove capable of surviving in unpredictable and possibly threatening environments (Meyer and Wilson, 1991; Meyer et al., 1993; Cliff et al., 1994; Maes et al., 1996).


The main problem facing any animat designer is that of figuring out what sensors and what actuators should be connected through what control architecture in order to solve what survival problem in what environment. Although the objectives, methods and results of animat research have been extensively reviewed elsewhere, the corresponding papers have mostly emphasized their similarities and differences with respect to those of traditional Artificial Intelligence (Meyer and Guillot, 1991, Wilson, 1991; Maes, 1993; Meyer and Guillot, 1994; Steels and Brooks, 1995; Meyer, 1995, 1996 ). Within the framework of such comparisons, the biological roots of the animat approach haven't been granted special attention. This paper will be centered on these roots and will argue that adaptive behaviors generated by biomimetic mechanisms are promising solutions to various survival problems, and that such behaviors are likely to provide efficient substitutes for the purely engineering solutions to these problems that are widely used at present. Current research efforts representing three domains of application will illustrate this point, and will deal successively with problems of *navigation*, of *action selection*, and of *evolutionary design*.

## 2 Navigation

The ability to navigate - i.e., to reach any goal place from any starting place, while avoiding passing through other places - is probably the most basic requisite for an animat's survival. Indeed, without such an ability, the animat wouldn't be able to reach energy sources, to avoid bumping into damaging obstacles, or to escape from dangerous hazards. Engineering solutions to the navigation problem usually resort to internal geometrical representations of the environment (Schwartz and Yeap, 1987; Canny, 1988; Latombe, 1991; Hwang and Ahuja, 1992) that are confronted with various implementation difficulties, due to memory or time requirements, as well as sensory and motor errors (Nehmzow, 1995). Therefore, as many animals are very good at navigating in apparently simpler and more robust ways, it is not surprising that several research efforts have already been aimed at incorporating the corresponding current biological knowledge into animat navigation models (for a review, see Trullier et al., 1996).

The work of Franceschini (Franceschini et al., 1989; Franceschini, 1992), for example, calls upon knowledge gained about the compound eye of the fly and upon the development of an on-board electro-optical system for the guidance of a mobile robot. This robot moves about and, from its relative motion between itself and the environment, evaluates the distance to any potential obstacle. The control system of the robot includes a main visual system with compound optics and a panoramic field that calls upon a circular network of motion detectors. Each such detector measures the relative angular velocity of any contrast point that passes across its receptive field as a result of the robot's own motion and makes use of the motion parallax to evaluate the radial distance to that point (Figure 1). The information relayed by the motion detectors is then input to an analog parallel neural network designed to implement an obstacle-avoidance algorithm. This network transforms the polar map of the environment delivered by the detectors into two analog signals that modulate the speed of the drive and steering motors of the robot. The control architecture is also equipped with an accessory visual system that constantly surveys the horizon and informs the robot about the bearings of a target lamp.

Likewise, the navigation system described by Cartwright and Collett (1983) is inspired from the way bees use nearby landmarks to return to an already known food source. This ability relies upon the use of the remembered image of what was present on their retina when they were at the food source: bees seem to adjust their flight path by continuously comparing their retinal image of the landmark configuration with their memorized snapshot and to lessen the



**Figure 1.** Principle of motion parallax. The distance  $D$  to a point  $P$  located at azimuth  $\Phi$  can be calculated by an animat translating at speed  $V_0$ , if it is equipped with sensors measuring the angular speed  $\Omega$  of  $P$ .


discrepancy between the two (Figure 2).

Other models involve spatial representations that seem to be encoded in the hippocampus and para-hippocampal formations in the brain of some animals (O'Keefe and Nadel, 1978) and to involve so-called *place cells* and *head direction cells*. In relation to the spatial layout of surrounding landmarks, the activity of the former depends on the animal's current location, whatever the direction faced (for a review, see Muller et al., 1991), whereas the activity of the latter depends upon the absolute orientation of the head, whatever the animal's location (Taube et al., 1990a, 1990b). Various more or less realistic equivalents of place cells and head direction cells have been used in several navigation systems. For example, when the animat of Zipser (1986) explores its environment, it recruits at regular time intervals a so-called *view cell*<sup>1</sup> whose activity is tuned to the distance and bearings of given landmarks seen from the current position. If the animat moves slightly from this position, the activity of the corresponding view cell diminishes according to a Gaussian law so that, when the distances and bearings of the landmarks significantly differ from those to which the view cell has been tuned, the cell will no longer be activated. Likewise, every time a view cell is recruited, a so-called goal cell is also recruited, whose output codes the direction of the goal from the current position (Figure 3).


After exploration, when the animat moves in its environment, the degree of activation of each view cell it has previously recruited will depend upon the position and orientation of the animat relative to the landmarks. Such an activity pattern uniquely characterizes the place the animat is currently situated in and, therefore, constitutes what, since Tolman (1948), has come to be called a *cognitive map*. Insofar as the activity level of each view cell is used to weight the direction information provided by its associated goal cell, the direction in which the animat should move in order to reach its goal from the current position is given by the weighted average of the directions coded by each goal cell (Figure 4).

Such an implementation endows an animat with minimal adaptive capacities. In particular, its ability to self-estimate its position in the environment relies upon information that

<sup>1</sup>It is not called a place cell because its activity also depends upon the current orientation of the animat.



**Figure 2.** Cartwright and Collett's snapshot model. The model bee is represented by two concentric circles, the outer circle represents its current retinal image and the inner one represents its memorized snapshot. (a) At food location, the images of the landmarks (filled circles) project onto the retina and the snapshot, as shown. To return to the feeding site (b), the model bee tries to match its retinal image (dark areas on the outer circle) and the remembered snapshot (dark areas on the inner circle). Each dark area on the snapshot is paired with the closest dark area on the retinal image (dashed lines). Each pair yields a radial vector (forward/backward motion) for the size adjustment and a tangential vector (left/right rotation) for the orientation adjustment (unfilled arrows). The resulting movement, which globally reduces the discrepancy between the snapshots, is computed (c) from the summed vector (filled arrow).



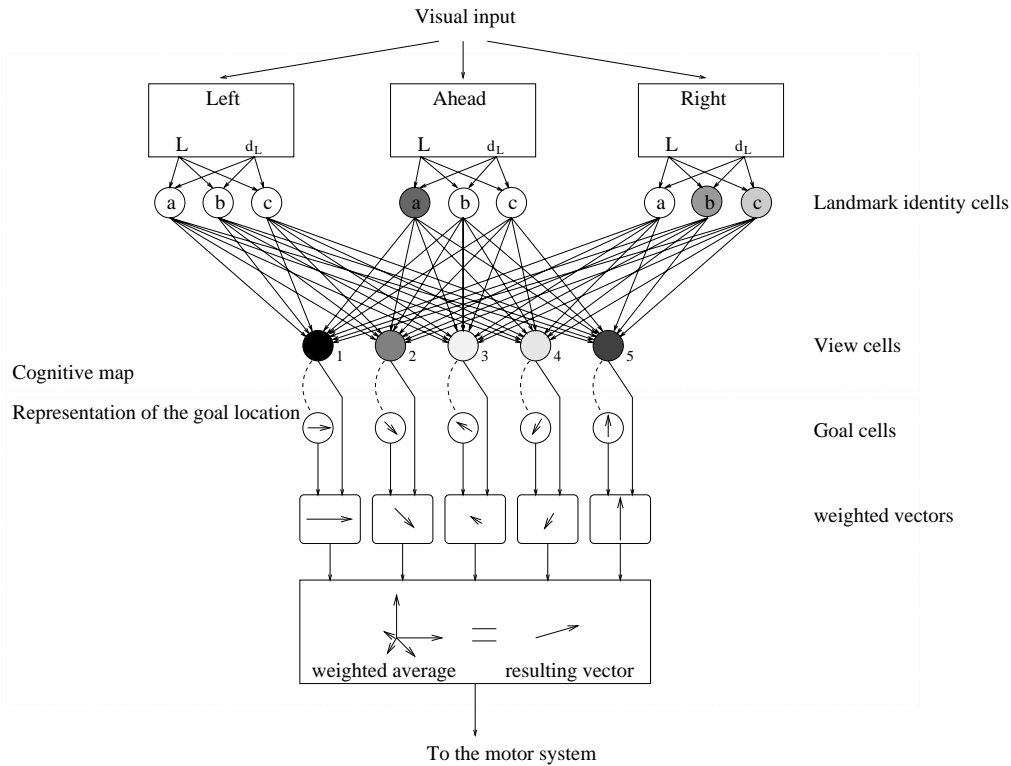
**Figure 3.** Zipser's view-field model. Each view-cell in the model is associated with a view-field, defined by the distances from the center of the field to three landmarks and by an orientation. This orientation, which serves as a local reference direction, corresponds to the orientation the animat had at its first visit to the place, and is defined by where the three landmarks are with respect to the animat's head-direction (left, ahead, or right). The direction to the goal from the view-fields is coded as vectors, oriented with respect to each local reference direction. The direction to the goal is interpolated from the activities and the vectors associated with all the view-cells (see text and Figure 4).

is distributed across numerous place cells in its cognitive map, and the accidental destruction of a few such place cells would cause a graceful degradation of the animat's navigation capacities.


The spatial representations built by other navigation systems also encode the topological links that characterize the animat's environment. During exploration, when the animat moves from one given place to another, a link is drawn between the place cells that are associated with these places, and this link is indexed by the associated move. Thus, subsequently, because the animat knows which moves are likely to lead to which place from which other place, it can plan a trajectory to the goal. This, for example, is the case with Mataric's robot (Mataric, 1990) that explores the environment by following walls. This robot uses the information provided by its sensors and its effectors to characterize different places in its environment and builds a topological cognitive map that records the possible moves from place to place. Thus, while the robot explores the environment shown on Figure 5, it records in its map that place C4 is passed through while it moves in the direction indicated by its compass and while its sonars detect similar nearby obstacles on its right and on its left - thus suggesting that place C4 is actually a corridor. Likewise, it records that a right turn leads from C4 to another place, RW6, which is passed through while the robot moves in a given direction and while it detects a nearby obstacle on its right - thus suggesting that place RW6 is adjacent to a right wall. In other words, places like C4 and RW6 on the robot's cognitive map play the role of place cells that become activated when the robot passes through the corresponding places in the environment, and such a map allows the robot to position itself correctly.

Because the map also provides information about the physical length of each place, the shortest path leading to any given goal from the current place can be generated by initiating a spreading activation process throughout the map, in all directions from the goal. Insofar as the speed of this process depends upon the length of the places through which it travels, the direction from which goal activation first arrives in each place indicates the direction in which to move in order to reach the goal.

Such a navigation system provides an animat with the highly adaptive capacity of planning its trajectory to the goal. According to the *experience in the head* metaphor ( Craik, 1943), trajectory planning makes it possible to run internal simulations that, being decoupled from overt behavior, avoid the hazards of dangerous encounters.



**Figure 4.** Zipser's navigation system. The upper module of the neural network (bordered in gray) is the oriented place representation by view cells. Each landmark (a, b or c), indicated by its identity L, is seen left, ahead or right relative to the head-direction and at a certain distance  $d_L$ . The landmark identity cells a, b, and c fire proportionally to the distances to the associated landmarks ( $d_L$ ) when the associated landmarks are in their gross angular field (L is 0 or 1). For instance, in the position of the animat in Figure 3, landmark identity cells a and b for the "Ahead" visual field and landmark identity cell c for the "Right" visual field are active (shading) and all the other landmark identity cells remain silent. The combination of the landmark identity cells activates the view cells. Here, view-cells 1 and 5 fire the most and view-cell 3 fires the least (shading). The lower module represents the subsequent population coding of the direction to the goal. Each view-cell is associated with the direction to the goal, and the corresponding vector is represented by a goal cell. The vectors representing these directions are weighted by the activities of their associated view cells and the weighted average yields the estimate of the direction to the goal from the animat's current position.

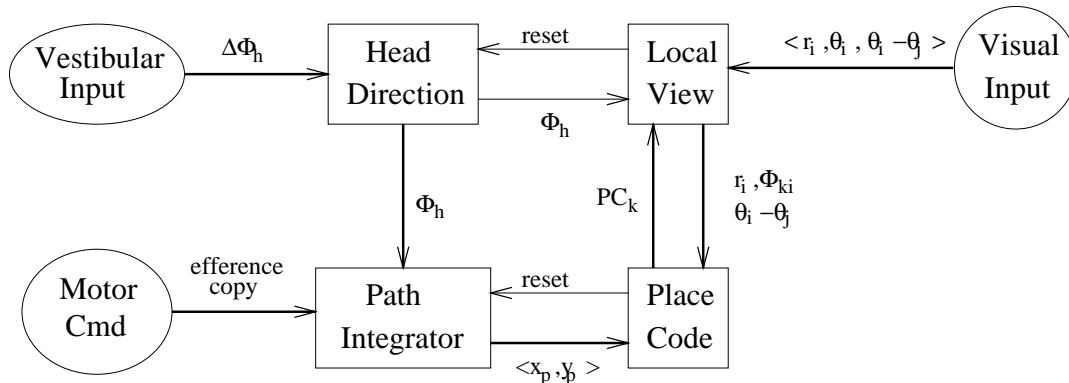


**Figure 5.** Example of an environment and of the associated graph in Mataric’s model. A place is labeled according to the characteristics of the range sensor readings (RW: right wall, LW: left wall, C: corridor) and to the compass reading (from 0 to 15) when the robot, driven by its lower-level guidance strategy (a wall-following procedure), travels through it. In the corresponding cognitive map, each place is represented by a node, and any two nodes can be linked dynamically to code the adjacency relationship.

Finally, another variety of navigation model calls upon the representation of metric spatial relations between two or more places on a plane and upon vector manipulations. Usually these models involve a system of coordinates by reference to which the positions of places are recorded. This, for example, is the case with the model of Wan et al. (1994) that integrates multi-modal information involving kinesthetic signals provided by *efferent copies*<sup>2</sup>, proprioceptive signals supplied by the integration of the angular velocity of the animat’s head, and exteroceptive signals from its visual sensory system (Figure 6).

In this model, the correspondence between the animat’s cognitive map and its position in the environment is coded by the activity of so-called *place code* units. This activity is determined by the product of seven Gaussian functions, respectively tuned, at exploration time, to the distances and allocentric bearings of two selected landmarks, to the egocentric bearing difference between two other selected landmarks, and to the estimated Cartesian coordinates of the animat’s position relative to an identified reference point. Should one or more items taken into account in such a calculation be missing, the corresponding term or terms would simply be excluded from the product. This model thus accounts for the experimental persistence of place cell activity in darkness (Quirk et al., 1992), because information provided by the path integrator mechanism is capable of substituting for the missing visual information. In this model, the local view module (McNaughton, 1991) transforms the egocentrically sensed visual information into an allocentric representation that is stored in the place code module, by simply summing the head direction and the egocentric bearing. Conversely, by subtracting the egocentric bearing provided by the visual input to the allocentric bearing recorded in the place code module, the local view module makes it possible to update the estimate of the head direction (Figure 7).

<sup>2</sup>Motor subsystems that drive muscles can map an *efference copy* or *corollary discharge* back onto the perceptual system, which can thus distinguish self-induced movements from world-induced movements (McCloskey, 1981).



**Figure 6.** Functional block diagram of Wan et al.’s model. It shows how multimodal information might be combined and mutually influenced. Place is represented as the fuzzy conjunction of multimodal information, coupling path integration ( $\langle x_p, y_p \rangle$  Cartesian coordinates with respect to an a priori global reference frame), visual information (distance  $r_i$  and relative bearing  $\theta_i$  of each visible landmark  $i$ , transformed into absolute bearing  $\Phi_{ki}$  that depends on the current place  $k$ ), and head-direction ( $\Phi_h$ ). The head-direction representation is updated through the integration of angular velocity  $\Delta\Phi_h$  and reset by place recognition when there is some drift ( $PC_k$ ).

Likewise, the place code module makes it possible to reset the Cartesian coordinates of the animat, which can then be used by a vector manipulation module to plan a trajectory to any given goal.

Besides the particular adaptive capacities that the integration of multimodal signals and the reset of the error-accumulating information given by the kinesthetic and proprioceptive integrators provide to the above specific implementation, such a navigation system affords an animat other important potentials for survival, since additional vector manipulation abilities make metric *detours* and *shortcuts* possible.

### 3 Action selection

The chances of survival of any animal are obviously heavily dependent upon its ability to solve the action selection problem (McFarland, 1981; Maes, 1994), i.e., upon its ability to choose, at any time, the proper action to perform in order to optimize the achievement of its goals. Since the days of the robot Shakey (Nilson, 1983), traditional engineering solutions to this problem have relied upon the assumption that the environment can be accurately and exhaustively modeled by the animat and have been reduced to the problem of planning a sequence of actions that will, if performed, achieve a single and fixed goal imposed by the experimenter. Other engineering solutions (Kaelbling and Rosenschein, 1990) have been designed to deal with sets of fixed goals by automatizing the process of designing the arbitration circuitry that selects the proper action. These solutions require the experimenter to specify the animat’s repertoires of goals and actions, and how these goals can be reduced to other goals or actions. Then, a compiler analyses this specification and generates a circuit that implements the desired goal-seeking behavior.

Numerous research efforts have been directed at dealing with the more challenging action selection problems that any animal experiences in nature and at unraveling the basic mecha-



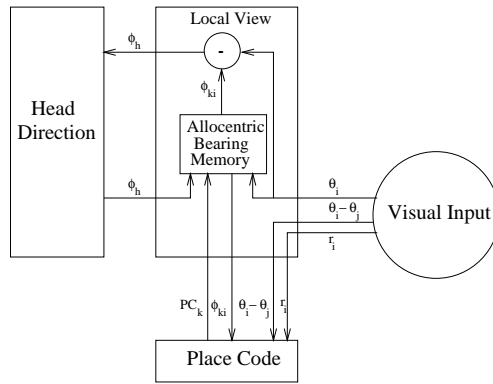
nisms that allow him to survive (for a review, see Guillot and Meyer, 1994). The corresponding models usually involve several endogenously-generated time-varying goals and implement a *motivational system* (McFarland, 1974; Toates, 1986; Colgan, 1989). The Booker model (Booker, 1988), for instance, uses a *classifier system* - i.e., a rule-based structure designed to reproduce various cognitive processes in animals and men (Holland et al., 1986) - in order to allow an animat to learn to produce goal-seeking sequences through the use of rules that manipulate objects, goals, and object/goal associations, rather than rules that code simple stimulus/response associations. Such a structure allows the animat to move about on a grid containing food objects and noxious objects, these objects being located in the middle of stimulus auras whose intensity diminishes with the distance from the object in question. Contact with food gives rise to a reward, while contact with a noxious object entails a punishment. To survive, the animat must learn which actions are appropriate in which situations; it, therefore, must build a model of its environment. To this end, its control architecture mimics the hierarchical structure described by Tinbergen (1951) to model instinctive behavior in animals. Basically, it involves three instinctive centers concerned with locomotion, food-seeking, and pain aversion, as well as innate releasing mechanisms that prompt the animat to pursue two objectives: to eat in the presence of food and to flee in the presence of a noxious object (Figure 8).

The animat's default activity consists in exploring its environment. This activity is interrupted when the strength of the sensory or motivational signals associated with a lower instinctive center allow it to be activated and to trigger an approach or avoidance behavior. These behaviors can be expected to prompt the animat to feed or to flee, according to what kind of object it has come in contact with, and to shut off the activity of the corresponding center, thus allowing the animat to resume its exploration.

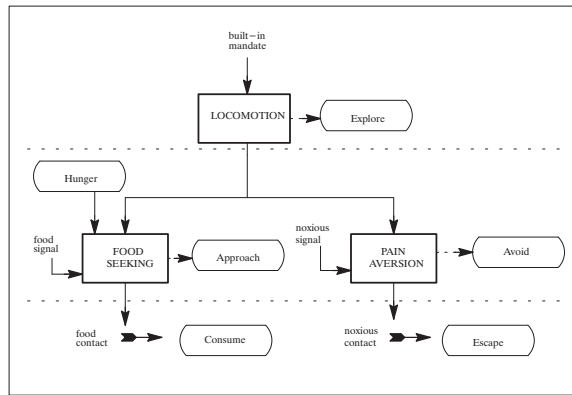
In this approach, the animat's adaptive capacities derive from its ability to learn representations. These representations are used to classify the objects encountered in the environment into categories that have affective significance and accordingly to trigger the appropriate behavior. In other words, Booker's animat learns to respond to the regularities hidden behind the equivocal nature of its sensory cues and to take advantage of the useful and structured information afforded by the environment (Gibson, 1986).

According to Tyrrel (1993), such a hierarchical control architecture gives rise to several problems when dealing with more complex environments that pose more challenging survival problems and that entail the management of more than three elementary behaviors. To demonstrate this point, he devised a simulated environment which posed 13 different survival problems to an animat, i.e. obtaining food, obtaining water, keeping clean, regulating body temperature, avoiding predators, being vigilant for predators, staying close to cover, avoiding the boundaries of the territory, avoiding dangerous places, avoiding irrelevant creatures, sleeping at night in a den, not getting lost and reproducing. To solve these problems the animat has a choice of 35 actions that it can undertake, such as drink, clean self, sleep, move north, move south-west and look around. By choosing appropriate actions in appropriate situations, the animat can exert some control over the values of its internal variables - like its levels of food and water - and over its environment - like how well it perceives its local environment or can be perceived by predators. The overall quality of a given action selection controller is measured by an equivalent of genetic fitness, i.e. by the number of times the animat manages to reproduce before it dies.

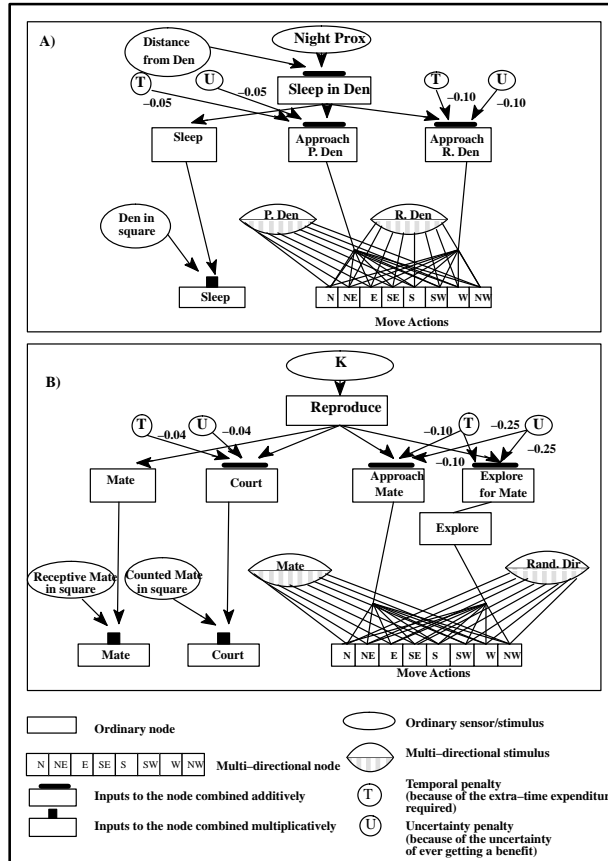
Comparative results obtained by Tyrrell on various control architectures suggest that the most efficient design is that of a set of overlapping *free flow hierarchies*, each hierarchy being devoted to each survival problem that the animat is confronted with. Such hierarchies are made up of nodes interconnected via weighted connections, which are similar to standard artificial



**Figure 7.** In Wan et al.’s model, the animat learns the relationship between egocentric and allocentric bearings of specific cues as viewed from a set of locations. Information about the current (estimated) head-direction (*input*  $\Phi_h$ ), the current (estimated) position ( $PC_k$ ) and the egocentric bearings of landmarks ( $\theta_i$ ) converge into the Allocentric Bearing Memory, which computes the head-direction independent (allocentric) bearings of these landmarks ( $\Phi_{ki}$  for landmark  $i$  at place  $k$ ). This module can also correct the head-direction representation by using the memorized allocentric bearings, the representation of the current place and visual information.



**Figure 8.** Control architecture of Booker’s animat. The rectangles correspond to instinct centers and arrows to innate releasing mechanisms. Each instinctive center is characterized by a releasing stimulus and an associated action, and can be submitted to two sources of control: the current motivational state and other instinctive centers higher in the hierarchy. Here, the food-seeking and pain-aversion centers are located on the same hierarchical level, under the control of the locomotion center. The food-seeking center is subjected to an additional motivational control exerted by the degree of hunger the animat experiences.




**Figure 9.** Examples of Tyrrel's free flow hierarchies. A) Control of Sleep in Den behavior. The Night Prox stimulus is 0.0 at daybreak, then increases as nightfall approaches and stays at its maximum value through the night. The Den in Square stimulus is 1.0 if the animat is in his den, 0.0 otherwise. P.Den and R.Den stimuli respectively code for the perceived and remembered directions of the den. T and U are temporal and uncertainty penalties which tend slightly to inhibit the Approach Den nodes. B) Control of Reproduce behavior. K is a constant stimulus that provides an unvarying motivation to reproduce.

neurons except that their rule for combination of weighted inputs is not necessarily a strict summation. Nodes express multiple preferences for each of a set of lower-level alternatives and the spread of activation between nodes propagates combining evidence from the upper node - which codes for the survival problem to be solved - down to the lowest nodes - which code for the actions which can be executed. Because actions are mutually exclusive, it is the action-level node receiving the most activation at every moment which is activated. Figure 9A, for example, describes the free flow hierarchy that causes the animat to move toward its den as nightfall approaches and then to sleep in the den for the rest of the night. Likewise, Figure 9B describes the free-flow hierarchy that makes the animat move toward a mate, court and reproduce.

The architecture of Blumberg's (1994) model is a compromise between the strict and loose hierarchies implemented respectively in the Booker and Tyrrell models. Although it implements a *winner-take-all* system such that only one behavior system is expressed at a given time in the movements of the animat, it nonetheless allows losing activities to express their preferences in the form of recommendations to the winning activity. The winner may use this information as it sees fit. For example, an anti-predator system can detect the presence of a distant predator and, while still staying inactivated, post recommendations against movements that would bring the animat closer to the predator. In addition, the Blumberg model includes a biologically plausible mechanism of activity-specific fatigue that has been suggested by Ludlow (1976). According to this mechanism, which implements the winner-take-all arbitration among activities at every level of the control hierarchy, the animat can avoid the mindless pursuit of an unattainable goal to the detriment of other goals. Finally, proper calibration of the effects of innate releasing mechanisms allow the animat to opportunistically switch to another action, even if the one it was currently indulging in has not yet fulfilled its endogenous need. For example, on its way to find food, the animat can take advantage of a nearby water source, although its endogenous need to feed is greater than its need to drink.

Although the preceding models exhibit impressive adaptive abilities, they all belong to the category of *motivated automata* in terms of the classification proposed by McFarland and Bösser (1993). According to these authors, the motivational state of an animal at any particular time ideally depends on its physiological or internal state, on the cue state arising from its perception of the external world, on the consequences of its current behavior, and on the expected consequences of its future behavior. If the decision to perform a given action doesn't take into account the expected consequences of the alternatives, the animal behaves like a motivated automaton, otherwise it behaves like a *motivationally autonomous agent*. To do this, the animal requires knowledge of the probable consequences of its actions: in other words, it must have some memory of the past consequences of similar activities, and it must be capable of planning.

Donnart and Meyer (1994, 1996a, 1996b) devised the control architecture of a motivationally autonomous animat. Basically, it implements a hierarchical classifier system that calls upon three pools of production rules - action rules, planning rules and mapping rules - that are learned as the animat experiences its environment. Action rules choose the next action to perform according to the current sensory information provided by the animat's proximate sensors and according to the current goal or sub-goal. Planning rules use the sensory information provided by the proximate sensors, augmented by the animat's current coordinates and orientation, to decompose the current goal into sub-goals. Finally, mapping rules record topological links between landmarks that are indexed by the animat's current control strategy. Such a control architecture enables an animat equipped with a rudimentary and noisy sensory-motor apparatus to explore its environment, to extricate itself from dead-ends with complicated shapes, to build a cognitive map of its environment, to accurately estimate its current position, and to plan

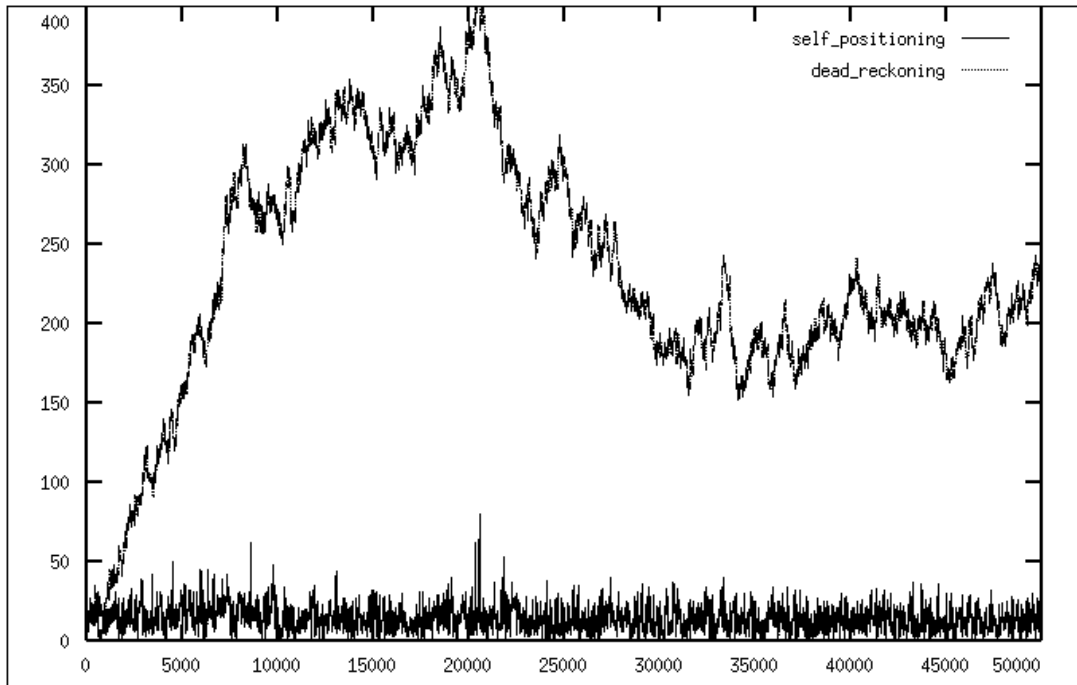


**Figure 10.** A square environment explored by the animat of Donnart and Meyer. The size of the animat is 15 pixels, the size of each side of the environment is 735 pixels. Parts of the cognitive map elaborated by the animat are shown on the picture, in the form of vectors coding topological links between pairs of landmarks. Most often links are short and concern landmarks encountered while the animat skirts around an obstacle. Random wanderings through free space eventually lead to the recording of links between more distant landmarks.

trajectories that avoid obstacles and lead to a given goal. Figure 10 shows parts of the cognitive map that the animat builds when left free to explore its environment. Such a map can be used by the animat to correctly position itself even when its faulty dead-reckoning mechanism generates an error that increases with the distance traveled, as demonstrated on Figure 11.

Finally, Figure 12 shows that a few iterations of the same navigation task allow the animat to adapt its action rules and to optimally skirt around an obstacle. The animat is also able to detect several salient landmarks along the corresponding trajectory. These landmarks can then be used to generate planning rules defining a plan that allows the animat to avoid the obstacle from a distance in future iterations of the navigation task. This plan can be quickly adapted to successive modifications of the environment because, every time an unexpected obstacle is encountered, the animat can resort to its action rules to skirt around it and can generate new planning rules that define an avoidance sub-plan.

It should be noted that, in the present implementation, the animat is seeking a single




**Figure 11.** Evolution of the distance between the animat's current position estimate and its actual position, during the construction of the cognitive map of Figure 10. The upper curve corresponds to results obtained without using the map, the lower one corresponds to results obtained when the map was used. Distances in pixels (ordinate) versus successive elementary moves (abscissa).

overall goal that is imposed by the experimenter, namely that of reaching a given place in its environment, and that it selects its action according to plans that aim at optimizing the expected length of the corresponding trajectory. However, these plans are modular and refined as needed. Each sub-plan is coded in a separate planning rule that takes the current sensory information into account and whose operational value is continually reevaluated. Therefore, several such rules are likely to be activated at any one moment, thus allowing the animat to use them as one resource among others to decide which action to perform and to adaptively switch between plans or sub-plans in a reactive manner (Suchman, 1987; Agree, 1988). Minor modifications of its control architecture will allow the animat to pursue several endogenously-generated goals in the future, depending upon plans that take into account evaluation criteria dedicated to the accomplishment of each specific goal.

## 4 Evolutionary design

Following the work of Holland (1975), a number of research efforts have addressed the simulation of processes that improve the behavior of animats in a population across successive generations. These efforts involve the implementation of selection procedures that eliminate individuals with ill-adapted behaviors and favor the reproduction of individuals displaying behaviors that are well adapted. Most often, they call upon a variety of evolutionary algorithm - like a *genetic algorithm* (Goldberg, 1989) or a *genetic programming* (Koza, 1992) algorithm - which allows



**Figure 12.** Adaptive plans generated by Donnat and Meyer's animat. Top left: Reactive trajectory in an environment containing a dead-end. This trajectory is obtained after 15 iterations, i.e., after 15 successive experiments during which the animat has reached the goal when starting from the initial state. Top right: At this stage, the animat has memorized two plans for avoiding the dead-end, of which only the best is shown. From now on, the animat will follow this plan and reach its goal by turning left to avoid the obstacle. Bottom left: If, at iteration 16, a new obstacle is added to the environment along the animat's optimum path, the animat is able to reactively skirt around it and reach the goal. In the process of updating its estimates of the lengths of the various trajectories leading to the goal, it takes six additional iterations to discriminate in favor of the second plan, which becomes more likely to govern the animat's path from the 22th iteration on. The animat will accordingly turn right to avoid the obstacle. Bottom right: Likewise, introducing a new obstacle into the environment at iteration 30 gives the advantage to the modified version of the first plan.

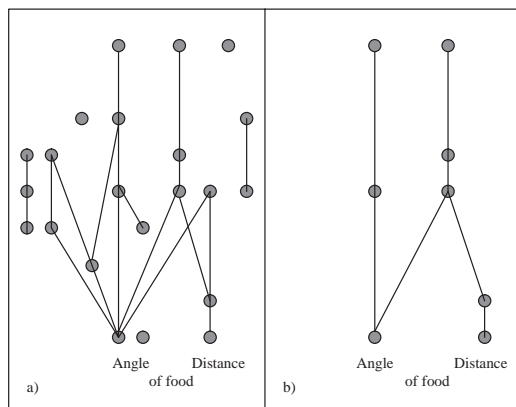
the parallel management of a population of *chromosomes*, such that each one codes the control architecture of a given animat. Each of these chromosomes can be assigned a *fitness* that assesses the ability of the corresponding control architecture to solve a given problem. The application of the evolutionary algorithm accordingly consists in causing the population to evolve from generation to generation, while rendering the probability of reproduction of each chromosome proportional to its fitness, and using genetic operators such as *mutation* or *crossing-over* to give rise to new control solutions in the population. Under these circumstances, such an evolutionary process generally causes control architectures of ever-increasing fitness to be generated until it converges to some local or global optimum.

Numerous engineering approaches to evolutionary design are described in the literature. They are usually applied to the design of neural networks (Harp et al., 1989; Miller et al., 1989; Montana and Davies, 1989; Cliff et al., 1993) or to the design of control programs that involve production rules (Grefenstette, 1992) or Lisp-like instructions (Koza et al., 1992). However, it can be argued that such approaches, because they entail direct mapping of the animat's *genotype* - i.e., the information coded in a chromosome - into its *phenotype* - i.e., the corresponding control architecture - probably lack many of the advantages that the natural process of development affords to living systems. Such a process, which is responsible for decoding the genotype into the phenotype, is not instantaneous, nor inflexible, and probably offers many opportunities for the environment to shape the adaptive solutions that the final control architecture will exhibit.

Among the models already published that explicitly incorporate a developmental process (for a review, see Kodjabachian and Meyer, 1995), few call upon biomimetic mechanisms like cell division, cell migration, or axon and dendrite growth. The work of Nolfi and Parisi (1991), for instance, is one such exception and 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 to the right or to the left, or to move forward. Likewise, the work of Miglino et al. (1996) involves an animat that must reach a target area in its environment using sensors, which react to the angle and distance of two landmarks, and actuators, which allow forward moves, right and left turns, and the possibility of remaining still. The nervous systems of these animats are bidimensional networks 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 seven 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* control respectively 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. Lastly, the *neuron-type gene* specifies to which input a sensory neuron will be active or which move a motor neuron will generate. Figure 13 provides an example of the kind of neural network that such a model develops. There are some indications that such a coding scheme presents interesting capacities from the standpoint of adaptive behavior: it favors the structuring of the





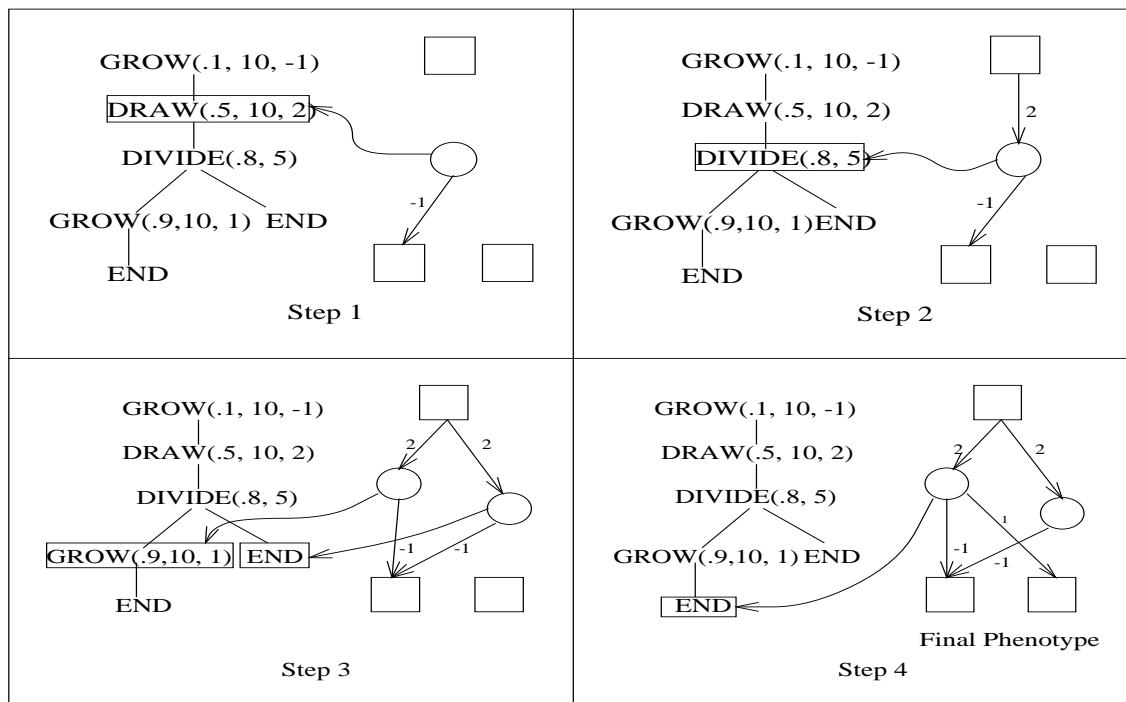
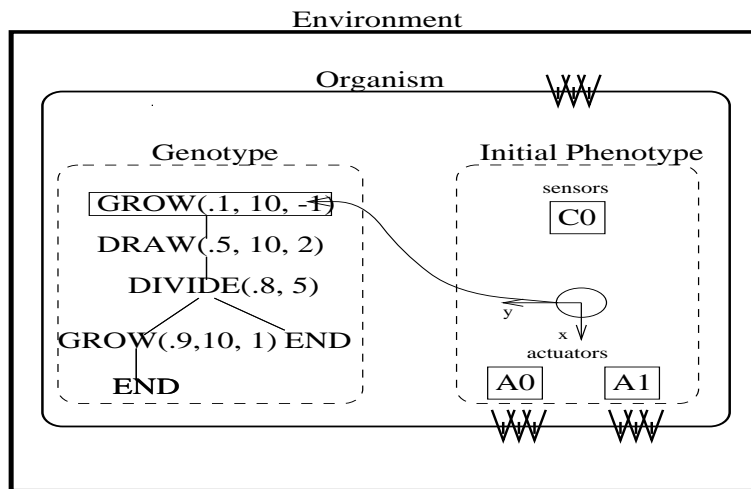
**Figure 13.** Growth process resulting from a randomly generated genotype in Nolfi and Parisi’s approach. The lowest layer corresponds to sensory neurons, central layers to internal neurons and the upper layer to motor neurons. a) Connections established during the developmental process. b) Functional network mapping sensory input into motor output obtained after elimination of redundancies or unused neurons.

corresponding architectures into functional sub-networks, and it allows non adaptive mutations - i.e., mutations that do not affect the animat’s fitness - to set the stage for future adaptive mutations by inducing changes in the control architecture of individuals of the preceding generations.

In Kodjabachian and Meyer (1996) a geometry-oriented variant of the *cellular* encoding scheme devised by Gruau (1994) is used to evolve the developmental program of a neural network that controls the locomotion of a six-legged animat. Each leg is associated with an angle sensor neuron - whose activity depends upon the leg’s angle with the vertical - and with three motor neurons - that are respectively responsible for moving the leg forward or backward, and for lifting the foot. Thus the evolutionary problem to be solved is that of designing a nervous system that will coordinate the activities of its sensor, motor and internal neurons, such as to ensure walking.


Within such a framework, each cell in a developing network has a copy of the chromosome that codes the developmental process, and each cell reads the chromosome at a different position. The chromosome is represented as a grammar tree, with ordered branches whose nodes are labeled with character symbols. These symbols represent instructions for cell development that act on the cell or on its connections to other cells. During a given developmental step, a cell executes the instruction referenced by the symbol it reads, and moves its reading head down in the tree. Figure 14 illustrates the corresponding process.

The developing architectures are afforded the possibility to draw connections with sensory and motor neurons whose functionalities and spatial organization are provided by the experimenter, depending upon the body plan of the animat (Figure 15). The development of the animat’s nervous system starts with six neurons that all execute the instructions of the same



**Figure 14.** The developmental encoding scheme used by Kodjabachian and Meyer. The genotype that specifies an organism's nervous system is encoded as a grammar tree whose nodes are specific developmental instructions. Within such chromosomes, mutations change one instruction into another, and crossovers swap branches. Each cell in the developing network reads the chromosome at a different position. The `DIVIDE` instruction causes a cell to divide and generate a daughter cell in a given direction and at a given distance - according to the arguments of the instruction. Likewise, `GROW` and `DRAW` instructions cause a cell to establish respectively efferent and afferent connections with other cells, in a given direction, at a given distance, and with a given synaptic weight. `END` instructions cause a cell to stop developing and become a mature neuron. Each neuron is modeled as a leaky-integrator, characterized by a time constant and a bias that can be set by instructions `SETTAU` and `SETBIAS`. More or less developmental steps are required to generate a phenotype, depending upon the length of the corresponding genotype.

genetically encoded program and that grow into six sub-networks, which may be interconnected.



**Figure 15.** The arrangement of sensors, actuators and initial neurons used by the developmental program of a six-legged animat. S: angle sensor neuron, which detects the angle a given leg makes with the vertical; F: foot motor neuron, which lifts a given foot when active; R: return strike motor neuron, which propels a given leg forward; P: power strike motor neuron, which propels a given leg backwards.

Figure 16 provides an example of a developmental program obtained after 700 generations, when the initial population was made up of 200 randomly-generated chromosomes, and when the fitness of any animat was given by the distance it covered during a given evaluation time.

Figure 17 and Figure 18 respectively illustrate the walking behavior of the corresponding animat and the activity of its sensors and actuators during walking.


Such results demonstrate that quite complex control architectures can be generated by much simpler developmental programs, thus tremendously reducing the size of the genetic search space that the evolutionary process explores. Hopefully, they open the way for the automatic

```

G(S(D(I(I(I(I(S(I(D(D(I(E,I(G(E),G(I(E,I(E,D(G(G(E)))))))))),
G(I(E,I(D(D(I(E,D(I(G(D(E)),E))))),D(G(S(D(G(D(D(E)))))))))),
D(D(G(G(E))))),D(I(G(I(I(E,G(E)),E)),G(G(G(G(E)))))),
G(D(D(I(I(I(G(D(G(E))),D(I(I(D(G(E)),E),I(I(G(G(S(S(E))),D(I(G(D(E)),E))),
E))),D(E)),G(G(G(E)))))),D(G(D(G(D(G(D(I(D(D(I(D(G(D(D(E))),D(G(E))))),
G(I(G(D(E)),G(G(G(G(S.(E))))))))))))))

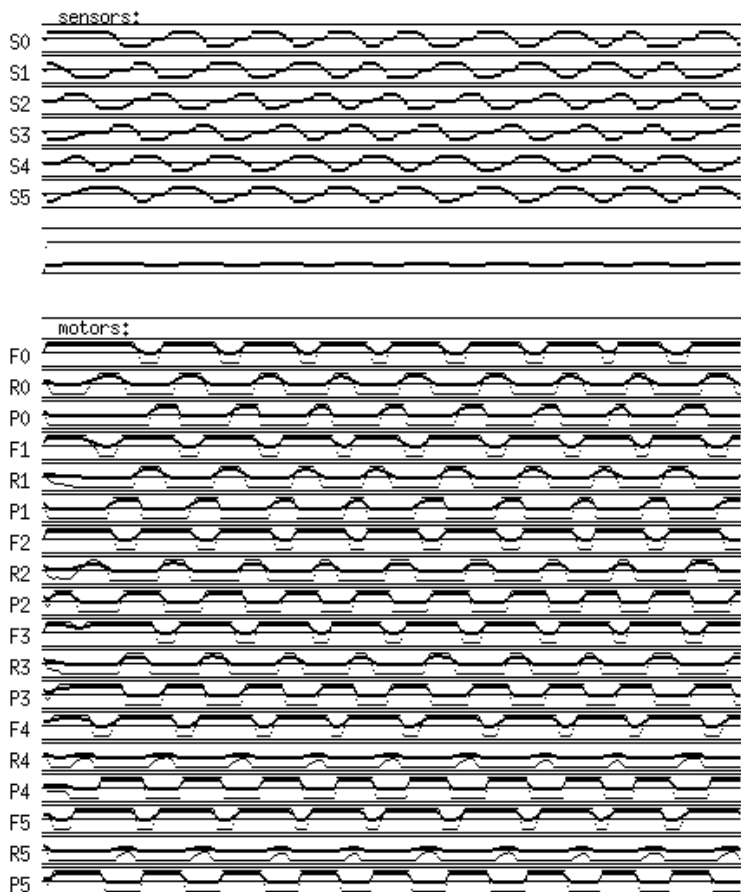
```

**Figure 16.** An evolved developmental program (instruction arguments excluded) of a six-legged animat. This program is expressed as a parenthesized expression, i.e., in a formalism that is equivalent to the grammar trees used in Figure 14. I: DIVIDE; G: GROW; D: DRAW; S: SETBIAS; E: END. SETTAU was not used in this run. Besides the 24 sensory and motor neurons, the neural network generated by this program contains 168 inter-neurons and 2222 connections after development.



**Figure 17.** The upper part of the figure shows a top view and a side view of the six-legged animat, at a given time and a given distance from start. The lower part illustrates the walking rhythm of the animat. Dotted lines indicate periods of protraction (leg off ground and moving forward relative to the body and ground) for each leg. By convention, these lines refer to legs in the following vertical order, from top to bottom: Rear Right, Middle Right, Front Right, Rear Left, Middle Left and Front Left. The rhythm exhibited by this animat is a variety of so-called *Wilson's rhythms* that are known to be exhibited by real insects.

generation of more cognitive control architectures than those that less biologically inspired methods have evolved so far.



**Figure 18.** Activities of the sensors, the motors and a given inter-neuron in the neural network that controls walking. The symbols identifying each neuron are the same as those on Figure 15.

## 5 Discussion

If the introduction of biomimetic mechanisms into the control architecture of an animat is likely to endow this animat with various important adaptive capacities that its engineered counterparts might lack, it is nevertheless clear that the corresponding improvements are restricted by the current limitations of biological knowledge.

In the case of navigation, for example, these limitations concern both the overall strategy a given animal uses in given set of circumstances, and the specialized mechanisms that implement the corresponding strategy. In this perspective, it is enlightening that several models (Zipser, 1986; Wilkie and Palfrey, 1987; Schmajuk, 1993; Blum and Abbott, 1996), implementing different navigation strategies, have been proposed in the literature to explain the same experimental results obtained with rats in the Morris water maze (Morris, 1981). Likewise, as far as the mechanisms are concerned, it is clear that, if the anatomy of the hippocampus is fairly well known,

its functional role - whether it is specialized in the storage of spatial representations or whether it acts as a general associative memory - is still under debate (Rolls, 1991), without mentioning the uncertainties about what sort of signals it receives as input and delivers as output. Likewise, the nature of the landmarks a given animal uses to localize or orient itself in its environment is seldom perfectly known, as it is still not established that such landmarks are actually recognized individually.

Concerning action selection mechanisms in animals, biological knowledge is also limited in many respects. In particular, it is still unclear whether the motivational systems of animals are organized in strict hierarchies, in flat structures, or in organizations somewhere between the two. Likewise, although it is usually admitted that these systems integrate both exteroceptive and interoceptive signals, possibly together with some memorized information, the exact way such an integration is performed is largely unknown. There is also an endless debate over the optimization of behavioral sequences in animals (Gould and Lewontin, 1979; Dennett, 1983; Dupre, 1987). In cases where an optimization mechanism seems to be at work, there is still the necessity of specifying whether it is genetically encoded in some behavioral rules or whether it is the end-product of some deliberate planning process. In both cases, the corresponding evaluation criteria remain entirely to be determined.

Finally, much remains to be understood about how the processes of evolution and development shape the control mechanisms of animals and about how they might interact with learning. In particular, considerable uncertainties remain about how developmental rules are encoded in a given genotype and about how the triggering of such rules depends upon the animal's interactions with its environment and upon its internal biochemistry (Gilbert, 1991). Likewise, although there is some evidence since Baldwin (1896) that individual learning can accelerate evolution in a non-Lamarckian manner, the corresponding mechanisms, and notably how they interact with development, are seldom investigated. Obviously such investigations are all the more difficult because it is hard to obtain relevant data about behaviors and mechanisms that characterized animals that lived thousands or even billions of generations ago.

## 6 Conclusions

The traditional top-down, analytical and reductionist approach of biology has undoubtedly led to impressive improvements in our understanding of many specific mechanisms that underlie the adaptive behaviors of animals. This paper provided several examples demonstrating how the corresponding knowledge may be useful when one wishes to integrate such mechanisms into a coherent functional model, capable of controlling animats that are simple, but complete, and that survive in environments as realistic as possible, in which they can feed, reproduce, escape predators, etc. However, this paper also stressed the shortcomings of current biological knowledge, which often lacks the bottom-up and synthetic perspective mandatory in animat modeling. Therefore, if artificial life has up to now gained more benefits from natural life than conversely, it can be hoped that both approaches will ultimately avail themselves of their complementary methods and objectives.

## 7 Bibliography

- Baldwin, J.M. (1896). A new factor in evolution. *The American Naturalist*. 30, 441-451.  
Blum, K.I. and Abbott, L.F. (1996). A model of spatial map formation in the hippocampus of the rat. *Neural Computation*. 8, 85-93.

- Blumberg, B. (1994). Action-Selection in Hamsterdam: Lessons from Ethology. In Cliff, D., Husbands, P., Meyer, J.A. and Wilson, S.W. (Eds.) From animals to animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior. The MIT Press.
- Blumberg, B. (1996). No Bad Dogs: Ethological Lessons for Learning in Hamsterdam. In Maes, P., Matarić, M., Pollack, J., Meyer, J.A. and Wilson, S. (Eds.) From animals to animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior. The MIT Press.
- Booker, L.B. (1988). Classifier systems that learn internal world models. *Machine Learning*. 3, 161-192.
- Brooks, R. and Maes, P. (Eds.) (1994). Artificial life IV. The MIT Press.
- Canny, J.F. (1988). The Complexity of Robot Motion Planning. The MIT Press.
- Cartwright, B.A. and Collett, T.S. (1983). Landmark learning in bees. Experiments and models. *Journal of Comparative Physiology A*. 15, 521-543.
- Cliff, D., Harvey, P. and Husbands, P. (1993). Explorations in evolutionary robotics. *Adaptive Behavior*. 2, 1, 73-110.
- Cliff, D., Husbands, P., Meyer, J.A. and Wilson, S.W. (Eds.) (1994). From animals to animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior. The MIT Press.
- Colgan, P. (1989). Animal Motivation. Chapman and Hall.
- Craik, J. (1943). The Nature of Explanation. Cambridge University Press.
- Dennett, D. (1983). Intentional systems in cognitive ethology: The Panglossian paradigm defended. *Behavioral and Brain Sciences*. 6, 343-390.
- Donnart, J.Y. and Meyer, J.A. (1994). A hierarchical classifier system implementing a motivationally autonomous animat. In Cliff, D., Husbands, P., Meyer, J.A. and Wilson, S.W. (Eds.) From animals to animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior. The MIT Press.
- Donnart, J.Y. and Meyer, J.A. (1996a). Reactive and planning rules learning in a motivationally autonomous robot. *IEEE Transactions on Man, Systems and Cybernetics. Part B: Cybernetics*. 26, 3, 381-395.
- Donnart, J.Y. and Meyer, J.A. (1996b). Spatial exploration, map learning, and self-positioning with MonaLysa. In Maes, P., Matarić, M., Pollack, J., Meyer, J.A. and Wilson, S. (Eds.) From animals to animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior. The MIT Press.
- Dupre, J. (1987). The Latest on the Best. The MIT Press.
- Franceschini, N., Riehle, A. and Le Nestour, A. (1989). In Stavenga, D.G. and Hardie, R. (Eds.). Facets of vision. Springer Verlag.
- Franceschini, N., Pichon, J.M. and Blanes, C. (1992). From insect vision to robot vision. *Phil. Trans. R. Soc. Lond. B*. 337, 283-294.
- Gibson, J. (1986). An Ecological Approach to Visual Perception. Lawrence Erlbaum.
- Gilbert, S.F. (1991). Developmental Biology. Sinauer.
- Goldberg, D. (1989). Genetic Algorithms in Search, Optimization, and Machine Learning. Addison Wesley.
- Gould, S. and Lewontin, R. (1979). The spandrels of San Marco and the Panglossian paradigm: A critique of the adaptationist programme. *Proceedings of the Royal Society of London, B*, 205, 581-598.
- Grefenstette, J. J. (1992). The Evolution of Strategies for Multiagent Environments. *Adaptive Behavior*. 1,1, 66-90.
- Gruau, F. (1994). Automatic definition of modular neural networks for adaptive behavior.

- Adaptive Behavior. 3,2, 151-184.
- Guillot, A. and Meyer, J.A. (1994). Computer Simulations of Adaptive Behavior in Animats. In Thalman, N. and Thalman, G. (Eds.) IEEE Proceedings on Computer Animation'94. IEEE Computer Society Press.
- Guillot, A. and Meyer, J.A. (1996). Synthetic animals in synthetic worlds. In Kunii, T. and Luciani, A. (Eds.). Synthetic Worlds. Wiley. In press.
- Harp, S., Samad, T. and Guha, A. (1989). Towards the genetic synthesis of neural networks. In Schaffer, J.D. (Ed.) Proceedings of the Third International Conference on Genetic Algorithms. Morgan Kaufmann.
- Holland, J.H. (1975). Adaptation in Natural and Artificial Systems. The University of Michigan Press.
- Holland, J.H., Holyoak, K.J., Nisbett, R.E. and Thagard, P.R. (1986). Induction. Processes of inference, learning, and discovery. The MIT Press.
- Hwang, Y.K. and Ahuja, N. (1992). Gross Motion Planning. ACM Computing Surveys. 24, 3, 219-291.
- Kaelbling, L.P. and Rosenschein, S. (1990). Action and Planning in Embedded Agents. In Maes, P. (Ed.). Designing Autonomous Agents: Theory and Practice. From Biology to Engineering and Back. The MIT Press.
- Kodjabachian, J. and Meyer, J.A. (1995). Evolution and development of control architectures in animats. Robotics and Autonomous Systems. 16, 161-182.
- Kodjabachian, J. and Meyer, J.A. (1996). Evolutionary Design of an Artificial Neural Network Controlling the Locomotion of a Six legged Animat. Proceedings of the ICML96 Workshop on Evolutionary Computation and Machine Learning. To appear.
- Koza, J.R. (1992). Genetic programming: On the programming of computers by means of natural selection. The MIT Press.
- Koza, J.R., Rice, J.P. and Roughgarden, J. (1992). Evolution of Food-Foraging Strategies for the Caribbean Anolis Lizard Using Genetic Programming. Adaptive Behavior. 1,2, 171-199.
- Langton, C.G. (Ed.). (1989). Artificial Life. Addison Wesley.
- Langton, C.G. (Ed.). (1992). Artificial Life III. Addison Wesley.
- Langton, C.G., Taylor, C., Farmer, J.D. and Rasmussen, S. (Eds.) (1992). Artificial Life II. Addison Wesley.
- Latombe, J.C. (1991). Robot Motion Planning. Kluwer.
- Ludlow, A.R. (1976) The behaviour of a model animal. Behaviour. 58, 131-172.
- Maes, P. (Ed.) (1990). Designing Autonomous Agents. Theory and Practice from Biology to Engineering and Back. The MIT Press.
- Maes, P. (1993). Behavior-Based Artificial Intelligence. In Meyer, J.A., Roitblat, H.L. Wilson, S.W. (Eds.). From animals to animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior. The MIT Press.
- Maes, P. (1994). Modeling Adaptive Autonomous Agents. Artificial Life. 1,135-162.
- Matarić, M.J. (1992). Integration of Representation Into Goal-Driven Behavior-Based Robots. IEEE Transactions on Robotics and Automation. 8, 3, 304-312.
- McCloskey, D.I. (1981). Corollary discharges: Motor commands and perception. In Brooks, V.B. (Ed.). Handbook of Physiology - The Nervous System. II. Motor control. American Physiological Society.
- McFarland, D. (1974). Motivational Control Systems Analysis. Academic Press.
- McFarland, D. (1981). The Oxford Companion to Animal Behaviour. Oxford University Press.
- McFarland, D. and Bösser, T. (1993). Intelligent behavior in Animals and Robots. The MIT



Press.

McNaughton, B.L., Chen, L.L. and Markus, E.J. (1991). Dead reckoning, landmark learning, and the sense of direction: A neurophysiological and computational hypothesis.

Meyer, J.A. (1995). The Animat Approach to Cognitive Science. In Roitblat, H. and Meyer, J.A. (Eds.). *Comparative Approaches to Cognitive Science*. The MIT Press.

Meyer, J.A. (1996). Artificial Life and the Animat Approach to Artificial Intelligence. In Boden, M. (Ed.). *Computational Psychology and Artificial Intelligence*. Academic Press.

Meyer, J.A. and Guillot, A. (1991). Simulation of adaptive behavior in animats: review and prospect. In Meyer, J.A. Wilson, S.W. (Eds.). *From animals to animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*. The MIT Press.

Meyer, J.A. and Guillot, A. (1994). From SAB90 to SAB94: four years of animat research. In Cliff, D., Husbands, P., Meyer, J.A. and Wilson, S.W. (Eds.) *From animals to animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. The MIT Press.

Meyer, J.A. Wilson, S.W. (Eds.) (1991). *From animals to animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*. The MIT Press.

Meyer, J.A., Roitblat, H.L., Wilson, S.W. (Eds.) (1993). *From animals to animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior*. The MIT Press.

Miglino, O., Nolfi, S. and Parisi, D. (1996). Discontinuity in evolution: how different levels of organization imply pre-adaptation. In Belew, R. and Mitchell, M. (Eds.). *Plastic Individuals in Evolving Populations*. Addison Wesley.

Miller, G.F., Todd, P.M. and Hedge, S.U. (1989). Designing neural networks using genetic algorithms. In Schaffer, J.D. (Ed.) *Proceedings of the Third International Conference on Genetic Algorithms*. Morgan Kaufmann.

Montana, D.J. and Davies, L. (1989). Training feedforward neural networks using genetic algorithms. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*. Morgan Kaufmann.

Moran, A., Moreno, J., Merelo, J. and Chacon, P. (Eds.) (1995). *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life*. Springer Verlag.

Morris, R.G. (1981). Spatial localization does not require the presence of local cues. *Learning and Motivation*. 12, 239-260.

Muller, R.U., Kubie, J.L., Bostock, E.M., Taube, J.S. and Quirk, G.J. (1991). Spatial firing correlates of neurons in the hippocampal formation of freely moving rats. In Paillard, J. (Ed.) *Brain and space*. Oxford University Press.

Nehmzow, U. (1995). Animal and Robot Navigation. *Robotics and Autonomous Systems*. 15, 71-81.

Nilsson, N.J. (1984). Shakey the robot. Technical Report 323. SRI AI Center.

O'Keefe, J. and Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford University Press.

Nolfi, S. and Parisi, D. (1992). Growing neural networks. Technical Report PCIA-91-15. Institute of Psychology. Rome.

Quirk, G.J., Muller, R.U. and Kubie, J.L. (1990). The firing of hippocampal place cells in the dark depends on the rat's recent experience. *Journal of Neuroscience*. 10, 6, 2008-2017.

Rolls, E.T. (1991). Functions of the primate hippocampus in spatial and nonspatial memory. *Hippocampus*. 1,3, 258-261.

Schmajuk, N.A., Thieme, A.D. and Blair, H.T. (1993). Maps, routes, and the hippocampus: a

neural network approach. *Hippocampus*. 3,3, 387-400.

Schwartz, J.T. and Yeap, C.K. (1987). *Advances in Robotics, Vol.I, Algorithmic and geometric Aspects of Robotics*. Lawrence Erlbaum.

Steels, L. and Brooks, R. (Eds.) (1995). *The Artificial Life Route to Artificial Intelligence. Building Embodied, Situated Agents*. Lawrence Erlbaum.

Taube, J.S., Muller, R.U. and Ranck, J.B. (1990a). Head direction cells recorded from the postsubiculum in freely moving rats. I. description and quantitative analysis. *J. Neurosci.* 10, 420-435.

Taube, J.S., Muller, R.U. and Ranck, J.B. (1990b). Head direction cells recorded from the postsubiculum in freely moving rats. II. Effects of environmental manipulations. *J. Neurosci.* 10, 436-447.

Taylor, C.E. (1992). *Fleshing Out*. In Langton, C.G., Taylor, C., Farmer, J.D. and Rasmussen, S. (Eds.) *Artificial Life II*. Addison Wesley.

Tinbergen, N. (1951). *The study of instinct*. Clarendon Press.

Toates, F. (1986). *Motivational Systems*. Cambridge University Press.

Tyrrell, T. (1993). The use of hierarchies for action selection. *Adaptive Behavior*. 1, 387-420.

Varela, F.J. and Bourgine, P. (Eds.) (1992). *Toward a Practice of Autonomous Systems*. Proceedings of the First European Conference on Artificial Life. The MIT Press.

Wan, H., Touretzky, D. and Redish, A. (1994). Towards a computational theory of rat navigation. In Mozer, M. , Smolensky, P., Touretzky, D., Elman, J. and Weigend, A. (Eds.). *Proceedings of the 1993 Connectionist Models Summer School*. Lawrence Erlbaum.

Wilson, D.M. (1966). Insect walking. *Annual Review of Entomology*. 11, 103-122.

Wilson, S.W. (1991). *The Animat Path to AI*. In Meyer, J.A. and Wilson, S.W. (Eds.) *From animals to animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*. The MIT Press.