Animats: computer-simulated animals in behavioral research

J. M. Watts

J Anim Sci 1998. 76:2596-2604.

Animats: Computer-Simulated Animals in Behavioral Research[1,2]

J. M. Watts

Department of Herd Medicine and Theriogenology, Western College of Veterinary Medicine, University of Saskatchewan, Saskatoon, Saskatchewan, Canada S7N 5B4

ABSTRACT: The term *animat* refers to a class of simulated animals. This article is intended as a nontechnical introduction to animat research. Animats can be robots interacting with the real world or computer simulations. In this article, the use of computer-generated animats is emphasized. The scientific use of animats has been pioneered by artificial intelligence and artificial life researchers. Behavior-based artificial intelligence uses animats capable of autonomous and adaptive activity as conceptual tools in the design of usefully intelligent systems. Artificial life proponents view some human artifacts, including informational structures that show adaptive behavior and self-replication, as animats may do, as analogous to biological organisms. Animat simulations may be used for rapid and inexpensive evaluation of new livestock environments or management techniques. The animat approach is a powerful heuristic for understanding the mechanisms that underlie behavior. The simple rules and capabilities of animat models generate emergent and sometimes unpredictable behavior. Adaptive variability in animat behavior may be exploited using artificial neural networks. These have computational properties similar to natural neurons and are capable of learning. Artificial neural networks can control behavior at all levels of an animat's functional organization. Improving the performance of animats often requires genetic programming. Genetic algorithms are computer programs that are capable of self-replication, simulating biological reproduction. Animats may thus evolve over generations. Selective forces may be provided by a human overseer or be part of the simulated environment. Animat techniques allow researchers to culture behavior outside the organism that usually produces it. This approach could contribute new insights in theoretical ethology on questions including the origins of social behavior and cooperation, adaptation, and the emergent nature of complex behavior. Animat studies applied to domestic animals have been few so far, and have involved simulations of space use by swine. I suggest other applications, including modeling animal movement during human handling and the effects of environmental enrichment on the satisfaction of behavioral needs. Appropriate use of animat models in a research program could result in savings of time and numbers of animals required. This approach may therefore come to be viewed as both ethically and economically advantageous.

Key Words: Animat, Behavior, Computer Simulation

 J. Anim. Sci. 1998. 76:2596–2604

Introduction

Until recently, ethology, the scientific study of animal behavior, has been practiced entirely through observation of living animals and their interactions with the environment. Recent advances in the study of artificial intelligence have presented ethologists with an opportunity to develop a new approach to animal behavior research. This article is intended to provide a nontechnical introduction to animat research. The term *animat* (Wilson, 1991) refers to a class of computer-simulated animals, or robots. These can exist within an artificial environment, either physical or computer-generated, devised by a researcher, or else can be designed to interact with the real world. The animat research paradigm has its origins in behavior-based artificial intelligence and artificial life research. This article examines these disciplines and some of the key concepts that they employ. The present and possible future application of these ideas for the study of behavior is examined. In particular, I consider some possible uses of computer-simulated

[1]Invited paper presented at a symposium titled "The Use of Computers for Modeling, Recognizing, and Recording Behavior" at the ASAS 89th Annu. Mtg., July 1997, Nashville, TN.

[2]I would like to thank Joseph Stookey for inviting me to present an early version of this paper at the Nashville symposium. Helpful comments on the manuscript were given by K. S. Schwartzkopf-Genswein, G. Flannigan, and C. S. Waltz. Thanks are also due to W. R. Stricklin and an anonymous reviewer. This article has benefited considerably from their suggestions.

Received October 12, 1997.
Accepted April 21, 1998.

animats as tools in applied animal behavior research. The development of animat techniques and their use in modeling animal behavior has previously been discussed by Stricklin et al. (1995), as background to their own animat-based studies of space use by swine. However, with the exception of this work, the considerable potential of the animat approach to animal behavior research remains largely unexplored to date.

Is Artificial Intelligence Intelligent?

The traditional goal of artificial intelligence research has been the creation of systems that exhibit humanlike intelligent behavior (Brooks, 1991). As the speed, power, and availability of computer hardware have increased, some systems have been developed to perform tasks that, if done by a human, would seem to require significant cognitive abilities. These include expert systems to provide decision support in areas such as medical diagnosis (Szolovits and Pauker, 1993). A recent, well-publicized, triumph of artificial intelligence is embodied in a chess-playing computer, Deep Blue, that is able to defeat the best of human players (Michie, 1997).

Most successful artificial intelligence applications exploit the ability of digital computers to store huge databases of information and to perform repetitive operations on those data at superhuman speed. The dominant approach to artificial intelligence until recently, also called the *knowledge-based*, or *top-down*, approach (Maes, 1993), has been to design systems that take advantage of this capability. Knowledge-based systems typically display great competence within a narrow specialty. The problem domain is precisely defined, and the structure of the program reflects these constraints. Likewise, the extent of knowledge required for the task is predefined. Machines like the chess-playing Deep Blue exemplify this approach.

However, Deep Blue plays chess quite differently than a human does (Michie, 1997). Expert players use pattern recognition to assess the configuration of pieces and the general directions the game might take from a small number of likely moves. Deep Blue lacks the intuitive skill of a grandmaster and instead relies on brute computing power to evaluate billions of future positions. Researchers in the artificial intelligence community recognize that knowledge-based artificial intelligence systems are not intelligent in the same sense that animals, including humans, are. Moreover, a goal of the field has been to rectify this discrepancy.

Behavior-Based Artificial Intelligence

Behavior-based, or *bottom-up*, artificial intelligence (Maes, 1993) has arisen as an alternative to the knowledge-based approach, and in some respects may be considered its antithesis. Behavior-based artificial intelligence exploits the complexity of behaviors that can arise in systems with very few rules and few restrictions on the type of output they can produce. Much research has focused on the interaction between the system and its environment, and the goal is to develop agents that are generally capable and self-directed within the environment, rather than specialized in a tiny domain and otherwise helpless. Such agents are often conceived as simulating the activity of an animal, whether real or hypothetical. Thus, behavior-based artificial intelligence has also become known as the *autonomous agent*, or *animat*, approach (Maes, 1995). Animat research may refer to work with embodied robots in the physical world or to simulated animals operating in virtual environments that are designed to present appropriate challenges (Tyrell and Mayhew, 1991).

In behavior-based artificial intelligence work there is developing an emphasis on robots (e.g., Brooks, 1990; Husbands et al., 1997). Some researchers feel that embodiment within a physical environment is vital to the development of meaningful intelligence (Steels, 1995). Robots have also proved useful in modeling the behavior of real animals (e.g., Webb, 1996; Kuwana et al., 1997). For most conceivable ethological investigations applied to farm animals, however, it seems likely that computer-generated animats, rather than robots, will be the synthetic subjects of choice.

Animats as Artificial Life

Animats are human-engineered behavioral systems that may be considered to have some attributes in common with living organisms. The new field of artificial life has been called *the biology of the possible* (Langton, 1996). Artificial life researchers suggest that conventional biology is limited in that it is concerned only with products of “natural” evolutionary processes. If it is possible to engineer artifacts or computer programs, which by some reasonable definition could be called “living,” then, as artificial life proponents argue, a comparative biology should develop to include study of such artificial organisms. Artificial life is closely related to, but not synonymous with, behavior-based artificial intelligence in that it emphasizes bottom-up methodology (Moreno et al., 1997). An artificial life approach could be used to investigate problems in biology (Dyer, 1995; Taylor and Jefferson, 1995) and agriculture (Mueller, 1995). The evolution of artificial organisms and the emergence of competing “species” have been modeled by Tom Ray (Kawata and Toquenaga, 1994; Ray, 1995). Parasitism and counteradaptations by hosts have evolved in these simulations. Biological processes have also inspired developments in artificial life. Computer viruses are replicating programs that infect and interfere with the function of host computers. Computer viruses are viewed by some as a form of artificial life (Spafford, 1995). A biologically inspired immune system to combat computer viruses has been developed by Kephart (1994). The study of self-replicating packages of information in computer systems may suggest new approaches in immunology and genetics. Boden (1996) introduces a comprehensive collection of essays on the philosophical issues in artificial life.

Simple Rules, Complex Behavior

A recurrent theme in animat studies is the generation of complex patterns of behavior by agents with limited capabilities defined by only simple sets of rules. Animats often behave in ways that are not explicitly specified in their programming. Such activities are said to be *emergent*. Emergence is particularly interesting when the activity appears to be adaptive for the animat in the particular context of the simulation. Experimenting with the conditions and rules that lead to emergent functionality in animat systems is valuable to artificial intelligence researchers in their quest for useful intelligence. For ethologists, there is a need to consider the role that emergence plays in the ability of animals to display adaptive and apparently sophisticated behavior. Recently, it has been shown (Schaub and Korol, 1996; Webb, 1996) that simple electronic and electromechanical robots can deal effectively with complex environments. Mindful of Lloyd Morgan's canon, researchers that study the cognitive bases of behavior should consider how effective behavior could occur without mental representations or other higher processes, before invoking these as explanations.[3]

The emergence of complex behavior from simple mechanisms has been considered in thought experiments by Braitenberg (1984). In one example (Figure 1), an imaginary vehicle has two light sensors on the front, each controlling the speed of the motor driving the opposite side rear wheel. As the amount of light falling on the sensor increases, so does the speed of the connected wheel. This vehicle should move always in the direction of the greatest light intensity and away from shadows, eventually colliding with the source; this behavior Braitenberg labels "aggressive."

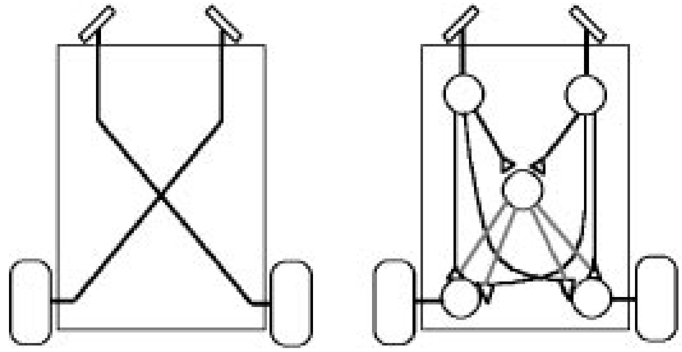

Figure 1. Two types of light-seeking vehicle. The hardwired vehicle (left) has two light sensors, each connected to the opposite-side rear wheel so that the wheel turns at a speed proportional to the intensity of light falling on the sensor (after Braitenberg, 1984). The vehicle on the right has a simple neural network between the sensors and the wheels. The central facilitatory neuron reinforces turns toward the light, allowing the vehicle to "learn" light-seeking behavior (after Scutt, 1994).

Real vehicles wired in this way have been constructed (Scutt, 1994), and they were capable of negotiating light-occluding obstacles en route to a light source. Scutt also describes a variation (Figure 1) in which the vehicle "learns" the light-seeking behavior itself. The connections between sensors and motors are represented as artificial neurons, initially set up to be fully and neutrally interconnected. The addition of a central facilitatory interneuron reinforces connections that cause it to be more active, increasing the chance of subsequent activity in those pathways. Initially turning at random, the vehicle gradually learns to seek the light. Learning occurs on the boundaries between light and shadow, where frequent changes of direction are caused. Vehicles learn most easily in environments containing obstacles and shadows, so that the light intensity varies from one point to another. Vehicles in barren environments with just an unshielded light source never learn light-seeking behavior. It is interesting to note that in rodents, complex environments are associated with enhanced learning (Rosenzweig and Bennett, 1996) and neural development (Carughi et al., 1990; Kempermann et al., 1997), compared with barren environments.

Some of the most interesting emergent behavioral phenomena, in both natural and artificial environments, occur when animals or agents work together as parts of a social whole, as for example, in the trail-following behavior of simulated ants (Colorni et al., 1992). The ants move back and forth along a trail, following traces of a pheromone that each deposits as

[3]Conwy Lloyd Morgan (1894) suggested parsimony in interpreting behavior. His principle, known as *Lloyd Morgan's canon*, urges that "in no case may we interpret an action as the outcome of the exercise of higher psychical faculty, if it can be interpreted as the outcome of the exercise of one which stands lower in the psychological scale." Of course, as some critics (e.g., Bateson, 1991) have observed, blind adherence to this rule is unimaginative and yields only the simplest interpretations, not necessarily the best. Despite criticism, the canon remains a valid caution against overly complex interpretations of behavior. In the present context it suggests the alternative that some "mindlike" behavior may be an emergent effect of sensorimotor mechanisms, or lower levels of neural processing.

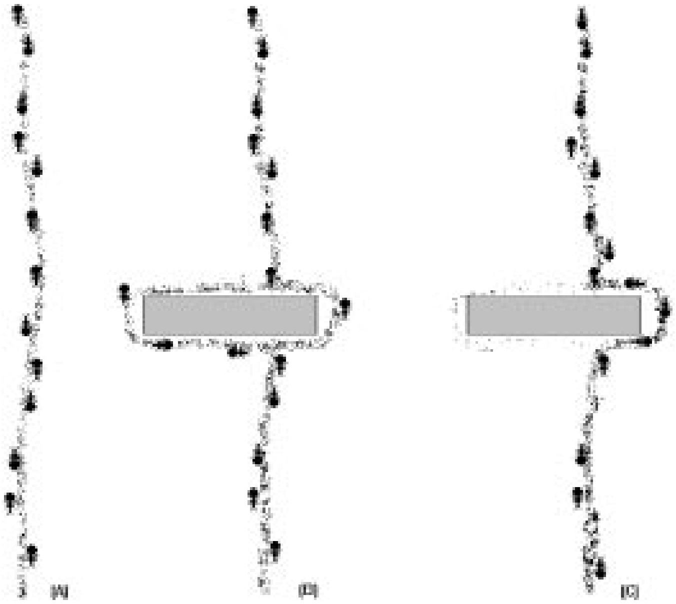

Figure 2. Emergent problem-solving by simulated ants (after Colorni et al., 1992).

it moves (Figure 2a). If the trail is blocked by a large object, the ants have a rule that tells them to go around the object until they encounter the trail again. They do not know which direction is shorter, so they turn randomly left or right (Figure 2b). However, if ants set off on left or right-handed journeys around the obstacle with equal frequency, in a given time more of the ants complete the short journey than the long one. For an ant coming the other way, this has the effect that the scent will be more concentrated on the short side, making the choice of this direction more attractive. Eventually, the short trail becomes much more attractive, and the long route is abandoned (Figure 2c). Selection of the short route is achieved without any knowledge that it is shorter. A single ant would not be able to do this, except by chance. Rather, it would continue to reinforce its original choice, long or short. Simulated ants following similar rules can solve traveling salesman problems. The traveling salesman problem is the task of finding the shortest route between several hypothetical cities. A recent ant program tackled sets of up to 1,577 cities (Dorigo and Gambardella, 1997). The solution provided in most cases was not a mathematically proven optimum, but was rather a close approximation obtained by progressive refinement over many trail-following journeys, an interesting example of emergent heuristic computation. Understanding the formation of trails used by free-ranging livestock to travel between resource locations may provide an opportunity to influence their movement patterns and improve the use of pasture resources.

On occasions, the activity of social groups of animats that follow simple rules can be strikingly reminiscent of the collective behavior of their living counterparts. A good example is the dynamics of flocking in birds. Craig Reynolds (1987) devised a simulation using animats he called *boids* to represent birds. These had three rules governing their spatial behavior in relation to nearby boids: 1) stay together (i.e., move toward nearby boids); 2) match speed to nearby boids; and 3) avoid collisions. These produced lifelike group movement, resembling that of real bird flocks. Reynolds experimented with placing obstacles in the path of his "flock", which had no specific rule to maintain group cohesion while negotiating obstacles. The animat flock was undaunted by the obstacles, simply dividing to avoid them and re-forming afterward. Bedau (1997) argues that the study of emergent, supple dynamics on a large scale from simple, rule-based systems, as exemplified by Reynolds's boids has relevance to the understanding of human mental phenomena. The dynamics of group movement in herds of hypothetical terrestrial agents have been investigated by Hodgkins and Brogan (1994). Werner and Dyer (1993) have modeled the evolution of herding behavior and predator avoidance in a simple simulated ecosystem.

Exploiting Variable Success in Virtual Environments

Animals in their natural environment show variation in form and behavior. Individual learning and natural selection exploit this variation, resulting in animals that are better adapted, phenotypically and behaviorally, to their environments. Animat models may make use of artificial neural networks and genetic programming to achieve comparable ends.

Artificial Neural Networks

Artificial neurons are analogous to the neurons of animals in terms of their connectivity and computational properties. Arrays of such neurons, known as *artificial neural networks*, have been used for all levels of behavioral control in animats from decision making to control of individual limb movements, as well as for the simulation of simple, but entire nervous systems (e.g., Beer, 1990). Gardner (1993) edits a useful guide to the structures and uses of neural networks. Figure 3 illustrates a common type of neural network. Stimulation of a neuron in the input layer is passed to neurons of subsequent layers to which it is connected via synapses that vary in strength. The probability of a neuron firing is proportional to the sum of the synaptic potentials it receives from neurons in the previous layer to which it is connected. The strength of these connections is assigned an initial value, which varies as the network "learns." Training is accomplished by shaping the network output through "reward," or reinforcement. Neural pathways that connect an input to an appropriate, or correct, output

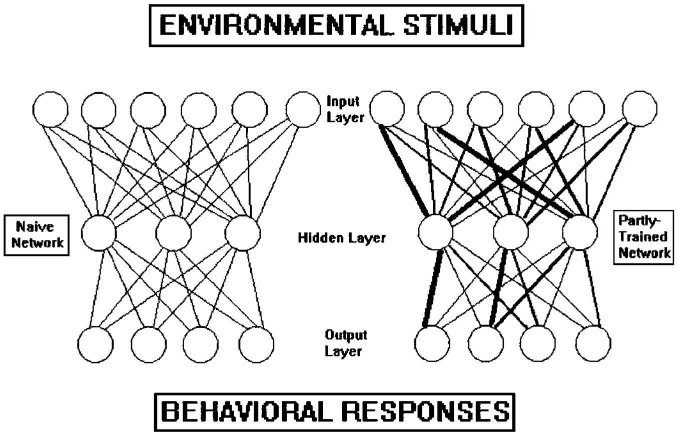

Figure 3. Diagram representing a typical neural network. Three layers of neurons (circles) are joined by synaptic connections (lines). Sensory information activates neurons in the input layer and is processed by one or more hidden layers, which in turn activate output neurons representing behaviors or choices. Initially, the network is fully and neutrally connected (left). As training proceeds, some connections are reinforced (represented by thicker lines), increasing the probability that these pathways will be used again (right).

are strengthened by a process of back-propagation, which acts on each connection from output through hidden layers back to the input layer. This increases the probability that a particular route through the network will be used again. This learning process may be directed by a human overseer, but in some cases networks can learn from interacting with their environment without active human intervention required.

Genetic Programming

The defining characteristics of an animat can be represented as lines of computer instructions or data within the simulation program. If the parameters of all animats are not identical, there may be variations in individual success in the simulation. Genetic programming acts on this variation to preserve and magnify adaptive traits. The set of instructions that defines a particular animat correspond to the genome of an animal. In a process analogous to biological reproduction, the genotypes of successful individuals are replicated in later generations of animats. Individuals with advantageous characteristics can be copied exactly, or the system can incorporate a rate of mutation to maintain some variability for future exploitation. Alternatively, as in Figure 4, two sets of instructions can be combined to form offspring with both parents' characteristics (for relevant discussions on genetic algorithms, see Mitchell and Forrest, 1995 and Ch. 19 in Haefner, 1996). As with neural networks, the selection process can be implemented by a person (corresponding to selective breeding of a domestic animal) or can be integral to the simulated environment.

Applied Animat Research

There has been almost no attempt thus far to use animats in applied animal behavior research. This is probably due to widespread unfamiliarity with the concepts. The exception has been work on the space requirements of pigs by Stricklin et al. (1995, 1997, 1998). In an interesting series of experiments, they have used an animat approach to model the effects of group size and pen shape and size on behavior. The authors believe that animat modeling has considerable potential in applied ethology as a means of investigating space requirements. They also propose that artificial intelligence research in general can contribute insights into cognition, awareness, and other psychological processes as well as philosophical issues of mind. Thus, they believe that artificial intelligence research is relevant to animal welfare issues.

One important advantage of computer simulation, which Stricklin et al. (1995) have discussed, is the ability to try out, in the initial development stages, the effect of pen design on the movement of pigs before actually constructing them. It might be productive to extend this idea to facilities for moving and handling animals on a large scale. One could envisage an animat-based computer-aided design system for developing efficient handling facilities for cattle and other livestock. This would probably entail a mixture of top-down and bottom-up thinking. We know that factors such as human presence, chute width and curvature, and changes in light intensity can influence animal movement (Grandin, 1993). However, the understanding of such complex relationships tends to be intuitive and is only somewhat empirical. This very complexity can make the subject of animal handling and chute design an ideal candidate for an animat approach.

A possible research strategy for developing a practical, animat-based, handling facility design simulator is summarized in Figure 5. An early task would be to devise a feature set for animats that resulted in a convincing simulation of animal movement. One could start by drawing up a list of all influences likely to affect the direction and speed of movement of an animal. These would include proximity of other animals and physical barriers, light levels, presence of a person, and other factors. Sophisticated models might incorporate influences such as hunger, effects of previous experience, and perhaps weather conditions. The causal links between these factors and the

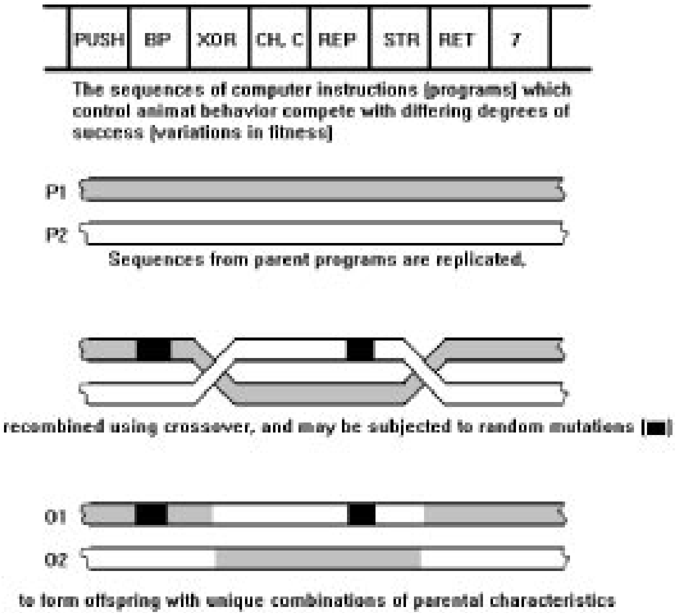

Figure 4. Genetic programming. In this case, sexual reproduction by recombination of copies of two parental algorithms.

movement of the animats could be set initially to arbitrary weights. One should not try to anticipate these relationships by imposing a fixed set of rules from the outset. The detailed behavior of a group of animats, just as with their live counterparts, is likely to be emergent and, thus, more or less unpredictable from initial conditions. It will be more productive to allow animats to learn to be reasonable mimics of real livestock. This could be achieved by using an artificial neural network to represent the relations between environmental influences and the resultant behavior. This would require a supervised training procedure in which animat behavior is shaped toward a lifelike approximation of animal movement. Alternatively, a realistic system could be selectively "bred" using a genetic algorithm approach that treats the weights of the various parameters as a genotype to be replicated. A combination of these two methods could also be used as Werner and Dyer (1993) have done in their modeling of predator-prey interactions. The most lifelike models would need to be validated against observations of live animals in chutes and alleys. Once one had arrived at an apparently valid model, it would be possible to experiment with altering features of the physical environment and study the ability of the model to demonstrate convincing animal movement in differently shaped facilities. Such a system could bring about considerable improvements in facility design, even if it were only used to eliminate hopelessly poor designs before committing resources to building them.

It might also be possible to use animat-based software to predict pasture usage and competition for resources by livestock. There have been attempts before to model space and resource usage of cattle at pasture using computers (e.g., Senft et al., 1983, 1985). These have been done using a top-down approach and have been moderately successful, if simplistic. They do not seem to have found any practical application to date. It should be possible to do better using a bottom-up approach. A good deal of research would be involved in constructing a behavior-based model of pasture usage by cattle or sheep, but I think the effort would be very worthwhile. A useful predictive model would require a substantial program of behavioral observation of animals at pasture. This is required not only for selecting parameters to incorporate into the model but also for subsequent validation and refinement. Such a project would be of ethological value in its own right. It is unlikely that a model could be developed that would predict the actions of individual animals with any degree of precision. However, this is not necessary in order for the system to be of practical use. The most useful application would probably be the ability to predict, at a herd or flock level, the effects that changing watering points or fencelines, improving drainage, felling woodland, and other management considerations have on the use of space and forage by the animals.

Animat-based investigations might also contribute to improvements in animal welfare by facilitating the development of enriched environments that help to meet the ethological needs of captive animals. Techniques such as the application of consumer demand economics (Nicol and Dawkins, 1990; Dawkins, 1983) allow us to measure the relative importance, or absolute necessity, of various environmental features and resources from an animal's point of view. This information should allow the development of animats with more or less realistic motivational states and goal-directed behavior. These models could be used to predict the effect on behavior and satisfaction of ethological needs that could be expected from any proposed changes to housing systems. This approach could be used in poultry housing to model space use and the welfare value of perches, different substrates, and other refinements. Animats could be used to model enriched environments for zoo animals. Effectively enriched zoo enclosures are not only of benefit to animal well-being, but they also offer the public a more interesting exhibit (e.g., Markowitz et al., 1995).

The animat approach could be used by behavioral ecologists to examine the applicability, to farm animals, of mathematical models of the optimization of resource usage. It is also possible to use some of the associated techniques to investigate cognitive processes (e.g., the processes of recognition and representation that subserve intraspecific communication). Recently, Reby et al. (1997) demonstrated the training of an artificial neural network to recognize the voices of individual deer. They taught the network using iterations of a number of vocalizations of four fallow deer. The network learned to achieve 100% recognition of the training samples and was able to predict which animal an unheard sample came from with over 90% success rate. This is significant because it models an active cognitive process, the representation of individual identity, which is an important element in the ethology of communication and the social context in which communicative exchanges take place. This kind of study can yield insights into the development and maintenance of individual recognition among real animals, and, thus, helps us to learn more about the cognitive basis of communication.

Value of Animats in Animal Behavior Research

The extent to which artificial life research will become recognizable as a distinct scientific discipline within the life sciences is difficult to predict. It is apparent though that computer simulation and modeling have achieved respectability as research tools in such disciplines as ecology, population biology, and neuroscience. The approach of modeling the behavior of individual organisms is, at least partly, a spin-off from conceptual developments in artificial intelligence research. This author contends that animats have considerable potential as research tools for the broad-minded behavior scientist.

The range of questions that could be tackled using an animat-based approach is extensive. Unlike some other biological disciplines, ethology lacks a technology for culturing the phenomenon of interest (i.e., the behavior) outside of the organism that normally produces it. The animat approach allows for the *ex vivo* generation of behavioral phenomena and the opportunity to manipulate the expression of the behaviors and the environment with precision, flexibility, and freedom from extraneous influences that are unlikely ever to be matched in studies with living animals. Modeling behavior in this way could facilitate the development of hypotheses that could be tested by further studies on the species represented by

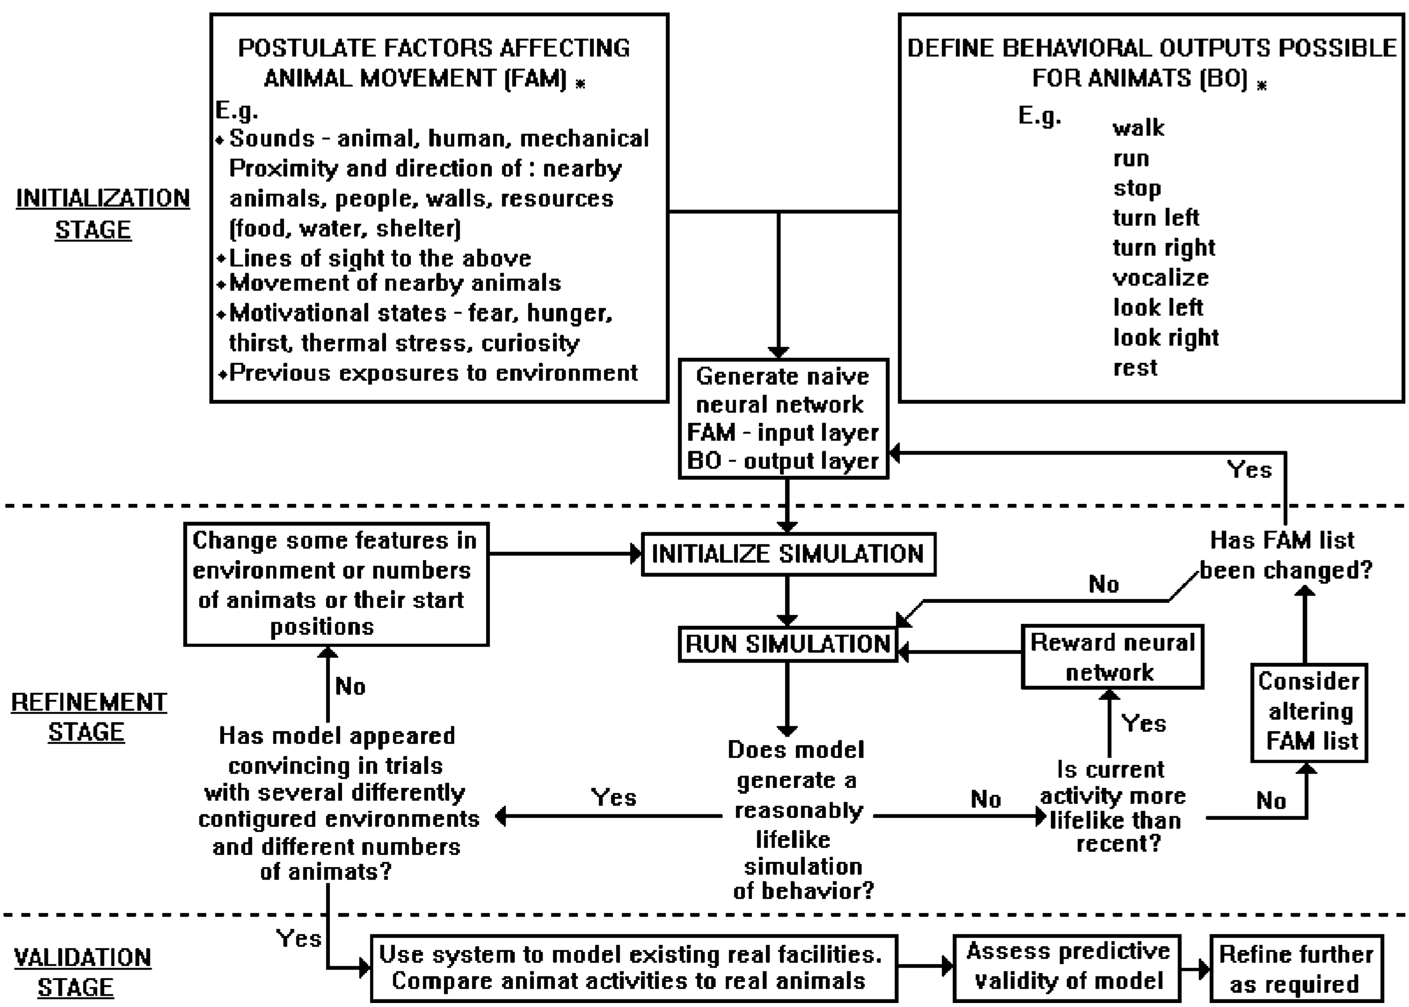

Figure 5. A possible approach to developing an animat-based system to model and predict livestock movements in handling facilities. (*The lists of factors that affect movement and behavioral outputs are not intended to be exhaustive.)

the model. This process would reduce the number of animals required for research or teaching purposes. It may therefore come to be regarded as an ethically advantageous and cost-effective strategy. As a heuristic, animat modeling offers a chance to generate some important insights into a variety of questions including the origins of sociality, cooperative behavior (including altruism), and adaptation through emergent functionality. Behavioral ecology can exploit this method to help validate mathematical models of optimality, game theory, signaling, and predator-prey interactions. A major advantage of the approach in this regard would be the ability to run many iterations of a theoretical model in a reasonable time while making small adjustments to the parameters. Animat simulations may offer something qualitatively different in that they permit, in principle at least, empirical evaluation of hypotheses that defy algorithmic expression.

Implications

The ideas and technologies that make animat-based behavior research possible are unfamiliar to many animal scientists. The use of animats could make useful contributions to animal science. There have been few efforts thus far to use animat methods in animal science research. The potential uses are many. Besides artificial intelligence, evolutionary biology, and neuroscience, there is ample scope for applications in livestock management. The extent to which animat methods will be used by applied animal behavior workers will depend on their willingness to embrace unfamiliar ideas and their imagination in identifying specific problems that could be tackled using animat simulations. From an animal welfare viewpoint, and as a cost-reducing measure, animat studies could be rewarding, because they may allow a reduction in the number of animals required in a research program. This consideration alone is sufficient to warrant a closer look at the potential of animat studies.

Literature Cited

Bateson, P. 1991. Assessment of pain in animals. Anim. Behav. 42: 827–839.

Bedau, M. A. 1997. Emergent models of supple dynamics in life and mind. Brain Cogn. 34:5–27.

Beer, R. D. 1990. Intelligence as Adaptive Behavior: An Experiment in Computational Neuroethology. Academic Press, San Diego, CA.

Boden, M. A. 1996. The Philosophy of Artificial Life. Oxford University Press Inc., New York.

Braitenberg, V. 1984. Vehicles: Experiments in Synthetic Psychology. MIT Press, Cambridge, MA.

Brooks, R. A. 1990. Elephants don't play chess. In: P. Maes (Ed.) Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back. pp 3–15. MIT Press, Cambridge, MA.

Brooks, R. A. 1991. Intelligence without representation. Artif. Intell. 47:139–159.

Carughi, A., K. J. Carpenter, and M. C. Diamond. 1990. Effect of environmental enrichment during nutritional rehabilitation on body growth, blood parameters and cerebral cortical development of rats. J. Nutr. 119:2005–2016.

Colorni, A., M. Dorigo, and V. Maniezzo. 1992. Distributed optimization by ant colonies. In: F. J. Varela, and P. Bourgine (Ed.) Toward a Practice of Autonomous Systems: Proceedings of the 1st European Conference on Artificial Life. pp 134–142. MIT Press, Cambridge, MA.

Dawkins, M. S. 1983. Battery hens name their price: Consumer demand theory and the measurement of ethological needs. Anim. Behav. 31:1195–1205.

Dorigo, M., and L. M. Gambardella. 1997. Ant colonies for the travelling salesman problem. Biosystems 43:73–81.

Dyer, M. G. 1995. Toward synthesizing artificial neural networks that exhibit cooperative intelligent behavior: Some open issues in artificial life. In: C. G. Langton (Ed.) Artificial Life: An Overview. pp 111–134. MIT Press, Cambridge, MA.

Gardner, D. 1993. The Neurobiology of Neural Networks. MIT Press, Cambridge, MA.

Grandin, T. 1993. Behavioral principles of cattle handling under extensive conditions. In: T. Grandin (Ed.) Livestock Handling and Transport. pp 43–57. CAB International, Wallingford, U.K.

Haefner, J. W. 1996. Modeling Biological Systems: Principles and Applications. Chapman and Hall, New York.

Hodgkins, J. K., and D. C. Brogan. 1994. Robot herds: Group behaviors for systems with significant dynamics. In: R. A. Brooks and P. Maes (Ed.) Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems. pp 319–324. MIT Press, Cambridge, MA.

Husbands, P., I. Harvey, D. Cliff, and G. Miller. 1997. Artificial evolution: A new path for artificial intelligence? Brain Cogn. 34:130–159.

Kawata, M., and Y. Toquenaga. 1994. From artificial individuals to global patterns. Trends Ecol. Evol. 9:417–421.

Kempermann, G., H. G. Kuhn, and F. H. Gage. 1997. More hippocampal neurons in adult mice living in enriched environments. Nature (Lond.) 386:493–495.

Kephart, J. O. 1994. A biologically inspired immune system for computers. In: R. A. Brooks, and P. Maes (Ed.) Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems. pp 130–139. MIT Press, Cambridge, MA.

Kuwana, Y., I. Shimanoyama, Y. Sayama, and H. Miura. 1997. A robot that behaves like a silkworm moth in the pheromone stream. In: C. G. Langton and K. Shimohara (Ed.) Artificial Life V: Proceedings of the Fifth International Workshop on the Synthesis and Simulation of Living Systems. pp 370–376. MIT Press, Cambridge, MA.

Langton, C. G. 1996. Artificial Life. In: M. A. Boden (Ed.) The Philosophy of Artificial Life. pp 39–94. Oxford University Press Inc., New York.

Maes, P. 1993. Behavior-based artificial intelligence. In: J-A. Meyer, H. L. Roitblat, and S. W. Wilson (Ed.) From Animals to Animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior. pp 2–10. MIT Press, Cambridge, MA.

Maes, P. 1995. Modeling adaptive autonomous agents. In: C. G. Langton (Ed.) Artificial Life: An Overview. pp 135–162. MIT Press, Cambridge, MA.

Markowitz, H., C. Aday, and A. Gavazzi. 1995. Effectiveness of acoustic prey: environmental enrichment for a captive African leopard (*Panthera pardus*). Zoo Biol. 14:371–379.

Michie, D. 1997. Slaughter on seventh avenue. New Sci. 154:26–29.

Mitchell, M., and S. Forrest. 1995. Genetic algorithms and artificial life. In: C. G. Langton (Ed.) Artificial Life: An Overview. pp 267–289. MIT Press, Cambridge, MA.

Moreno, A., J. Umerez, and J. Ibañez. 1997. Cognition and life: The autonomy of cognition. Brain Cogn. 34:107–129.

Morgan, C. L. 1984. An introduction to comparative psychology. P53. London: Walter Scott Limited.

Mueller, R.A.E. 1995. Artificial life: Frequently asked questions (FAQs) and partially baked ideas (PBIs) about its uses in agricultural research. In: A. J. Udink Ten Cate, R. Martin-Clouaire, A. A. Dijkhuizen, and C. Lokhorst (Ed.) Artificial Intelligence in Agriculture. pp 151–159. Elsevier, Oxford, U.K.

Nicol, C., and M. S. Dawkins. 1990. Homes fit for hens. New Sci. 125:46–51.

Ray, T. S. 1995. An evolutionary approach to synthetic biology: Zen and the art of creating life. In: C. G. Langton (Ed.) Artificial Life: An Overview. pp 179–210. MIT Press, Cambridge, MA.

Reby, D., S. Lek, I. Dimopoulous, J. Joachim, J. Lauga, and S. Aulagnier. 1997. Artificial neural networks as a classification method in the behavioural sciences. Behav. Processes 40:35–43.

Reynolds, C. W. 1987. Flocks, herds and schools: A distributed behavioral model. Comput. Graphics 21:25–34.

Rosenzweig, M. R., and E. L. Bennett. 1996. Psychobiology of plasticity: Effects of training and experience on brain and behavior. Behav. Brain Res. 78:57–65.

Schaub, B., and T. Korol. 1996. Junkyard 'bots. Equinox 90(Nov/Dec):38–45.

Scutt, T. 1994. The five neuron trick: Using classical conditioning to learn how to seek light. In: D. Cliff, P. Husbands, J-A. Meyer, and S. W. Wilson (Ed.) From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior. pp 364–370. MIT Press, Cambridge, MA.

Senft, R. L., L. R. Rittenhouse, and R. G. Woodmansee. 1983. The use of regression models to predict spatial patterns of cattle behavior. J. Range Manage. 36:553–557.

Senft, R. L., L. R. Rittenhouse, and R. G. Woodmansee. 1985. Factors influencing patterns of cattle grazing behavior on shortgrass steppe. J. Range Manage. 38:82–87.

Spafford, E. H. 1995. Computer viruses as artificial life. In: C. G. Langton (Ed.) Artificial Life: An Overview. pp 249–266. MIT Press, Cambridge, MA.

Steels, L. 1995. The artificial life roots of artificial intelligence. In: C. G. Langton (Ed.) Artificial Life: An Overview. pp 75–110. MIT Press, Cambridge, MA.

Stricklin, W. R., J. Z. Zhou, and H. W. Gonyou. 1995. Selfish animats and robot ethology: using artificial animals to investigate social and spatial behavior. Appl. Anim. Behav. Sci. 44: 187–203.

Stricklin, W. R., J. Z. Zhou, and H. W. Gonyou. 1997. Computer-based investigations of the social and spacing behavior of confined animals. J. Anim. Sci. 75(Suppl. 1):137 (Abstr.).

Stricklin, W. R. et al. 1998. Artificial pigs in space: Using artificial intelligence and artificial life techniques to design animal housing. J. Anim. Sci. 1998. 76:2609–2613.

Szolovits, P., and S. G. Pauker. 1993. Categorical and probabilistic reasoning in medicine revisited. Artif. Intell. 59:167–180.

Taylor, C., and D. Jefferson. 1995. Artificial life as a tool for biological inquiry. In: C. G. Langton (Ed.) Artificial Life: An Overview. pp 1–13. MIT Press, Cambridge, MA.

Tyrell, T., and J.E.W. Mayhew. 1991. Computer simulation of an animal environment. In: J-A. Meyer, and S. W. Wilson (Ed.) From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior. pp 263-273. MIT Press, Cambridge, MA.

Webb, B. 1996. A cricket robot. Sci. Am. 275 (Dec):94–99.

Werner, G. M., and M. G. Dyer. 1993. Evolution of herding behavior in artificial animals. In: J-A. Meyer, H. L. Roitblat, and S. W. Wilson (Ed.) From Animals to Animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior. pp 393–399. MIT Press, Cambridge, MA.

Wilson, S. W. 1991. The animat path to AI. In: J-A. Meyer and S. W. Wilson (Ed.) From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior. pp 15–21. MIT Press, Cambridge, MA.