From SAB94 to SAB2000: What's New, Animat?

Agnès Guillot **Jean-Arcady Meyer**
AnimatLab-LIP6
8, rue du capitaine Scott
75015 Paris, France
{agnes.guillot, jean-arcady.meyer}@lip6.fr

Abstract

This paper is complementary to a previous review of significant research on adaptive behavior in animats. It summarizes the current state-of-the art and outlines directions for possible progress.

1. Introduction

In the proceedings of SAB94, we published a review of significant research on adaptive behavior in animats since the first SAB conference, held in 1990 (MEYE94). This review summarized the state-of-the art, insofar as the proceedings of three dedicated conferences could help delineate it. Now that three other SAB conferences have been held, we considered that it would be useful to update that earlier review, in order to assess the corresponding progress, to infer the directions in which interesting developments are likely to be expected, and to stress needs for specific additional research efforts.

As in the preceding review, this one makes reference only to SAB conference proceedings (SAB96, SAB98, SAB00), on the premise that this perspective, although voluntarily limited, does nevertheless afford a comprehensive view of the field of adaptive behavior in animats. It will first classify the main results that have been obtained since 1994 according to the general organization of the current book, an organization that seems to have almost converged in the course of six conferences. Then it will outline possible directions for future work.

2. Perception and motor control

Since SAB94, new sensory modalities have been incorporated into models of sensori-motor coordination and have been occasionally used for actual robot control. For instance, biomimetic models for odor-guided movements (BELA96, MALA96) have been described, some of which have been implemented on a Lego robot that detects alcohol (SHAR98), or on a robot-lobster that may use the temporal pattern of odors produced by turbulent dispersal processes to locate a distal food source (GRAS96). Likewise, a biomimetic model for phonotaxis has been described and implemented on a cricket-robot (WEBB96), while a biomimetic model of sound diffraction and reflections in the human concha has been applied to bat pinna design for echolocating animals (CARM00). In CHAP00, a biologically-inspired wind sensor is mounted on a Khepera robot to perform a reactive maze solving task.

Vision has been extensively studied in this perspective and many research efforts have dealt with visio-motor coordination. For instance, a biomimetic model of depth perception using peering in insects has been implemented on a Khepera robot (LEWI98). Visual looming has been implemented on a Pioneer robot as a robust and inexpensive range sensor to be used as a complement to a sonar (SAHI98; see also NEVE96 and DAHM98). A system for maze navigation using optical flow has been demonstrated on a simulated robot (DUCH96). A simple vision system has been used to modulate the type of gait and the direction of motion produced by a locomotor circuit in a simulated salamander, thereby conferring the capacity to track a randomly moving target both in water and on dry land (IJSP00). Likewise, specific contributions have dealt with course stabilization and fixation behavior in flies (HUBE98), as well as with adaptive image stabilization, visually-guided pointing and orientation behavior in Cog and other robots with moving eyes (FERR96, MARJ96, PANE00).

Motor control for motion has also been extensively studied. For instance, Yokoi et al. designed a so-called Morpho-Functional Machine, i.e., a deformable robot that exhibits an amoeba-like motion (YOKO98). Likewise, Kawai and Hara describe a machine that changes its morphology in order to carry an object toward a goal while passing through a narrow corridor (KAWA98). Walking controllers for four-legged robots are described in ILG98 and ITO98. Walking controllers for simulated six-legged insects are described in CRUS96. More generally, a mathematical framework that helps design networks of neural oscillators capable of controlling the rhythmic motion of an animat is given in ARSE00.

In ZIEM96 and ZIEM00, several recurrent connectionist control architectures likely to make possible efficient coupling between sensory inputs and motor outputs are described. The latter work introduces the distinction between synchronically and diachronically structured control mechanisms, and describes how an animat can actively and selectively decide when to use feedback to revise its sensorimotor mapping. According to the author, this mechanism allows the animat to flexibly assign varying meanings to environmental stimuli.

Finally, two specific aspects of the perception-action coordination problem are dealt with in SCHE96 and WILS00. Scheier and Lambrinos describe how categorization must be treated as a sensori-motor coordination problem, not as a perceptual one. Their system learns to recognize an object with visual and haptic modalities and to discriminate it from other objects (SCHE96). Wilson and Neal, using a model of interactions between a shepherd, his dog and a sheep, study how the behavior repertoire of the dog robot impacts on the number of interactions required from the shepherd to control the sheep (WILS00).

3. Action selection and behavioral sequences

Several research efforts contributed to delineating the methodological issues underlying the action-selection problem and the design of motivational systems. For instance, Spier and McFarland use a biologically-inspired motivational system to demonstrate that tradeoff between opportunities in response to the environment is often sufficient for generating behavioral sequences that an external observer might attribute to an underlying planning process (SPIE96). Aube and Senteni suggest that emotions are motivations that ensure management and regulation of ressources (AUBE96). Steinhage and Bergener demonstrate how action selection can be implemented in a dynamical system as the result of nonlinear phase transitions (STEI98). Lastly, Seth describes how action selection and selective attention can be exhibited by a simple animat with direct sensori-motor links, thus challenging the concepts of action, attention, and selection (SETH98).

Various architectures for action selection have been compared within the context of reinforcement learning (HUMP96) or from the point of view of hierarchical versus parallel organizations (BRYS00). In GONZ00, a biomimetic basal ganglia model of action selection has been embedded within the control architecture of a Khepera robot, and shown to exhibit nice properties of clean switching, lack of distortion and persistence. In particular, interesting similarities to what is observed on animals have been obtained through the effects on a robot model of varying simulated dopamine levels. In CHAO00, Chao, Panangadan and Dyer describe a connectionist architecture that enables animats to navigate efficiently and learn to build specified structures within an artificial environment. This approach calls upon an external teacher to learn an action-selection architecture that mediates between reactive and planning behaviors, a problem also tackled in REVE98.

Witkowski (WITK00) describes the role extinction mechanisms play in the context of action selection. Such extinction mechanisms contribute to the protection of the animat against the potentially fatal consequences of unattainable high-priority goal-driven activities. His Dynamic Expectancy Model is one of the contemporary learning action selection models that are based on explicit use of prediction to drive the learning process (TANI98, STOL00).

4. Internal world models for navigation

Simple internal world models are elaborated by animats that are able to categorize their environment when moving through it. This is exemplified by the work of Marsland et al. (MARS00) in which a novelty filter using a model of habituation allows a robot operating in an unstructured environment to produce a self-organized model of its surroundings, and to detect deviations from the learned model. Likewise, various systems for self-categorization of sensori-motor patterns can be found in BERT98, SCHE98, TANI98 and LINA00. In the latter case, the technique facilitates the understanding of the concepts abstracted from the animat's sensori-motor flow, and can be used for automatic map-building.

Another extremely simple internal model of the environment is described in PIAG00, where the borderline between behavior-based and representation-based navigation is investigated. This approach calls upon a minimal internal representation to solve local navigation problems induced by local minima in artificial potential fields. Likewise, in CORB96, an analogous potential field approach is described, where a biomimetic model of detour behavior in frogs calls upon generalized schema-based learning. An architecture that combines fuzzy and reactive techniques for obstacle-avoidance is described in GHAN96, where two coupled mobile robots have to move in an environment with obstacles.

Other internal models have been used in several biomimetic approaches to animal homing behavior. In DICK96, alternative modeling approaches to how insects learn about the sun's course are described. In NEHM00, a numerical simulation of Kramer's "Map and Compass" model of long-range pigeon navigation is performed. This model postulates that pigeons use naturally occurring gradients to determine the course to the loft, and compass senses (sun and magnetic) to establish and maintain this direction. Likewise, to simulate homing behavior in desert ants, Moller et al. implemented on a mobile robot a path-integration system using a polarized light compass, in conjunction with a visual piloting system (MOLL98). In KIM00, a circular neuron cell structure, in which each neuron accumulates distance traveled in a particular direction, is suggested as a suitable computational structure for finding a proper homing vector.

The way rodents encode spatial representations of their environment has been exploited in several landmark-based navigation models. This is shown for instance in GAUS98 and in TRUL98. In ARLE00, two biomimetic models of the operation of head-direction cells and place-cells are combined and implemented in a Khepera robot for navigation. In FILL00, such a biomimetic approach is combined with a traditional POMDP (Partially Observable Markovian Decision Process) model that implements an active perception mechanism for map learning and reliable localization in a simulated robot. Another engineering model of landmark-based navigation is given in OWEN98, whereas DONN96 describes how an animat that doesn't use vision to categorize landmarks, but calls on proprioception only, is nevertheless capable of building a cognitive map of its environment and of using it to accurately position itself.

These navigation models could benefit from the work of Balkenius and Morén, who demonstrate that a stable context representation can be learned from a dynamic sequence of attentional shifts between various stimuli in the environment. Such a system can be used for novelty detection and, more specifically, can be used in models where place-cell firing must be associated with specific landmarks (BALK00). Likewise, interesting suggestions are to be found in TOOM98, where a biomimetic model of landmark learning in gerbils is used to demonstrate that complex spatial navigation behavior does not need to be predicated on complex and navigation specific computations. Other interesting suggestions are also to be found in WIER98, where a method for collecting useful experiences through exploration in stochastic environments is described (see also WILS96). Finally, an hybrid model that learns continuously from ongoing experience without preconstructed data-sets and that learns both procedural knowledge - through Q-learning - and declarative knowledge - through propositional rules - is described in SUN98. This model is implemented on a simulated robot required to navigate towards a target through a minefield.

5. Learning

Conditioning is a variety of implicit learning in animals that improves their perceptual or motor skills by repetition, without calling on awareness or higher cognitive processes. Classical conditioning allows an animal to recognize cues for biologically significant events, while operant conditioning allows to change its voluntary behavior according to the outcome of its actions.

A wide variety of models implement such abilities in animats. For instance, models of classical conditioning are described in SALO98, BALK98 and HALL00. A model of operant conditioning has been used by Touretzky and Saksida (TOUR96) to implement chaining capacities in a robot, according to which complex behavioral routines are built up from smaller action segments, the response of the first one being the stimulus for the next. Stolzmann et al. (STOL00) describe another model of operant conditioning based upon Hoffmann's learning theory of anticipatory behavioral control. This model reproduces some of the experimental results that have been obtained on rats in a Skinner box: It is notably capable of distinguishing between different reaction-effect relations and of relating them to different stimuli. Lastly, models that combine classical and operant conditioning are to be found in BALK96, GAUD96 and BLUM96. The latter case demonstrates how a virtual dog can acquire new behaviors like taste aversion, outcome devaluation, habit formation, and superstitious behavior.

In the field of reinforcement learning, improvements to the traditional Q-learning algorithm have been demonstrated in ARAU96, DIGN96, DIGN98, MINA98, MORE98, NAKA98, IIJI00 and MOTO00. Several research efforts have been devoted to non-Markovian problems: ARAU96, MCCA96, SUNa00, SUNb00. A unified approach to perceptual aliasing is presented by Lanzi (LANZ00) who introduces the so-called "on the payoffs" aliasing problem and suggests that, to achieve proper performance, an animat does not need to learn the whole mapping scheme from perception-action pairs to payoffs. To this end, non-tabular reinforcement learning schemes (e.g., LCS) may be more effective than tabular techniques inspired from Dynamic Programming (e.g., Q-learning). An alternative to the usual state-action evaluation approach to reinforcement learning is suggested by Porta and Celaya in the case of categorizable environments, i.e., environments where the effects of a given action can be foreseen through a limited number of the animat's sensors. Here, the problem is to determine the relevance of the sensors with respect to each action and to the corresponding reward. The corresponding paper (PORT00) describes an application to step coordination in a simulated 6-legged robot walking either in flat and rough terrain.

Learning by being taught or by imitation has also received special emphasis, notably in RAO96, ANDR00 and COLL00. In CRAB00, it is shown how observation and imitation of a teacher can be used by a learning agent to satisfy a sequence of goals. Learning of goal sequences differs from usual action-learning in that the order of individual actions is left open, but the order of the goals that these actions achieve is fixed. This approach is applied to animats that can perform construction tasks while maintaining their survival in a complex and hazardous environment.

Finally, several articles investigate how emotions might be involved in learning, including WRIG96, GADA98 and MORE00.

6. Evolution

While Lerena and Courant explored the relationships between sexual selection and natural selection in LERE98 and LERE00, mechanisms for artificial selection have been studied and put to work in a large number of applications. It thus has been possible to evolve animats that play hockey (BLAI98), that explore their environment (SMIT96), forage for food (BENN96) or collect garbage (CALA98), that visually track targets (JAKO98, KORT00) or discriminate landmarks (NOLF00), or that are capable of switching from swimming to walking (IJSP98). Likewise, the feasibility of evolving both the morphology and the control of animats has been investigated in VENT96, KIKU98 and BONG00.

To demonstrate the capacity of evolutionary approaches to generate more than mere reflexive behaviors, Beer and colleagues evolved a series of neural controllers that exhibit "minimally cognitive behaviors", i.e., the simplest behaviors that raise issues of genuine cognitive interest. In BEER96, animats are evolved for orientation and reaching objects, as well as for discrimination between objects (see also BIRO98). In SLOC00, animats are evolved that can judge the passability of openings relative to their own body size, that can distinguish between visible parts of themselves and other objects in their environment, that can predict and remember the future location of objects in order to catch them blind, and that can switch their attention between multiple distal objects. Very often, such functionalities rely on mechanisms for active scanning and sensory-motor coordination.

In nature, evolution concurs with development and learning in animal adaptation, and one main objective of animat research is to understand the corresponding synergies. In this perspective, relationships between evolution and learning have been investigated in FLOR96 and MAYL96. In DIPA00, rules of plastic change at synaptic level within neural controllers are genetically encoded. Robots are evolved to perform phototaxis and to recover after the inversion of their visual field and other disruptions. Likewise, interactions between development and evolution have been investigated in EGGE96 and DELL96.

Mechanisms of co-evolution have been extensively studied through a variety of pursuit-evasion games, as in CLIF96, WAHD98, FICI98, FLOR98. The latter two papers describe Red-Queen effects that prevent regular fitness increase, according to which co-evolution is not automatically better than simple evolution. In FUNE98 and FUNE00 a statistical method is proposed that serves to evaluate the fitness of each individual in such co-evolutionary experiments.

Another methodological contribution is that of Jakobi, whose "minimal simulation" approach is designed to help transfer to the real world controllers or morphologies that have been evolved in simulation (JAKO98).

It may likewise be instructive to refer to ZAER96 for a case-study where artificial evolution failed to produce controllers analogous to those hand-crafted by humans, for reasons that the authors think are rooted in the difficulties of formulating an effective evaluation function.

7. Collective behaviors

Several collectivities of animats have been involved in research efforts that dealt with foraging or related behaviors. In this context, the mandatory tradeoff between exploration and exploitation is investigated in BONA96, while DEB096 and SETH00 deal with optimal foraging theory. In WERG96, it is shown how robots with local sensing and action form a system that dynamically and globally adapts to environmental changes. This robotic system, inspired from the natural phenomenon of ant pheromone trail formation, encodes information in its physical environment in order to reduce sensing, actuation, and computational requirements for gathering metal pucks. In FONT96, a territorial principle that implements a division of labor into exclusive spatial areas is used for the same task. In MELH98 a collective sorting and segregation task is performed by a system of simple homogeneous autonomous robots which have no capacities for spatial orientation or memory. Finally, the role of social development in the evolution of cooperation has been explored by Di Paolo (DIPA98), who suggests that the role played by natural selection be reconsidered as the main explanatory factor in the determinants of social behaviors.

The way signaling fighting ability can help solve conflicts has been explored in several contexts. For instance, Noble describes an evolutionary simulation that challenges Enquist's assumption that weak animals will signal their fighting ability honestly because they have so much to lose by bluffing (NOBL00). Likewise, Vaughan et al. implement stylized fighting behavior in a collectivity of robots to solve spatial interference problems. In case of space conflict between two robots, these robots compare their apparent levels of aggression and the more aggressive robot takes precedence over the less aggressive one (VAUG00). A related work is that of Hemelrijk (HEME96), who studies dominance interactions, spatial dynamics and emergent reciprocity in a virtual world. Another related work is that of Noble (NOBL98), who describes intention movements and the evolution of signaling in animal contests.

Other varieties of communication have been studied in the SAB context (NOBL96, SAUN96). In particular, Reznikova and Ryabko apply Information Theory to the study of communication in ants and demonstrate that, in the communication system of these insects, the frequency of use of a message correlates with its length. The authors also demonstrate that the numerical competence of ants calls on adding or subtracting small numbers in preference to large ones (REZN00).

More sophisticated varieties of communication are explored in the "language games" initiated by Steels and his colleagues. STEE96 focusses on emergent adaptive lexicons, while STEE98 deals with the structural coupling of cognitive memories. Moukas and Hayes describe how a movement-based language, like that of bees, can be implemented in Lego robots. In this work, a robot learns the language elements, what they mean, and how to reproduce them, by observing a teacher robot that performs a "dance" indicating the presence of a particular type of "food" at a particular distance and bearing (MOUK96).

Finally, several research efforts address multi-agent pursuit games. This is the case with ONO96, ZHAO96 and ARAI00.

8. Applied adaptive behavior

The concepts and techniques in favor in the SAB community have been used in diverse applications, ranging from market trading (CLIF98), to traffic system modeling (MORI98), environmental monitoring (COST98), and software agent learning (RAMA98). In DAUT00, Dautenhahn and Werry describe how mobile robots can play a therapeutic role in the rehabilitation of children with autism. In SKLA00, Sklar and Pollack describe an evolutionary algorithm that is used to select content for keyboarding educational games in a web-based learning community.

9. Prospects

As emphasized by Clark and Miller (CLARK98), recent research into animats suggests that a great deal of their adaptive capacities are grounded, not in the systematic activity of internal representations, but in complex interactions involving neural, bodily and environmental factors. In his quest for understanding intelligence, Pfeifer, for instance, has long advocated the design of complete and embodied animats (PFEI96, SCHE98). In PFEI00, he introduces the concept of "ecological balance" which means that, given a particular task environment, there must be a harmonious relationship between an animat's morphology, materials and control. In HARA00, he elaborates on the relation among morphology, material and control in morpho-functional machines. The role of embodiment is also stressed by Krichmar et al. (KRIC00; see also ALMA98), who demonstrate the role of early sensory experience for the development of perceptual categories in Darwin VI. This process appears to be highly dynamic and to depend strongly upon the actual sequence and content of sensory experience, and upon individual histories of stimulus encounters. Therefore, because of its embodiment, a robot never experiences a stimulus in exactly the same way.

However, Clark and Miller challenge the above-mentioned pessimistic view about the explanatory role of representations in the determinants of adaptive behavior (CLARK98). Their work, together with that of researchers who explicitly dealt with cognitive processes (see, e.g., BEER96, SPIE98, SLOC00, STOL00), raises the questions of how far animat designers will be able to raise the cognitive capacities of their creatures, and which role internal representations will play in the corresponding achievements. No doubt numerous empirical answers will be brought to bear on these issues in the near future.

Another contribution that gives food for thought is that of Keijzer (KEIJ98). According to this author, the difference is considerable between movement and behavior: Animats and robots move, while animals behave. Such a distinction is related to that between proximal and distal stimuli. The former directly impinge on the animat's body, but this stimulation in itself tends to be neutral as far as adaptation is concerned. The latter are provided by the adaptively relevant elements in an environment that are usually at some distance from the animat, and provide the possible nourishment, threats or mates on which survival and reproduction depend. In this context, the behavior of animals, but not that of robots, consists of a process of self-organization according to which variable proximal (fast and short-term) sensory-motor coupling maintains stable distal (slower and longer-term) perception-action coupling. Whatever the case, many lessons are still to be drawn from comparisons of natural and artificial adaptive behaviors. In particular, an animat certainly has a long way to go before it might be taken for an animal, in any sort of Turing-like test that may be invented in future.

Several applications of dynamical systems theory to animat design were mentioned above (e.g., BEER96, NEVE96, JAEG98, RYLA98, MENZ00), whose authors seem convinced that this approach will scale with more complex behaviors and survival problems. Likewise, although Aubin's viability theory has just been introduced to the SAB community (AUBI00), it appears in our opinion to be too closely related to animat concerns (see, e.g., MEYE94) for not soon start guiding the intuition of animat designers. There is therefore definitely some hope that such efforts to bring mathematics and their deductive power into the empiricism of animat research will sooner or later generate the sort of stability theorem or convergence proof that the field definitely lacks at present. Advances in these directions will undoubtedly foster numerous practical applications, over those already described above.

Several methodological contributions to the field of adaptive behavior have already been listed (e.g., PFEI96, JAKO98, FUNE00). Another such contribution worth mentioning is that of Bakker and de Jong (BAKK00), who provide a means for counting the number of states in an animat's behavior and for counting the number of states required to perform a particular task in an environment. Such state counts provide a measure of the complexity of agents and environments. Likewise, in a reinforcement learning framework, Wilson (WILS96) reviews ten strategies for the autonomous control of the explore-versus-exploit decision, and pleas for a better understanding of how a system can tell how well it is doing. Lastly, in the work of Fleming et al. (FLEM00), the brain of a lamprey is used to control a Khepera robot. The observed artificial behaviors help extract information about information processing in the neural tissue connected to the robot. Such exciting work probably paves the way for the numerous "hybrid" approaches - where artificial and biological materials will be merged - the fast expansion of which can be readily foreseen, hopefully to the best avail. Be that as it may, methodological contributions like these reinforce the experimental bases of animat research, and hopefully many more will come to complete them.

Finally, in our previous review of animat research (MEYE94), we stressed the need for comparisons that would allow understanding what architectures and working principles can allow an animat to solve what kind of problem in what kind of environment. Although several papers have been focussed on such comparisons (e.g., DICK96, HUMP96, MAYL96, SMIT96, BALK98, BLAI98, CALA98, HEME98, WIER98), we are obliged to observe that the number of architectures and working principles has grown much faster than the number of comparisons. In this sense, the present situation is worse than it was three SAB conferences ago. Nevertheless, it can hardly be concluded that the field of animat research is not healthy and productive.

10. Conclusion

Since SAB94, the field of animat research has broadened and deepened. New machines, new mechanisms, new methods, and new concepts have been described in this paper. Beyond mere reflexes, significant steps have been taken towards implementing higher cognitive mechanisms in the control architectures of animats. In the near future, a number of the theoretical and methodological advances outlined herein should provide the sort of generalizations that we eagerly wished for in our previous review.

References

[ALMA98] N. Almassy and O. Sporns. Role of Behavior in the Development of Complex Neuronal Properties. In [SAB98].

[ANDR00] P. Andry, S. Moga, P. Gaussier, A. Revel and J. Nadel. Imitation: learning and communication. In [SAB00].

[ARAI00] S. Arai and K. Sycara. Effective Learning Approach for Planning and Scheduling in Multi-Agent Domain. In [SAB00].

[ARAU96] E. G. Araujo and R. A. Grupen. Learning Control Composition in a Complex Environment. In [SAB96].

[ARLE00]. A. Arleo and W. Gerstner. Modeling Rodent Head-direction Cells and Place Cells for Spatial Learning in Bio-mimetic Robotics. In [SAB00].

[ARSE00] A. M. Arsenio. Neural Oscillator Networks for Rhythmic Control of Animats. In [SAB00].

[AUBE96] M. Aube and A. Senteni. What Are Emotions For? Commitments Management and Regulation Within Animals/Animats Encounters. In [SAB96].

[AUBI00] J.P. Aubin. Elements of Viability Theory for Animat Design. In [SAB00].

[BAKK00] B. Bakker and M. de Jong. The Epsilon State Count. In [SAB00].

[BALK96] C. Balkenius. Generalization in Instrumental Learning. In [SAB96].

[BALK98] C. Balkenius and J. Moren. Computational Models of Classical Conditioning: A Comparative Study. In [SAB98].

[BALK00] C. Balkenius and J. Moren A Computational Model of Context Processing. In [SAB00].

[BEER96] R. D. Beer. Toward the Evolution of Dynamical Neural Networks for Minimally Cognitive Behavior. In [SAB96].

[BELA96] J. H. Belanger and M. A. Willis. Centrally-Generated and Reflexive Control Strategies in the Adaptive Behavior of Real and Simulated Animals. In [SAB96].

[BENN96] F. H. Bennett III. Emergence of a Multi-Agent Architecture and New Tactics for the Ant Colony Food Foraging Problem Using Genetic Programming. In [SAB96].

[BERT98] L. Berthouze, Y. Shigematsu and Y. Kuniyoshi. Dynamic Categorization of Explorative Behaviors for Emergence of Stable Sensimotor Configurations. In [SAB98].

[BIRO98] Z. Biro and T. Ziemke. Evolution of Visually-Guided Approach Behaviour in Recurrent Artificial Neural Network Robot Controllers. In [SAB98].

[BLAI98] A. D. Blair and E. Sklar. The Evolution of Subtle Maneuvers in Situated Hockey. In [SAB98].

[BLUM96] B. M. Blumberg, P. M. Todd and P. Maes. No Bad Dogs: Ethological Lessons for Learning in Hamsterdam. In [SAB96].

[BONA96] E. Bonabeau and F. Cogne. Oscillation-Enhanced Adaptability in the Vicinity of a Bifurcation: The Example of Foraging in Ants. In [SAB96].

[BONG00] J. C. Bongard and C. Paul. Investigating Morphological Symmetry and Locomotive Efficiency using Virtual Embodied Evolution. In [SAB00].

[BRYS00] J. Bryson. Hierarchy and Sequence vs. Full Parallelism in Action Selection. In [SAB00].

[CALA98] R. Calabretta, S. Nolfi, D. Parisi and G. P. Wagner. Emergence of Functional Modularity in Robots. In [SAB98].

[CARM00] J. M. Carmena, D. Kim and J. C.T. Hallam. Designing artificial ears for animat echolocation. In [SAB00].

[CHAO00] G. Chao, A. Panangadan and Michael G. Dyer. Learning to Integrate Reactive and Planning Behaviors for Construction. In [SAB00].

[CHAP00] T. Chapman, A. Hayes and M. Tilden. Reactive Maze Solving with a Biologically-Inspired Wind Sensor. In [SAB00].

[CLAR98] A. Clark and M. Wheeler. Bringing Representation Back to Life. In [SAB98].

[CLIF96] D. Cliff and G. F. Miller. Co-evolution of Pursuit and Evasion II: Simulation Methods and Results. In [SAB96]

[CLIF98] D. Cliff and J. Bruten. Market Trading Interactions as Collective Adaptive Behavior. In [SAB98].

[COLL00] D. Collins and G. Wyeth. Utilising a Cerebellar Model for Mobile Robot Control in a Delayed Sensory Environment. In [SAB00].

[CORB96] F. J. Corbacho and M. A. Arbib. Learning to Detour & Schema-based Learning. In [SAB96].

[CRAB00] F. L. Crabbe and M. G. Dyer. Observation and Imitation: Goal Sequence Learning in Neurally Controlled Construction Animats: VI-MAXSON. In [SAB00].

[COST98] J. Costa Seco, C. Pinto-Ferreira and L. Correia. A Society of Agents in Environmental Monitoring. In [SAB98].

[CRUS96] H. Cruse, C. Bartling, J. Dean, T. Kindermann, J. Schmitz, M. Schumm and H. Wagner. Coordination in a Six-Legged Walking System: Simple Solutions to Complex Problems by Exploitation of Physical Properties. In [SAB96].

[DAHM98] P. Dahm and C. Bruckhoff. Autonomous Decision Making in Local Navigation. In [SAB98].

[DAUT00] K. Dautenhahn and I. Werry. Issues of Robot-Human Interaction Dynamics in the Rehabilitation of Children with Autism. In [SAB00].

[DEBO96] P. de Bourcier. Using A-Life to Study Bee Life: The Economics of Central Place Foraging. In [SAB96].

[DELL96] F. Dellaert and R. D. Beer. A Developmental Model for the Evolution of Complete Autonomous Agents. In [SAB96].

[DICK96] J. Dickinson and F. Dyer. How Insects Learn about the Sun's Course: Alternative Modeling Approaches. In [SAB96].

[DIGN96] B. L. Digney. Emergent Hierarchical Control Structures: Learning Reactive/Hierarchical Relationships in Reinforcement Environment. In [SAB96].

[DIGN98] B. L. Digney. Learning Hierarchical Control Structures for Multiple Tasks and Changing Environments. In [SAB98].

[DIPA98] E. A. Di Paolo. Assessing the Role of Social Development in the Evolution of Cooperation. In [SAB98].

[DIPA00] E. A. Di Paolo. Homeostatic adaptation to inversion of the visual field and other sensorimotor disruptions. In [SAB00].

[DONN96] J.Y. Donnart and J.A. Meyer. Spatial Exploration, Map Learning, and Self-Positioning with MonaLysa. In [SAB96].

[DUCH96] A. P. Duchon. Maze Navigation Using Optical Flow. In [SAB96].

[EGGE96] P. Eggenberger. Cell Interactions as a Control Tool of Developmental Processes for Evolutionary Robotics. In [SAB96].

[FERR96] C. Ferrell Orientation Behavior Using Registered Topographic Maps. In [SAB96].

[FICI98] S. G. Ficici and J. B. Pollack. Coevolving Communicative Behavior in a Linear Pursuer-Evader Game. In [SAB98].

[FILL00] D. Filliat and J.A. Meyer. Active Perception and Map Learning for Robot Navigation. In [SAB00].

[FLEM00]. In [SAB00].

[FLOR96] D. Floreano and F. Mondada. Evolution of Plastic Neurocontrollers for Situated Agents. In [SAB96].

[FLOR98] D. Floreano, S. Nolfi and F. Mondada. Competitive Co-Evolutionary Robotics: From Theory to Practice. In [SAB98].

[FONT96] M. Schneider Fontan and M. J. Mataric. A Study of Territoriality: The Role of Critical Mass in Adaptive Task Division. In [SAB96].

[FUNE98] P. Funes, E. Sklar, H. Juille and J. Pollack. Animal-Animat Coevolution: Using the Animal Population as Fitness Function. In [SAB98].

[FUNE00] P. Funes and J. B. Pollack. Measuring Progress in Coevolutionary Competition. In [SAB00].

[GADA98] S. C. Gadanho and J. Hallam. Emotion-Driven Learning for Animat Control. In [SAB98]. [GAUD96] P. Gaudiano, E. Zalama, C. Chang and J. L. Coronado. A Model of Operant Conditioning for Adaptive Obstacle Avoidance. In [SAB96].

[GAUS98] P. Gaussier, C. Joulain and J.P. Banquet. Motivated Animat Navigation: A Visually Guided Approach. In [SAB98].

[GHAN96] R. Ghanea-Hercock and D. P Barnes. An Evolved Fuzzy Reactive Control System for Cooperating Autonomous Robots. In [SAB96].

[GONZ00] F. Montes Gonzalez, T. J. Prescott, K. Gurney, M. Humphries and P. Redgrave An embodied model of action selection mechanisms in the vertebrate brain. In [SAB00].

[GRAS96] F. Grasso, T. Consi, D. Mountain and J. Atema. Locating Odor Sources in Turbulence with a Lobster Inspired Robot. In [SAB96].

[HALL00] B. Hallam. Simulating Classical Conditioning using a Neuro-Connector Net. In [SAB00].

[HARA00] F. Hara and R. Pfeifer. On the relation among morphology, material and control in morpho-functional machines. In [SAB00].

[HEME96] C. K. Hemelrijk. Dominance Interactions, Spatial Dynamics and Emergent Reciprocity in a Virtual World In [SAB96].

[HUBE98] S. A. Huber and H. H. Bulthoff. Simulation and Robot Implementation of Visual Orientation Behaviors of Flies. In [SAB98].

[HUMP96] M. Humphrys. Action Selection Methods using Reinforcement Learning. In [SAB96].

[IIJI00] D. Iijima, W. Yu, H. Yokoi and Y. Kakazu. Obstacle avoidance for a distributed autonomous swimming robot by interaction-based learning. In [SAB00].

[IJSP98] A. J. Ijspeert, J. Hallam and D. Willshaw. From Lampreys to Salamanders: Evolving Neural Controllers for Swimming and Walking. In [SAB98].

[IJSP00] A. J. Ijspeert and M. Arbib. Visual tracking in simulated salamander locomotion. In [SAB00].

[ILG 98] W. Ilg, K. Berns, H. Jedele, J. Albiez, R. Dillmann, M. Fischer, H. Witte, J. Biltzinger, R. Lehmann and N. Schilling. BISAM: From Small Mammals to a Four-Legged Walking Machine. In [SAB98].

[ITO 98] S. Ito, H. Yuasa, Z.W. Luo, M.Ito and D. Yanagihara. Adaptive Behavior of Quadruped Walking Robot under Periodic Perturbation. In [SAB98].

[JAEG98] H. Jaeger. Multifunctionality: a Fundamental Property of Behavior Mechanisms Based on Dynamical Systems. In [SAB98].

[JAKO98] N. Jakobi. Evolving Motion-Tracking Behaviour for a Panning Camera Head. In [SAB98].

[KAWA98] N. Kawai and F. Hara. Formation of Morphology and Morpho-Function in a Linear-Cluster Robotic System. In [SAB98].

[KEIJ98] F.A. Keijzer. Some Armchair Worries about Wheeled Behavior. In [SAB98].

[KIKU98] K. Kikuchi and F. Hara. Evolutionary Design of Morphology and Intelligence in Robotic System Using Genetic Programming. In [SAB98].

[KIM00] D. Kim and J. C. T. Hallam. Neural Network approach to path integration for homing navigation. In [SAB00].

[KORT00] R. Kortmann, E. Postma and J. van den Herik. Visual resolution evolving to a trade-off curve. In [SAB00].

[KRIC00] J. L. Krichmar, J. A. Snook, G. M. Edelman and O. Sporns. Experience-Dependent Perceptual Categorization in a Behaving Real-World Device. In [SAB00].

[LANZ00] P. L. Lanzi. Adaptive Agents with Reinforcement Learning and Internal Memory. In [SAB00].

[LERE98] P. Lerena and M.Courant. Mate Choice and NK-Landscapes. In [SAB98].

[LERE00] P. Lerena. Sexual Preferences: Dimension and Complexity. In [SAB00].

[LEWI98] M. A. Lewis and M. E. Nelson. Look Before You Leap: Peering Behavior for Depth Perception. In [SAB98].

[LINA00] F. Linaker and L. Niklasson Extraction and inversion of abstract sensory flow representations. In [SAB00].

[MALA96] R. Malaka, S. Schmitz and W. M. Getz. A Self-Organizing Model of the Antennal Lobes. In [SAB96].

[MARJ96] M. Marjanovic, B. Scassellati and M. Williamson. Self-Taught Visually-Guided Pointing for a Humanoid Robot. In [SAB96].

[MARS00] S. Marsland, U. Nehmzow and J. Shapiro. Detecting Novel Features of an Environment Using Habituation. In [SAB00].

[MAYL96] G. Mayley. The Evolutionary Cost of Learning. In [SAB96].

[MCCA96] A. K. McCallum. Learning to Use Selective Attention and Short-Term Memory in Sequential Tasks. In [SAB96].

[MELH98] C. Melhuish, O. Holland and S. Hoddell. Collective Sorting and Segregation in Robots With Minimal Sensing. In [SAB98].

[MENZ00] R. Menzner, A. Steinhage and W. Erlhagen. Generating Interactive Robot Behavior: A Mathematical Approach. In [SAB00].

[MEYE94] J.A. Meyer and A. Guillot. From SAB90 to SAB94: Four Years of Animat Research. In [SAB94].

[MINA98] T. Minato and M. Asada. Skill Acquisition and Self-Improvement for Environmental Change Adaptation of Mobile Robot. In [SAB98].

[MOLL98] R. Moller, D. Lambrinos, R. Pfeifer, T. Labhart and R. Wehner. Modeling Ant Navigation with an Autonomous Agent. In [SAB98].

[MORE98] J. Moren. Dynamic Action Sequences in Reinforcement Learning. In [SAB98].

[MORE00] J. Moren and C. Balkenius. A Computational Model of Emotional Learning in the Amygdala. In [SAB00].

[MORI98] D. E. Moriarty, S. Handley and P. Langley. Learning Distributed Strategies for Traffic Control. In [SAB98].

[MOTO00] K. Motoyama, K. Suzuki, M. Yamamoto and A. Ohuchi. Evolutionary Design of Learning State Space for Small Airship Control. In [SAB00].

[MOUK96] A. Moukas and G. Hayes. Synthetic Robotic Language Acquisition by Observation. In [SAB96].

[NAKA98] T. Nakamura. Self-Organizing Internal Representation for Behavior Acquisition of Vision-Based Mobile Robots. In [SAB98].

[NEHM00] U. Nehmzow and R. Wiltschko. Computer Simulation of Long Range Bird Navigation. In [SAB00].

[NEVE96] H. Neven, A. Steinhage, M. Giese and C. Bruckhoff. Dynamics for Vision-Based Autonomous Mobile Robots. In [SAB96].

[NOBL96] J. Noble and D. Cliff. On Simulating the Evolution of Communication. In [SAB96].

[NOBL98] J. Noble. Tough Guys Don't Dance: Intention Movements and the Evolution of Signalling in Animal Contests. In [SAB98].

[NOBL00] J. Noble. Talk is cheap: Evolved strategies for communication and action in asymmetrical contests. In [SAB00].

[NOLF00] S. Nolfi and D. Marocco. Evolving Visually-Guided Robots able to Discriminate Beween Different Landmarks. In [SAB00].

[OWEN98] C. Owen and U. Nehmzow. Landmark-Based Navigation for a Mobile Robot. In [SAB98].

[ONO96] N. Ono, K. Fukumoto and O. Ikeda. Collective Behavior by Modular Reinforcement-Learning Animats. In [SAB96].

[PANE00] F. Panerai, G. Metta and G. Sandini. Adaptive image stabilization: aneed for vision-based active robotics agents. In [SAB00].

[PFEI96] R. Pfeifer. Building "Fungus Eaters": Design Principles of Autonomous Agents. In [SAB96].

[PFEI00] R. Pfeifer. On the role of morphology and materials in adaptive behavior. In [SAB00].

[PIAG00] M. Piaggio, A. Sgorbissa and R. Zaccaria. Micronavigation. In [SAB00].

[PORT00] J. M. Porta and E. Celaya. Learning in Categorizable Environments. In [SAB00].

[RAMA98] U. Ramamurthy, S. Franklin and A. Negatu. Learning Concepts in Software Agents. In [SAB98].

[RAO 96] R. P. N. Rao and O. Fuentes. Learning Navigational Behaviors Using a Predictive Sparse Distributed Memory. In [SAB96].

[REVE98] A. Revel, P. Gaussier, S. Lepretre and J.P. Banquet. Planification Versus Sensory-Motor Conditioning: What Are the Issues? In [SAB98].

[REZN00] Z. Reznikova and B. Ryabko. Using Information Theory Approach to study the Communication System and Numerical Competence in Ants. In [SAB00].

[RYLA98] M. Rylatt and C. Czarnecki. Beyond Physical Grounding and Naïve Time: Investigations Into Short-Term Memory For Autonomous Agents. In [SAB98].

[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.

[SAB96] P. Maes, M. Mataric, J.A. Meyer, J. Pollack and S.W. Wilson (Eds.). 1996. From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior. The MIT Press/Bradford Books.

[SAB98] R. Pfeifer, B. Blumberg, J.A. Meyer and S.W. Wilson (Eds.). 1998. From Animals to Animats 5: Proceedings of the Fifth International Conference on Simulation of Adaptive Behavior. The MIT Press/Bradford Books.

[SAB00]. J.A. Meyer, A. Berthoz, D. Floreano, H. Roitblat and S.W. Wilson (Eds.). 2000. From Animals to Animats 6: Proceedings of the Sixth International Conference on Simulation of Adaptive Behavior. The MIT Press/Bradford Books.

[SAHI98] E. Sahin and P. Gaudiano. Visual Looming as a Range Sensor for Mobile Robots. In [SAB98].

[SALO98] R. Salomon. Improving the DAC Architecture by Using Proprioceptive Sensors. In [SAB98].

[SAUN96] G. M. Saunders and J. B. Pollack. The Evolution of Communication Schemes over Continuous Channels. In [SAB96].

[SCHE96] C. Scheier and D. Lambrinos. Categorization in a Real-World Agent Using Haptic Exploration and Active Perception. In [SAB96].

[SCHE98] C. Scheier and R. Pfeifer. Exploiting Embodiment for Category Learning. In [SAB98].

[SETH98] A. K. Seth. Evolving Action Selection and Selective Attention Without Actions, Attention, or Selection. In [SAB98].

[SETH00] A. K. Seth. Unorthodox Optimal Foraging Theory. In [SAB00].

[SHAR98] T. Sharpe and B. Webb. Simulated and Situated Models of Chemical Trail Following in Ants. In [SAB98].

[SKLA00] E. Sklar and P. Pollack. An Evolutionary Approach to guiding Students in an Educational Game. In [SAB00].

[SLOC00] A.C. Slocum, D.C. Downey and R.D. Beer. Further Experiments in the Evolution of Minimally Cognitive Behavior: From Perceiving Affordances to Selective Attention. In [SAB00].

[SMIT96] R.E. Smith and H.B. Cribbs III. Cooperative Versus Competitive System Elements in Coevolutionary Systems. In [SAB96].

[SPIE96] E. Spier and D. McFarland. A Finer-Grained Motivational Model of Behaviour Sequencing. In [SAB96].

[SPIE98] E. Spier and D. McFarland. Learning To Do Without Cognition. In [SAB98].

[STEE96] L. Steels. Emergent Adaptive Lexicons. In [SAB96].

[STEE98] L. Steels. Structural Coupling of Cognitive Memories through Adaptive Language Games. In [SAB98].

[STEI98] A. Steinhage and T. Bergener. Dynamical Systems for the Behavioral Organization of an Anthropomorphic Mobile Robot. In [SAB98].

[STOL00] W. Stolzmann, M. V. Butz, J. Hoffmann and D. E. Goldberg. First Cognitive Capabilities in the Anticipatory Classifier System. In [SAB00].

[SUN98] R. Sun and T. Peterson. A Subsymbolic+Symbolic Model for Learning Sequential Navigation. In [SAB98].

[SUNa00] R. Sun and C. Sessions. Multi-Agent Reinforcement Learning with Bidding for Segmenting Action Sequences. In [SAB00].

[SUNb00] R. Sun and T. Peterson. Partitioning in Multi-Agent Reinforcement Learning. In [SAB00].

[TANI98] J. Tani and S. Nolfi. Learning to Perceive the World as Articulated: An Approach for Hierarchical Learning in Sensory-Motor Systems. In [SAB98].

[TOOM98] S.P. Toombs, W.A. Phillips and L.S. Smith. Reinforcement Landmark Learning. In [SAB98].

[TOUR96] D. S. Touretzky and L. M. Saksida. Skinnerbots. In [SAB96].

[TRUL98] O. Trullier and J.A. Meyer. Animat Navigation Using a Cognitive Graph. In [SAB98].

[VAUG00] R. T. Vaughan, K.Stoy, G. S. Sukhatme and M. J. Mataric. Go ahead, make my day: Robot conflict resolution by aggressive competition. In [SAB00].

[VENT96] J. Ventrella. Sexual Swimmers: Emergent Morphology and Locomotion without a Fitness Function. In [SAB96].

[WAHD98] M. Wahde and M. G. Nordahl. Evolution of Protean Behavior in Pursuit-Evasion Contests. In [SAB98].

[WEBB96] B. Webb and J. Hallam. How to Attract Females: Further Robotic Experiments in Cricket Phonotaxis. In [SAB96].

[WERG96] B. B. Werger and M. J. Mataric. Robotic ”Food” Chains: Externalization of State and Program for Minimal-Agent Foraging. In [SAB96].

[WIER98] M. Wiering and J. Schmidhuber. Efficient Model-Based Exploration. In [SAB98].

[WILS96] S. W. Wilson. Explore/Exploit Strategies in Autonomy. In [SAB96].

[WILS00] M. Wilson, M. Neal. Telerobotic Sheepdogs: how useful is autonomous behavior ? In [SAB00].

[WITK00] M. Witkowski. The Role of Behevioral Extinction in Animat Action Selection. In [SAB00].

[WRIG96] I. Wright. Reinforcement Learning and Animat Emotions. In [SAB96].

[YOKO98] H. Yokoi, W. Yu, J. Hakura and Y. Kakazu. Morpho-Functional Machine: Robotics Approach of Amoeba Model Based on Vibrating Potential Method. In [SAB98].

[ZAER96] N. Zaera, D. Cliff and J. Bruten. (Not) Evolving Collective Behaviours in Synthetic Fish. In [SAB96].

[ZHAO96] J. Zhao and J. Schmidhuber. Incremental Self-Improvement for Life-Time Multi-Agent Reinforcement Learning. In [SAB96].

[ZIEM96] T. Ziemke. Towards Adaptive Behavior System Integration. Using Connectionist Infinite State Automata. In [SAB96].

[ZIEM00] T. Ziemke. On ’Parts’ and ’Wholes’ of Adaptive Behavior: Functional Modularity and Diachronic Structure in Recurrent Neural Robot Controllers. In [SAB00].