From SAB90 to SAB94 : Four Years of Animat Research

Jean-Arcady Meyer and Agnès Guillot
Groupe de BioInformatique
Ecole Normale Supérieure, CNRS-URA 686
46, rue d'Ulm. 75230 Paris Cedex 05, France
(meyer@wotan.ens.fr - guillot@wotan.ens.fr)

Abstract

This paper builds on a previous review of significant research on adaptive behavior in animats. It summarizes the current state of the art and suggests some directions likely to provide interesting results in the near future.

1 Introduction

An animat is a simulated animal or a real robot whose rules of behavior are inspired by those of animals. It is usually equipped with sensors, with actuators, and with a behavioral control architecture that allow it to react or to respond to variations in its environment (internal or external), notably to those that might impair its chances of survival. The behavior of an animat is what the animat does. This is characterized by a sequence of actions which reflects the dynamic interplay between the animat and its environment, mediated through the animat's sensors and actuators.

The behavior of an animat is adaptive so long as it allows the animat to survive or to fulfill its mission. This requires that the animat's essential variables be monitored and maintained within their viability zone, an ability which can be enhanced, should the animat be capable of learning which actions elicit a positive or negative reward from the environment (Figure 1).

Since a previous review of significant research on adaptive behavior in animats (MEYE90), three conferences have been held on the subject. The present paper aims at updating the previous one and summarizes the current state of the art. For lack of space, it will only refer to the conference proceedings (SAB90, SAB92, SAB94), a strategical choice that, hopefully, should not leave any fundamental work aside. It is organized around the various components of Figure 1 and concludes with some possible directions for future work.

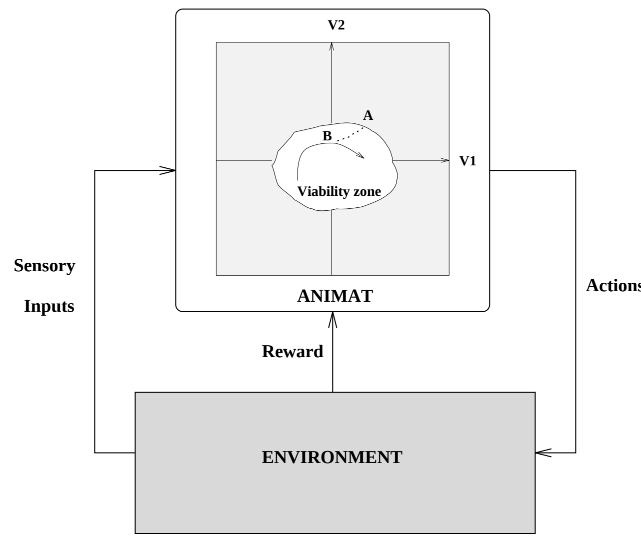

Figure 1: The interactions between an animat and its environment. The behavior of this animat is said to be adaptive because corrective action has been taken at point B so as to avoid crossing at point A the viability zone associated with the two essential variables V1 and V2. A corrective action can be preprogrammed or learned in the environment.

2 Actions

Not surprisingly - given the importance of robotic realizations in the animat community - the behavior most often exhibited by animats is motion. Thus, some animats move themselves in their environment like worms (HART92), others like bipeds (GREE94, VOGE92), quadrupeds (CHER92, CRUS90, DIGN94), hexapods (BEER90, CRUS90, CRUS92), or octopods (CRUS90). Other animats fly (MURA94), or swim (UTTA92), or even move themselves by means of brachiation (SAIT94). Still other animats are capable of moving an arm to explore their environment or to reach goals in it (FONE94, LIND92, MORA94).

If some robots are capable of pushing (KUBE92, NEHM94, SAHO94), or of grasping (MATA92) objects, their repertoire of actions is nevertheless quite limited. Animals exhibit far more diverse behaviors and many models aim at describing with as great an accuracy as possible the actual sensory-motor coordinating mechanisms which generate these behaviors. Such biomimetic models describe, for example, hierarchical dishabituation in toads (WANG90), prey-catching (ARBI90) or detour behavior (ARBI92) in anurans, visual saccades and head movements in the prey approach of salamanders (MANT90), spinal reflexes in the frog (GISZ92, GISZ94), human bodily adjustments in ball throwing tasks (ARBI94), phonotaxis in crickets (WEBB94), or fighting behavior in the Siamese fighting fish (HALP90).

Other behaviors are generated by models less dedicated to the description of precise biological mechanisms and are considered from a broader point of view. This is the case, for example, of various individual or collective behaviors involved in food seeking or foraging (ARKI92, ARKI94, BALL94, BURA94, CECC92, DENE90, DROG92, FLOR92, KUBO94, KURT90, MATA94, PARE90, SAUN94, SHAF94, STEE90, THER90, TODD92, TODD94), in herding and flocking (ARKI94, MATA92, WERN92), in predator-prey relationships (COLO92, DELU90, IBA92, MILLE94, REYN92), in mate finding (ROBB94), in exploring (GABO92) or in some sort of cleaning (ARKI94, MUNO94, PARK92) or sorting (DENEU90) tasks. Furthermore, various authors have devised special artificial environments (BERS94, BLUM94, LIN90, MAES90, TYRR90, WERN94) where animats must exhibit several of the above-mentioned behaviors in order to survive.

Within the context of social behavior, a few works are dedicated to the study of communication. For instance, de Bourcier and Wheeler (DEBO94) are interested in the signalling of aggressive intentions in animats and study how the cost of producing the corresponding signals influences the population dynamics of truth-tellers and bluffers. Likewise, McFarland (MCFA94) and Steels (STEE94) have devised a robotic ecosystem allowing the study of cooperative communications between truth-telling and bluffing robots. Yanco and Stein (YANC92) study how two robots can learn a private language in order to achieve coordinated movement and can adapt their language to cope with changing circumstances (see also PARK92 and ROBB94).

3 Perceptions

Several research efforts demonstrate the adaptive value of the long-evolved sensors of animals. For instance, Roitblat (ROIT90a, ROIT92) describes how the ear of the dolphin is used in solving intricate echolocation tasks, while Cliff (CLIF90), Lopez and Smith (LOPE92), and Mura and Franceschini (MURA94) show how the compound-eye of the fly can help tracking moving objects. Likewise, Osorio et al. (OSOR94) compare the architecture and function of the neural circuits controlling vision and olfaction in insects and suggest that these circuits implement fundamentally different computational strategies tailored to the complexities and unpredictabilities of the specific sensory signals they process. These authors conclude that the design of animats should take account of the environment in which an animat will act, and not be taken arbitrarily from a specific biological model.

Other aspects of the interaction between ad hoc sensors and perceptual strategies are also dealt with in the animat literature. For instance, Pierce and Kuipers (PIER90) studied how a robot with uninterpreted sensors and effectors could learn sequences of actions which can be taken to achieve a goal. In this case, a solution involves learning a function defined in terms of the sensory data. This function, which is maximized at the point corresponding to the goal, can be optimized by means of hill climbing. Parker (PARK92) studied how sensory filters could act as passive attention-focussing mechanisms in animats involved in a collective janitorial task. Such filters allow an animat to be more reactive to the actions of other agents. They convert communication about action of other animats into altered sensory readings, so that the animat ignores certain sensory readings or "hallucinates" others, and acts on the altered sensory feedbacks as if they were genuine. Another passive attention-focussing mechanism is studied by Foner and Maes (FONE94) whose animat learns an action model - i.e., what the perceptual consequences of a given action in a given sensory context will be. Instead of attempting to learn all that there is to know about experiencing the world, the system focusses its attention on the important aspects of its current experience and memory. The corresponding architecture implements both perceptual selectivity - which restricts the set of sensor data to which the agent attends - and cognitive selectivity - which restricts the set of internal structures that is adapted. An active attention-focussing mechanism is studied by Cliff (CLIF90) whose "computational Hoverfly" is capable of looking around and employs a foveal sampling strategy with gaze-control mechanisms repositioning the limited high-resolution area of the visual field (see also ARBI90, ARBI94, MANT90). Cohen and Atkin (COHE94) study an attention-focussing mechanism in time, instead of space. They show that paying attention to the environment according to an interval reduction monitoring strategy - i.e., monitor more frequently near a goal than far from it - is more efficient that periodic monitoring in many circumstances. They suggest that it might be worth implementing such a strategy in animats.

Finally other authors study how an evolutionary process could shape the sensory apparatus of an animat (CLIF92, HARV94, KURT90, REYN94). Kurtz (KURT90) describes the evolution, within a population of foraging animats, of the allele frequencies of three genes coding for internal, external, and relational information gathering abilities. Cliff et al. (CLIF92) describe how the angle of acceptance and the eccentricity of the photoreceptors of a visually guided robot could evolve - together with the architecture of its nervous system - and improve the robot's success at avoiding collisions with the walls surrounding its environment.

4 Architectures

In order to behave adaptively, an animat has to do the right thing at the right moment and, ideally, such a choice must depend upon the animat's perceptions of the external environment, on its physiological or internal state, on the consequences of its current behavior and on the expected consequences of its future behavior. In other words, the animat must want something and it must be endowed with a motivational system. Such a system, which selects at every moment a goal to be pursued and organizes the animat's commerce with it (TOAT90), necessitates a memory of the past consequences of the animat's activities and, eventually, a planner which relates the current behavioral choice to its future consequences (DONN94).

In several realizations, the overall architecture and the inner workings of the motivational system of an animat are fixed by its human designer or by an evolutionary process. In traditional robotics and AI, the motivational system of an intelligent agent is usually designed as a centrally controlled organization of functional modules - such as perception, modeling, planning and execution - which sequentially process information from sensors to actuators. Quite differently, animat designers conceive the motivational system of an intelligent and robust animat as a distributed agency of processing units which collectively exhibit the required functionalities. Each such module can be connected directly to sensors and actuators, and works in parallel with the others, thus endowing the animat with minimal sensory-motor faculties in the case of an internal breakdown affecting some other modules.

Often, the corresponding architectures are decisional hierarchies, more or less loosely inspired by the way many ethologists have described the action selection mechanisms of animals (BOOK90, CHER92, DIGN94, SCHN90). They involve several systems, each devoted to a specific function affecting the survival ability of the animat, which are organized as loosely overlapping hierarchies of decision centers. Each center, at each level, can take input from many different internal and external sensory stimuli and can send facilitating sensory signals to the centers of the level below. The activation of a center at a given level depends upon the combined effect of the sensory stimuli, of the facilitating control signals provided by the centers of the level above, and of inhibitory signals sent by the other centers of this level. These inhibitory signals implement some sort of winner-take-all mechanism, according to which only one center can be active at any given level. At the higher levels of these hierarchies, the activation of a center triggers a very general decision (i.e. reproduce), while such a decision becomes increasingly more specific as one goes down (i.e. court), until a final motor decision (i.e. mate) is made at the bottom level. The linear architectures in SAUN94 and HALP90 are special cases of such decisional hierarchies. Other special cases are those which implement an arbiter whose task is to merge in a single action the suggestions to perform one behavior or another afforded by several dedicated modules (BOOK90, COLL92, TENE92, YAMA94).

In free-flow hierarchies (TYRR92), the nodes at each level express multiple preferences for each of a set of lower-level candidates, rather than making a decision as to which one is most suitable. The flow of preferences travels through several nodes at each level, and it is only at the bottom level than an arbitration is done and an action selected. Another hierarchical architecture, which borrows characteristics from the two above logics, is described in BLUM94. Architectures based on feedforward neural networks (CECC92, COLL90, FLOR92, GAUS94, NEHM90, NEHM94, PARE90, VONK90, WERN92, WERN94) are special cases of free-flow hierarchies, where the top layer receives the sensory inputs and the bottom layer delivers the motor outputs.

The motivational systems of other animats are non-hierarchical and differ from each other by the degree of arbitration to which their behavioral modules are submitted. Thus, no arbitration occurs in force-field controlled architectures (ARBI90, ARBI92, ARKI92, ARKI94) whose behavioral modules each contribute a vector that is related to the animat's current goal. These vectors are normalized and summed before being transmitted to the final execution module. In subsumption architectures (KUBE92, MATA90, MATA92), each behavioral module is also active at any time, but can have its output suppressed by that of another module. In PARK92, the architecture comprises several "behavior sets" - each organized in a subsumption manner - which compete to control what the animat does. Once a behavior set is activated, other behavior sets are suppressed, so that only one behavior set is active at a time. In MAES90 or in GISZ92, the motivational system of an animat is organized as a spreading-activation network whose nodes trigger the various behaviors the animat can engage in.

These nodes are more or less active, depending upon the amount of "energy" that flows through each of them, but a winner-take-all procedure arbitrates among them and selects the behavior actually exhibited. In BEER90, no general arbitration mechanism is provided. Rather, the overall coherency of the animat's behavior depends upon specific inhibitory links which allow some active nodes to prevent others from being activated at the same time.

Although the animat community is still missing general knowledge about what kind of adaptive abilities are afforded by what kind of architecture, one can find useful contributions to such a knowledge in several papers (BROO90, BROO94, MAES92, PFEI92, ROIT90b, SCHN90, WALT90).

5 Learning

5.1 Reinforcement Learning

If the inner workings of an action selection mechanism are not fixed by the designer or by an evolutionary process, they must be learned. The corresponding animats must memorize that a given action is more rewarding than another in given circumstances, so that they will preferentially exhibit the former action instead of the latter, should the same circumstances be encountered again in the future. Furthermore, they must be capable of adjusting their behavioral policy if the relative merits of the various actions they can chose from change for whatever reason. In other words, a greater autonomy and considerable adaptive abilities are afforded to animats that are capable of learning by reinforcement.
Several systems, with increasingly adaptive functionalities, have been devised for such a purpose (see SUTT90 and LIN90 for comparisons). For instance, policy-only systems, stochastic automata (KUBO94), as well as some systems based on neural networks (GAUS94, NEHM92, NEHM94), learn a mapping from the animat's sensations to its actions and specify what the animat will do in each situation at its current stage of learning. Classifier Systems with internal messages, as well as some systems based on recurrent neural networks (LIN92, YAMA94), specify what to do according to current sensations and to a memory of past sensations and actions. Reinforcement-comparison systems and some simple Classifier Systems (VENT94) learn which immediate reward will be gained from doing a specific action in a specific sensory context. Temporal Difference algorithms (BAIR92, MILLA94, KLOP92, ROSE92), Adaptive Heuristic Critic systems (PIPE94), Q-learning systems (BERS94, DORI94, LIN92, MUNO94, SAIT94), Associative Control Process networks (BAIR92, KLOP92) and Classifier Systems (DONN94, DORI94) learn which long-term cumulative reward can be expected from a given action in a given sensory context. Finally, Dyna systems (PENG92, SUTT90), some dedicated neural networks (CHES94, SCHMI90) and some varieties of Classifier Systems (RIOL90) not only learn such a reward, but also learn a world model - that is, how the world changes as an animat acts in it. They accordingly predict which cumulative reward and which new sensory inputs are to be expected from a given action in a given sensory context.

Various reinforcement learning algorithms have been implemented within the above mentioned motivational systems. For instance, Digney and Gupta (DIGN94) use a Temporal Difference learning algorithm to let a hierarchical action selection mechanism coordinate the locomotion of a quadruped animat. Likewise, Gistzter (GISZ94) combines a spreading-activation network, a force-field mechanism for command fusion, and a reinforcement signal to simulate the wiping behaviors of the frog.

5.2 Associative Learning

Associative learning often occurs in animats within the context of a navigational task. Such a task is usually accomplished with the aid of landmarks, which must be first categorized and recognized, then used as stepping stones, or included into a cognitive map, in order to allow the animat to reach a goal. Smart and Hallam (SMAR94) discuss the role of proprioceptive and featural cues in location recognition tasks by rats. Such cues are used by the robot of Nehmzow and Smithers (NEHM90) to train an associative Kohonen network, by the robot of Yamauchi and Beer (YAMA94) to train a dynamical neural network, and by the robot of Mataric (MATA90) to train a spreading-activation network (see also KORT92). These networks allow landmark discrimination and the building of an internal representation of the environment as the robots move around (see also MORA94). Having learned such internal representations, the robots are able to recognize where they are situated and - at least in the case of Mataric's robot - to navigate from a starting position to a goal. Such a navigation, however, implies passing through places already visited and characterized (see also DONN94 and MAZE92). The ability to estimate the distances or the visual angles of some landmarks, together with the use of cognitive maps, allow other animats to orient themselves according to these landmarks and to navigate through places never visited before (YEAP90, PRES94). For instance, the navigation system of Benhamou et al. (BENH94) relies on the use of both place cells and direction cells, that of Prescott and Mayhew (PRES92) relies on landmarks characterized by their coordinates, and that of Schmajuk and Blair (SCHMA92) relies on landmarks characterized by their visual angle. The latter model also predicts the dynamics of several spatial learning tasks.

Associative learning in animats also occurs in a variety of situations incurring habituation (SCUT94, STAD92, WANG90), sensitization (SCUT94) or conditioning (HALP90, SCHMA94, SCUT94). Its adaptive value is studied in TODD90. In the work of Aitken (AITK94), an animat detects correlations between current motor actions and future sensory inputs, and it learns to act and perceive by exposure to the sensory consequences of desired motor sequences - rather than exposure to the desired motor sequences themselves. In the presence of a teacher that provides desirable sensory sequences - as a result of its own motor sequences - the animat is capable of learning complex behavioral sequences by imitation.

6 Planning

An animat that is capable of planning will exhibit a behavioral sequence in which a specific action succeeds another, not as a mere reaction to the new environmental conditions brought about by the previous action, but because the whole sequence seems likely to fulfill the animat's current goal. In other words, a planning animat is capable of assessing the future consequences of its actions, an ability that greatly enhances its chances of surviving beyond those afforded by mere trial-and-error learning.

One way to plan a behavioral sequence is to make use of a world model that behaves like the world. With such a model, an animat can indeed perform hypothetical experiments (that is, "experiments in its head") and search for a behavioral sequence likely to reach the goal. The Dyna architecture, and the related systems which have been cited above, all allow this sort of planning.

Other planning abilities are afforded to animats that are capable of deciding what to do according to their current sensory state and their current goal or subgoal. If such animats are also capable of transforming knowledge about previously learned action sequences into appropriate subgoals for new problems, then they will solve new tasks through composition of solutions for older tasks. Thus, actions performed at a given time depend upon both a subgoal that seems likely to contribute to the achievement of the overall goal and upon the current sensory input. Accordingly, the corresponding animats exhibit reactive planning (DONN94, SCHMI92; see also SAHO94).

Another approach is described in RIBE92, where a plan is collectively achieved by a group of interacting agents, each having only local information about the state of the world and having some means of transmitting limited information to the others (see also WEIS92).

7 Evolution

As mentioned above, the overall architecture and the inner workings of an animat are sometimes determined by an evolutionary process.

Some realizations (DELU90, KURT90) involve a traditional genetic simulation model, which allows the monitoring of the dynamics of allele frequencies and the study of speciation phenomena. Other realizations implement a genotype to phenotype mapping with the help of genetic programming techniques. Such approaches seek computer programs which, given some sensors and actuators, allow an animat to exhibit behavioral sequences that improve over time, according to a given fitness function. Thus Koza (KOZA90) evolved the behavioral programs of an artificial ant capable of following a particular pheromone trail, of a pursuer chasing an optimal evader, of an evader racing away from an optimal pursuer, and of two minimaxing players. Other applications of genetic programming are to be found in REYN92 and REYN94. In particular, Reynolds (REYN94) evolves a program which determines the optimal number and orientation of the sensors that allow an animat to follow a corridor while avoiding collisions with the walls. Elsewhere, the genotype-phenotype mapping is implemented in a classical neural network (COLL90, FLOR92, PARE90, TODD90, WERN92, WERN94) or in a dynamic neural network (CLIF92, HARV92, HARV94, YAMA94). For instance, in the work of Floreano (FLOR92), a feedforward neural network allows an evolved animat to exhibit a nest-based foraging behavior. This work shows that an evolutionary process may eventually not lead to an optimal behavior, because new architectures are necessarily built upon older ones which, although perfectly adapted to previous environmental conditions, may not be the best stepping-stones for reaching solutions to new environmental conditions.

The adaptive solutions sought by an evolutionary process usually involve predetermined sensors and actuators. As previously mentioned, however, several realizations permit the animat's sensors to evolve. Todd and Wilson (TODD92) describe a framework for exploring the evolution of adaptive behaviors in response to different environment structures, in which not only the sensors, but also the actuators and the behavioral rules of a given animat would all be allowed to evolve.

Although the improvement of the solutions to a specific problem of adaptation is usually dependent upon an explicit fitness function that helps select good solutions - that will be allowed to reproduce - and bad solutions - that will be eliminated - during the course of evolution, some realizations do not rely on such a function. They accordingly rely upon an implicit fitness and let

reproduction and selection depend upon the internal dynamics of an ecosystem (DELU90, PARE90, WERN92, WERN94). Likewise, although the majority of realizations involve a simulated animat, those described in HARV94 and FLOR94 involve real robots. Thus, Harvey et al. (HARV94) encode in a pair of chromosomes both the architecture of the neural network which controls the behavior of a gantry robot and the physical organization of the robot's visual system, in order to evolve various visually-guided behaviors. Likewise, Floreano and Mondada (FLOR94) evolve the neural architecture controlling the behavior of a real robot in a navigation and obstacle-avoidance task. Although the robot has a circular shape and is equipped with two wheels that rotate at equal speeds in each direction, the evolutionary process generates a frontal direction of motion that corresponds to the side where the robot has more sensors. Accordingly, the robot faces obstacles with the side that provides a finer resolution and a larger visual angle.
Several researchers have been interested in coevolutionary processes (DELU90, FLOR92, KOZA90, MILLE94, ROBB94, WERN92). For instance, Robbins (ROBB94) describes the evolution of a common mate-finding communication protocol within an ecosystem where immobile females guide blind moving males. Each animat is endowed with a message chromosome and an interpretation chromosome. Within such an ecosystem, it turns out that various parasites can accelerate the evolution of the protocol by increasing the evolutionary pressure put on the evolving animats.

Finally, two fundamental approaches to the mechanisms of evolution are described in TODD90 and MILLE92. The work of Todd and Miller (TODD90) helps to explore under what conditions the associative learning abilities of an animat could prove adaptive and evolve. Miller and Todd (MILLE92) compare the evolutionary forces of natural selection to those of sexual selection and explore how the dynamics of evolution may interact with the mechanisms of cognition - like those involved in directional mate preferences.

8 External environment

A given behavior might prove adaptive in one environment but not in another. Therefore, if one wants to give the animat community a chance to generate something other than fragmentary knowledge about ad hoc solutions to specific adaptive problems in specific environments (MEYE90), and if one seeks useful generalizations, one should subscribe to the general research program described by Todd and Wilson (TODD92). This program aims at characterizing the important features of environment structure in terms of the adaptive behavior they elicit and suggests studying how these features correlate with the sensors, the actuators, and the control structures that evolve in different environments. A specific application of this program (TODD94) demonstrates that very simple animats can survive, with neither sensors nor memory, in a variety of changing and unpredictable environments and that the challenges of different types of environments are best met by different probability distributions of blind actions.
Another approach to a characterization of environments is proposed by Wilson (WILS90) on the basis of the indeterminacy of an environment with respect to the sensory capabilities of its agents. Still another suggestion is made by Littman (LITT92) who characterizes an environment by the simplest agent that can possibly achieve optimal reinforcement in it. The corresponding results show that the degree to which the adaptive problem raised by a given environment can be partially solved by a suboptimal agent may also be a significant measure of environmental difficulty (see also LITT94).
Finally, the work of Horswill (HORS92) is not unrelated to the above concerns, insofar as it aims at characterizing the environmental features that simplify the computational problem facing an animat.

9 Prospects

The short term goal of animat research is to discover architectures and working principles that allow a real animal, a simulated animal, or a robot to exhibit a behavior that solves a specific problem of adaptation in a specific environment. Undoubtedly, the present state of the art already contributes usefully to this goal.
An intermediate goal of animat research is to generalize this knowledge and make progress towards understanding what architectures and working principles can allow an animat to solve what kind of problems in what kind of environments (MEYE90). Although current attempts at categorizing architectures and environments undoubtedly constitute valuable steps in this direction, other useful contributions can probably be expected from additional research into conceptual frameworks likely to help in abstracting over incidental differences between specific architectures, behaviors, and environments. Contributions in GALL92, KISS90, SLOT94 are probably useful in this respect.
The ultimate goal of animat research - at least in our opinion - is to contribute to our understanding of the adaptive value and working principles of human cognition (MAES92, WILS90). There is certainly a lot to be gained from asking how the highest cognitive abilities of man depend upon the evolution of the simplest cognitive abilities and adaptive behaviors of animals. From such a perspective, contributions in BROO94, BURA94, DEUG92, ROIT94, SHAF94, TOAT94, VERS92 should provide insights into cognition, representations, and the inner workings of the human mind. There is also a lot

to be gained from asking how processes of development and learning - those occuring at the level of the individual - interact with the process of evolution - occuring at the level of the population. Partial answers are to be found in MEYE90 and TODD90, but the field is lacking any simulation involving development, at least from this perspective (but see FONE94 and RUTK94). Finally, there is also a lot to be gained from asking how physical and social environments each contribute to the evolution of cognitive abilities. Several research efforts are aimed in such a direction (DEBO94, MCFA94, MILLE94) and might, in particular, help in studying the rise of a Machiavellian intelligence within a community where cheating animats can evolve.

10 Conclusion

Animat research is an active field of investigation which has already contributed several useful practical results and might provide valuable contributions to the understanding of human cognition in the future. However, the domain is in definite need of theoretical advances that could provide useful generalizations of still highly disparate pieces of knowledge.

References

[AITK94] A. M. Aitken. An architecture for learning to behave. In [SAB94].

[ARBI90] M.A. Arbib and A. Cobas. Schemas for prey-catching in frog and toad. In [SAB90].

[ARBI92] M. A. Arbib and H.-B. Lee. Anuran visuomotor coordination for detour behavior: From retina to motor schemas. In [SAB92].

[ARBI94] M. A. Arbib, N. Schweighofer, and W. T. Thach. Modeling the role of cerebellum in prism adaptation. In [SAB94].

[ARKI92] R. C. Arkin and J. D. Hobbs. Dimensions of communication and social organization in multi-agent robotic systems. In [SAB92].

[ARKI94] R. C. Arkin and K. Ali. Integration of reactive and telerobotic control in multi-agent robotic systems. In [SAB94].

[BAIR92] L. C. Baird III, and A. H. Klopf. Extensions of the associative control process (ACP) network: Hierarchies and provable optimality. In [SAB92].

[BALL94] N. Ball Organizing an animat's behavioral repertoire using kohonen feature maps. In [SAB94].

[BEER90] R. D. Beer and H. J. Chiel. The neural basis of behavioral choice in an artificial insect. In [SAB90].

[BENH94] S. Benhamou, P. Bovet, and B. Poucet. A place navigation algorithm based on elementary computing procedures and associative memories. In [SAB94].

[BERS94] H. Bersini. Reinforcement learning for homeostatic endogenous variables. In [SAB94].

[BLUM94] B. Blumberg. Action-selection in Hamsterdam: Lessons from ethology. In [SAB94].

[BOOK90] L. B. Booker. Instinct as an inductive bias for learning behavioral sequences. In [SAB90].

[BROO90] R. A. Brooks. Challenges for complete creature architectures. In [SAB90].

[BROO94] R. A. Brooks. Coherent behavior from many adaptive processes. In [SAB94].

[BURA94] S. Bura. Minimeme: of life and death in the noosphere. In [SAB94].

[CECC92] F. Cecconi and D. Parisi. Neural networks with motivational units. In [SAB92].

[CHER92] S. Cherian and W. O. Troxell. A neural network based behavior hierarchy for locomotion control . In [SAB92].

[CHES94] W. Chesters and G. M. Hayes. Connectionist environment modelling in a real robot. In [SAB94].

[CLIF90] D. Cliff. The computational hoverfly: A study in computational neuroethology. In [SAB90].

[CLIF92] D. Cliff, P. Husbands, and I. Harvey. Evolving visually guided robots. In [SAB92].

[COHE94] P. R. Cohen and M. Atkin. Preliminary evidence that the interval reduction monitoring strategy is general. In [SAB94].

[COLL90] R. J. Collins and D. R. Jefferson. Representations for artificial organisms. In [SAB90].

[COLO92] M. Colombetti and M. Dorigo. Learning to control an autonomous robot by distributed genetic algorithms. In [SAB92].

[CRUS90] H. Cruse. Coordination of leg movement in walking animals. In [SAB90].

[CRUS92] H. Cruse, U. Mueller-Wilm, and J. Dean. Artificial neural nets for controlling a 6-legged walking system. In [SAB92].

[DEBO94] P. de Bourcier and M. Wheeler. Signalling and territorial agression: An investigation by means of synthetic behavioral ecology. In [SAB94].

[DELU90] F. De Luigi and V. Maniezzo. The rise of interaction: Intrinsic simulation modeling of the onset of interacting behaviour. In [SAB90].

[DENE90] J. L. Deneubourg, S. Goss, N. Franks, A. Sendova-Franks, C. Detrain, and L. Chretien. The dynamics of collective sorting: Robot-like ants and ant-like robots. In [SAB90].

[DEUG92] D. Deugo and F. Oppacher. An evolutionary approach to cognition. In [SAB92].

[DIGN94] B. L. Digney and M. M. Gupta. A distributed adaptive control system for a quadruped mobile robot. In [SAB94].

[DONN94] J.-Y. Donnart and J.-A. Meyer. A hierarchical classifier system implementing a motivationally autonomous animat. In [SAB94].

[DORI94] M. Dorigo and H. Bersini. A comparative analysis of Q-Learning and classifier systems. In [SAB94].

[DROG92] A. Drogoul and J. Ferber. From Tom Thumb to the dockers: Some experiments with foraging robots. In [SAB92].

[FLOR92] D. Floreano. Emergence of nest-based foraging strategies in ecosystems of neural networks. In [SAB92].

[FLOR94] D. Floreano and F. Mondada. Automatic creation of an autonomous agent: Genetic evolution of a neural network Driven Robot. In [SAB94].

[FONE94] L. N. Foner and P. Maes. Paying attention to what's important: Using focus of attention to improve unsupervised learning. In [SAB94].

[GABO92] L. Gabora. Should I stay or should I go: Coordinating biological needs with continuously-updated assessments of the environment. In [SAB92].

[GALL92] J. C. Gallagher and R. D. Beer. A qualitative dynamical analysis of evolved locomotion controllers. In [SAB92].

[GAUS94] P. Gaussier and S. Zrehen. A topological neural map for on-line learning: Emergence of obstacle avoidance in a mobile robot. In [SAB94].

[GISZ92] S. Giszter. Behavior networks and force fields for simulating spinal reflex behaviors of the frog. In [SAB92].

[GISZ94] S. Giszter. Reinforcement tuning of action synthesis and selection in a 'virtual frog'. In [SAB94].

[GREE94] P. R. Green. How to watch your step: Biological evidence and a model. In [SAB94].

[HALP90] J. R. P. Halperin. Machine motivation. In [SAB90].

[HART92] R. Hartley. Propulsion and guidance in a simulation of the worm C. elegans. In [SAB92].

[HARV92] I. Harvey, P. Husbands, and D. Cliff. Issues in evolutionary robotics. In [SAB92].

[HARV94] I. Harvey, P. Husbands, and D. Cliff. Seeing the light: Artificial evolution, real vision. In [SAB94].

[HORS92] I. Horswill. A simple, cheap, and robust visual navigation system . In [SAB92].

[KISS90] G. Kiss. Autonomous agents, AI and chaos theory. In [SAB90].

[KLOP92] A. H. Klopf, J. S. Morgan, and S. E. Weaver. Modeling nervous system function with a hierarchical network of control systems that learn. In [SAB92].

[KORT92] D. Kortenkamp and E. Chown. A directional spreading activation network for mobile robot navigation. In [SAB92].

[KOZA90] J. R. Koza. Evolution and coevolution of computer programs to control independently-acting agents. In [SAB90].

[KUBE92] C. R. Kube and H. Zhang. Collective robotic intelligence. In [SAB92].

[KUBO94] M. Kubo and Y. Kakazu. Learning coordinated motions in a competition for food between ant colonies. In [SAB94].

[KURT90] C. Kurtz. The evolution of information gathering: Operational constraints. In [SAB90].

[LIN90] L.-J. Lin. Self-improving reactive agents: Case studies of reinforcement learning frameworks. In [SAB90].

[LIN92] L.-J. Lin and T. Mitchell. Reinforcement learning with hidden states. In [SAB92].

[LIND92] A. Linden and F. Weber. Implementing inner drive through competence reflection. In [SAB92].

[LITT92] M. L. Littman. An optimization-based categorization of reinforcement learning environments. In [SAB92].

[LITT94] M. L. Littman. Memoryless policies: Theoretical limitations and practical results. In [SAB94].

[LOPE92] L. R. Lopez and R. E. Smith. Evolving artificial insect brains for artificial compound eye robotics. In [SAB92].

[MAES90] P. Maes. A bottom-up mechanism for behavior selection in an artificial creature. In [SAB90].

[MAES92] P. Maes. Behavior-based artificial intelligence. In [SAB92].

[MANT90] G. Manteuffel. A biological visuo-motor system: How dissimilar maps interact to produce behavior. In [SAB90].

[MATA90] M. J. Mataric. Navigating with a rat brain: A neurobiologically-inspired model for robot spatial representation. In [SAB90].

[MATA92] M. J. Mataric. Designing emergent behaviors: From local interactions to collective intelligence. In [SAB92].

[MATA94] M. Mataric. Learning to behave socially. In [SAB94].

[MAZE92] E. Mazer, J.-M. Ahuactzin, E.-G. Talbi, and P. Bessiere. The Ariadne's clew algorithm. In [SAB92].

[MCFA94] D. McFarland. Towards robot cooperation. In [SAB94].

[MEYE90] J.A. Meyer and A. Guillot. Simulation of adaptive behavior in animats: Review and prospect. In [SAB90].

[MILLA94] J. Millan. Learning efficient reactive behavioral sequences from basic reflexes in a goal-directed autonomous robot. In [SAB94].

[MILLE92] G. F. Miller and P. M. Todd. Evolutionary interactions among mate choice, speciation, and runaway sexual selection. In [SAB92].

[MILLE94] G. Miller and D. Cliff. Protean behavior in dynamic games I: Arguments for the co-evolution of pursuit-evasion tactics. In [SAB94].

[MORA94] P. Morasso and V. Sanguineti. Self-organizing topographic maps and motor planning. In [SAB94].

[MUNO94] R. Munos and J. Patinel. Reinforcement learning with dynamic covering of state-space: Partitioning Q-learning. In [SAB94].

[MURA94] F. Mura and N. Franceschini. Visual control of altitude and speed in a flying agent. In [SAB94].

[NEHM90] U. Nehmzow and T. Smithers. Mapbuilding using self-organising networks in really useful robots

[NEHM92] U. Nehmzow, T. Smithers, and B. McGonigle. Increasing behavioural repertoire in a mobile robot. In [SAB92]. In [SAB90].

[NEHM94] U. Nehmzow and B. McGonigle. Achieving rapid adaptations in robots by means of external tuition. In [SAB94].

[OSOR94] D. Osorio, W. Getz, and J. Rybak. Insect vision and olfaction: Different neural architectures for different kinds of sensory signal? In [SAB94].

[PARE90] J. Paredis. The evolution of behavior: Some experiments. In [SAB90].

[PARK92] L. E. Parker. Adaptive action selection for cooperative agent teams. In [SAB92].

[PENG92] J. Peng and R. J. Williams. Efficient learning and planning within the Dyna framework . In [SAB92].

[PFEI92] R. Pfeifer and P. Verschure. Designing efficiently navigating non-goal-directed robots. In [SAB92].

[PIER90] D. Pierce and B. Kuipers. Learning hill-climbing functions as a strategy for generating behaviors in a mobile robot. In [SAB90].

[PIPE94] A. G. Pipe, T. C. Fogarty, and A. Winfield. A hybrid architecture for learning continuous environmental models in maze problems. In [SAB94].

[PRES92] T. J. Prescott and J. E. W. Mayhew. Building long-range cognitive maps using local landmarks. In [SAB92].

[PRES94] T. J. Prescott. Spatial learning and representation in animats. In [SAB94].

[REYN92] C. W. Reynolds. An evolved, vision-based behavioral model of coordinated group motion. In [SAB92].

[REYN94] C. W. Reynolds. Evolution of corridor following behavior in a noisy world. In [SAB94].

[RIBE92] F. Ribeiro, J.-P. Barthes, and E. Oliveira. Dynamic selection of action sequences. In [SAB92].

[RIOL90] R. L. Riolo. Lookahead planning and latent learning in a classifier system. In [SAB90].

[ROBB94] P. Robbins. Parasitism in an artificial world. In [SAB94].

[ROIT90a] H. L. Roitblat, P. W. B. Moore, P. E. Nachtigall, and R. H. Penner. Biomimetic sonar processing: From dolphin echolocation to artificial neural networks. In [SAB90].

[ROIT90b] H. L. Roitblat. Cognitive action theory as a control architecture. In [SAB90].

[ROIT92] H. L. Roitblat, P. W. B. Moore, D. H. Helweg and P. E. Nachtigall. Representation and processing of acoustic information in a biomimetic neural network. In [SAB92].

[ROIT94] H. L. Roitblat. Mechanisms and process in animal behavior: Models of animals, animals as models. In [SAB94].

[ROSE92] B. E. Rosen and J. M. Goodwin. Dynamic flight control with adaptive course coding. In [SAB92].

[RUTK94] J. Rutkowska. Emergent functionality in infant action and development. In [SAB94].

[SAB90] J.A. Meyer and S.W. Wilson (Eds). 1991. From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior. The MIT Press/Bradford Books.

[SAB92] J.A. Meyer, H.L. Roitblat, and S.W. Wilson (Eds). 1993. From Animals to Animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior. The MIT Press/Bradford Books.

[SAB94] D. Cliff, P. Husbands, J.A. Meyer, and S.W. Wilson (Eds). 1994. From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior. The MIT Press/Bradford Books.

[SAHO94] M. Sahota. Action selection for robots in dynamic environments through inter-behaviour bidding. In [SAB94].

[SAIT94] F. Saito and T. Fukada. Two-link robot brachiation with connectionist Q-learning. In [SAB94].

[SAUN94] G. Saunders, J. Kolen, and J. Pollack. The importance of leaky levels for behavior-based AI. In [SAB94].

[SCHMA92] N. A. Schmajuk and H. T. Blair. Dynamics of spatial navigation: An adaptive neural network. In [SAB92].

[SCHMA94] N. A. Schmajuk. Behavioral dynamics of escape and avoidance: A neural network approach. In [SAB94].

[SCHMI90] J. Schmidhuber. A possibility for implementing curiosity and boredom in model-building neural controllers. In [SAB90].

[SCHMI92] J. Schmidhuber and R. Wahnsiedler. Planning simple trajectories using neural subgoal generators. In [SAB92].

[SCHN90] U. Schnepf. Robot ethology: A proposal for the research into intelligent autonomous systems. In [SAB90].

[SCUT94] T. Scutt. The five neuron trick: Using classic conditioning to learn how to seek light. In [SAB94].

[SHAF94] S. Shafir and J. Roughgarden. The effect of memory length on the foraging behavior of a lizard. In [SAB94].

[SLOT94] J. J. Slotine. Stability in adaptation and learning. In [SAB94].

[SMAR94] W. D. Smart and J. Hallam. Location recognition in rats and robots. In [SAB94].

[STAD92] J. E. R. Staddon. A note on rate-sensitive habituation. In [SAB92].

[STEE90] L. Steels. Towards a theory of emergent functionality. In [SAB90].

[STEE94] L. Steels. A case study in the behavior-oriented design of autonomous agents. In [SAB94].

[SUTT90] R. S. Sutton. Reinforcement learning architectures for animats. In [SAB90].

[TENE92] J. Tenenberg, J. Karlsson, and S. Whitehead. Learning via task decomposition . In [SAB92].

[THER90] G. Theraulaz, S. Goss, J. Gervet, and J.-L. Deneubourg. Task differentation in Polistes wasp colonies: A model for self-organizing groups of robots. In [SAB90].

[TOAT90] F.Toates and P. Jensen. Ethological and psychological models of motivation – Towards a synthesis. In [SAB90].

[TOAT94] F. Toates. What is cognitive and what is not cognitive . In [SAB94].

[TODD90] P. M. Todd and G. F. Miller. Exploring adaptive agency II: Simulating the evolution of associative learning. In [SAB90].

[TODD92] P. M. Todd and S. W. Wilson. Environment structure and adaptive behavior from the ground up. In [SAB92].

[TODD94] P. M. Todd, S. W. Wilson, A. B. Somayaji, and H. Yanco. The blind breeding the blind: Adaptive behavior without looking. In [SAB94].

[TYRR90] T. Tyrrell and J. E. W. Mayhew. Computer simulation of an animal environment. In [SAB90].

[TYRR92] T. Tyrrell. The use of hierarchies for action selection. In [SAB92].

[UTTA92] W. R. Uttal, T. Shepherd, S. Dayanand, and R. Lovell. An integrated computational model of a perceptual-motor system. In [SAB92].

[VENT94] G. Venturini. Adaptation in dynamic environments through a minimal probability of exploration. In [SAB94].

[VERS92] P. F. M. J. Verschure and R. Pfeifer. Categorization, representations, and the dynamics of system-environment interaction: A case study in autonomous systems. In [SAB92].

[VOGE92] T. U. Vogel. Learning biped robot obstacle crossing. In [SAB92].

[VONK90] M. Vonk, F. Putters, and B.-J. Velthuis. The causal analysis of an adaptive system: Sex-ratio decisions as observed in a parasitic wasp and simulated by a network model. In [SAB90].

[WALT90] D. L. Waltz. Eight principles for building an intelligent robot. In [SAB90].

[WANG90] DeLiang Wang and M. A. Arbib. Hierarchical dishabituation of visual discrimination in toads. In [SAB90].

[WEBB94] B. Webb. Robotic experiments in cricket phonotaxis. In [SAB94].

[WEIS92] G. Weiss. Action selection and learning in multi-agent environments. In [SAB92].

[WERN92] G. M. Werner and M. G. Dyer. Evolution of herding behavior in artificial animals. In [SAB92].

[WERN94] G. M. Werner. Using second order neural connections for motivation of behavioral choices. In [SAB94].

[WILS90] S. W. Wilson. The animat path to AI. In [SAB90].

[YAMA94] B. Yamauchi and R. Beer. Integrating reactive, sequential, and learning behavior using dynamical neural networks. In [SAB94].

[YANC92] H. Yanco and L. A. Stein. An adaptive communication protocol for cooperating mobile robots. In [SAB92].

[YEAP90] W. K. Yeap and C. C. Handley. Four important issues in cognitive mapping. In [SAB90].