Neuroethology, Computational

Dave Cliff

Hewlett-Packard Laboratories
Bristol BS34 8QZ, England U.K.
dave_cliff@hp.com

Running head: Neuroethology, Computational

Contact details:

Phone: +44 117 312 8189
Fax: +44 117 312 8924
E-mail: dave_cliff@hp.com

To appear in M. A. Arbib (ed) *The Handbook of Brain Theory and Neural Networks,* 2nd edition, MIT Press.

As with other entries in the *Handbook*, this article was constrained to be no longer than 4500 words and to have no more than 15 references. A star (*) next to a reference indicates that it is expository. Within the text, cross-references IN CAPITALS refer to other articles in the *Handbook*.

INTRODUCTION

Over the past decade, a number of neural network researchers have used the term *computational neuroethology* to describe a specific approach to *neuroethology*. Neuroethology is the study of the neural mechanisms underlying the generation of behavior in animals, and hence it lies at the intersection of neuroscience (the study of nervous systems) and ethology (the study of animal behavior); for an introduction to neuroethology, see Simmons and Young (1999). The definition of computational neuroethology is very similar, but is not quite so dependent on studying animals: animals just happen to be biological *autonomous agents*. But there are also non-biological autonomous agents such as some types of robots, and some types of simulated embodied agents operating in virtual worlds. In this context, autonomous agents are self-governing entities capable of operating (i.e., coordinating perception and action) for extended periods of time in environments that are complex, uncertain, and dynamic. Thus, computational neuroethology can be characterised as the attempt to analyze the computational principles underlying the generation of behavior in animals and in artificial autonomous agents. For the sake of brevity in the rest of this article, autonomous agents will be referred to simply as ``agents", and computational neuroethology will be abbreviated to CNE.

CNE can be distinguished from classical computational neuroscience by its increased emphasis on studying the neural control of behavior within the context of neural systems that are both embodied and situated within an environment. The “computational” nature of CNE comes not so much from treating neural systems as inherently computational devices, but rather in the use of sophisticated computer-based simulation and visualization tools for exploring issues in neuroethology.

Put most simply, CNE involves the use of computational modelling in trying to understand the neural mechanisms responsible for generating `useful' behaviors in an agent. The word `useful' is rather imprecise: it is more common to talk of *adaptive* behaviors. In the ethology literature, an adaptive behavior is usually defined as a behavior that increases the likelihood that an animal will survive long enough to produce viable offspring. Often implicit in this definition is the assumption that the animal's environment is sufficiently unforgiving (or hostile) that if the animal does nothing, it will die before it can reproduce. In studying artificial agents, the utility of behavior is frequently evaluated by less harsh criteria, such as scoring observed behaviors according to some metric that indicates how close they come to satisfying some set of performance objectives or criteria.

Neural networks that generate adaptive behavior should not be confused with *adaptive* neural networks, where connection strengths may alter as a result of experience. Adaptation or plasticity may itself give rise to new or improved adaptive behaviors, but there are many cases of adaptive behaviors that are genetically determined (e.g., ``hard-wired'' behaviors such as reflexes and instincts).

When CNE is approached in the context of adaptive behavior research, it becomes clear that the neural system is just one component in the *action-perception cycle*, where an agent’s actions may alter what information it perceives concerning its environment, and where those alterations in perceived information may lead to changes in the agent's internal state, and where those changes in state may in turn affect further actions, thereby affecting what information is subsequently perceived, and so on. Thus, crucially, the agent's nervous system and body and environment all combine to form a tightly coupled dynamical system. This is a notion long stressed by Arbib:

> ``In speaking of human perception, we often talk as if a purely passive process of classification were involved -- of being able, when shown an object, to respond by naming it correctly. However, for most of the perception of most animals and much of human behavior, it is more appropriate to say that the animal perceives its environment to the extent that it is *prepared* to *interact* with that environment in some reasonably structured fashion."
> (Arbib, 1972, p.16)

As defined thus far, CNE may not seem to be particularly distinguishable from most work in neural network research. After all, many people in computational neural network research might argue that their work will, ultimately, lead to understanding of the neural mechanisms underlying the generation of (some) adaptive behaviors. For example, face recognition is an adaptive behavior in humans, and could probably be classed as an adaptive behavior in, say, a security robot. So why can't a back-propagation network that learns to distinguish between photographs of human faces (for example) be classed as work in CNE?

MOTIVATIONS

Typically, artificial neural network models employ homogeneous groups of highly idealised and simplified neuron models (called *units*), connected in a regular fashion, which exhibit some form of `learning' or adaptation. The large majority of such models can be described in essence as mapping or transforming between representations: input data is presented to the network in a particular format, and the network is judged successful when its outputs can be interpreted as a correct representation of the results of performing the desired transformation. In almost all cases, the input and output representation formats are pre-specified by the experimenter (although it should be acknowledged that this is not entirely true of unsupervised learning networks, and that there are a number of artificial neural network models that draw inspiration from biological data in their choice of input and output representations). If such networks are to be employed in artificial agents, or are to be of use in understanding biological agents, then this can only be so under the (often unspoken) assumption that, eventually, it will be possible to assemble a `pipeline' of such input-output transducer networks that links sensory inputs to motor outputs, and produce adaptive behavior. The most significant issue here is the heavy dependence on *a priori* intermediate representations, which may not be justifiable: neural sensory-motor pathways generating adaptive behaviors might not be neatly partitioned into representation-transforming modules; such pathways may not exhibit any patterns of activity identifiable as a representation in the conventional sense, and even if they do, there is no guarantee that they will be in strong accordance with representations chosen *a priori* by modellers.

This should not be mistaken for an argument against representation, nor for a denial of the vital role played by internal states in the generation of adaptive behaviors: it is simply an awareness of the dangers of being misled by *a priori* notions of representation. One of the safest ways of avoiding these dangers is to model, as far as is possible, *entire* sensory-motor pathways (i.e., the complete sequence of neural processing, from sensory input to motor output) involved in the generation of adaptive behavior. This requires that the agent is studied while *situated in an environment*: most sensory-motor processing for adaptive behavior involves dynamic interaction with the environment; and a situated agent is part of a closed-loop system, because certain actions can affect subsequent sensory inputs. The sensory-motor pathway should thus not be viewed as a `pipeline' transforming from a given input representation to a desired output representation, but rather as one link in the action-perception cycle.

When such an approach is adopted, the true nature of the representations and processing necessary for the generation of relatively complex adaptive behaviors is more likely to be revealed, and the validity of any *a priori* assumptions is clarified.

Naturally, it is beyond the state of the art to attempt to model complete sensory-motor pathways in humans or other large mammals, but experimental work in the neuroethology literature provides a wealth of data from less intellectually able animals, such as arthropods (the animal class which includes insects, spiders, and crustacea), amphibia, and other “simple” vertebrates such as eels or salamanders. Such animals are used as the domains of study in some CNE research, but in other work simple idealized models are rigorously studied, in a manner akin to Galileo’s models of perfect spheres rolling down inclined frictionless planes.

The argument that *a priori* commitment to certain representations or architectures for sensory-motor processing can lead to surprisingly wrong conclusions can be illustrated by reference to a classic series of thought-experiments devised by Braitenberg (1984). Braitenberg describes specifications for a series of simple mobile vehicles, operating in a world with simplified kinematics. The series of vehicles starts with an elementary device that performs primitive heat-seeking behavior; it progresses through vehicles that exhibit positive or negative taxes (i.e., orientation towards or away from a directional stimulus), and primitive forms of learning, pattern detection, and movement detection; culminating in vehicles that exhibit chaotic dynamics and predictive behavior. The internal control mechanisms of all the vehicles are rigorously minimal: the simpler vehicles contain nothing more than wires connecting sensors to actuators; while the more advanced ones employ nonlinear threshold devices with delays and pseudo-Hebbian adaptation.

Braitenberg notes that the psychological language indicative of intentional mental states has compelling intuitive appeal in describing the observed behavior of the vehicles. He ascribes *fear, aggression, love, values and taste, rules, trains of thought, free will, foresight, egotism* and *optimism* to his vehicles. But he also demonstrates that while such terms may be very useful at the level of description of an external observer, the internal causal mechanisms could be surprisingly simple and, crucially, could contain nothing that can meaningfully be said to either “represent” or “implement” these intentional mental states. That is to say: the intentionality is in the eye of the beholder, not in the workings of the agent. For further discussion of these issues, see Cliff and Noble (1997).

While Braitenberg's vehicles are nothing more than thought-experiments, they provide insight to possible organisational principles in natural and artificial creatures, and demonstrate the limits of applicability of intentional terminology. Further discussion of the utility of agent models in biology can be found in Dean (1998).

To summarise: research in CNE can be characterised as placing increased emphasis on modelling entire adaptive-behavior-generating sensory-motor pathways in embodied agents, where those agents are situated in environments that supply sensory-motor feedback. Such an approach lessens the chances of making untenable assumptions concerning issues of representation and processing. Moreover, in order to study such pathways where there is reliable biological data, it may often be necessary to focus attention on relatively simple animals such as arthropods or amphibia. For further discussion of the rationale for CNE, see Beer (1990), and Cliff (1990).

It is important to note that there is a tradition of related work in the artificial neural network literature: research in reinforcement learning for control tasks is most close; see REINFORCEMENT LEARNING FOR MOTOR CONTROL.

SELECTED CURRENT RESEARCH PROJECTS

Two specific long-standing CNE research programs are discussed in this section: the work of Arbib's research group on visuomotor behavior in simple vertebrates; and Beer's work on the neural foundations of adaptive behavior in even simpler agents; i.e. in cockroaches and in abstract idealized agents. Before delving into these bodies of work, it is useful to consider how they sit in the CNE canon, and to point to CNE research that resides at other points in that space.

For the purposes of framing, there are three major axes along which work in CNE can be categorised. In no particular order, they are: the degree of reliance on computer software simulation; the degree of concentration on a specific animal species or class; and the extent to which semi-automated design techniques are employed.

The degree of reliance on computer simulation in CNE research projects varies from the complete, where all work is carried out using software simulations; to the minimal, where the CNE model takes the form of an operational physical robot, with the model neurons (individually or at the network level) being constructed from electronic circuits. Examples of the former include work by Arbib (1987) and Beer (1990) while many examples of the latter are discussed by Webb (2002). Note also that both Arbib and Beer went on to use robot platforms in continuations of their work that was initially software-only (Beer et al 1992, Arbib & Liaw 1995). A comprehensive review of the merits of using physical robots (rather than computer software) as simulations of animals has recently been published by Webb (2001), which includes copious references to work in this field; and see also BIOLOGICALLY INSPIRED ROBOTS.

The extent to which CNE research projects concentrate on a specific animal species or class varies from, at one extreme, CNE studies of one specific species (e.g., Beer, 1990); through generic CNE studies of several species of animals within the same order (e.g., Arbib's 1987 work on anuran visual control of action); to the other extreme where neural mechanisms underlying the generation of adaptive behavior in wholly abstract and idealised agents is explored within the CNE methodology (e.g., Beer, 2002).

Finally, with the continuing falls in the real cost of processor power and memory and disk storage, there has been an increased tendency over the past decade to move away from hand-designed computational/robot models, toward models that are the product of automated or semi-automated design processes. The use of evolutionary computation techniques such as genetic algorithms in particular (see EVOLUTION OF ARTIFICIAL NEURAL NETWORKS) has proved fruitful. At the "hands-on" extreme, there are CNE models where each artificial neuron's parameters (e.g. its time constants, thresholds, and connectivity to other components) are specified by the designer of the model (e.g. Arbib, 1987; Beer, 1990). Whereas at the "hands-off" extreme the modeller sets up a (usually truly vast) space of possible network designs and then uses an evolutionary search process to identify points in that design space which best satisfy some performance metric (i.e., the fitness evaluation function). Examples of this latter approach include Isjpeert (2001) and Beer (2002).

Computational Frogs, Toads, and Salamanders

Probably the most mature body of work in CNE is the research program led by Arbib for two decades on a family of models of visually mediated behavior in simple vertebrates. In the initial years of this project the focus was on visuomotor activity in frogs and toads: see (Arbib, 1987) for a review of the project with peer commentary; and Arbib (1997) for discussion of how this work integrates with studies of monkeys and of rats. Arbib named his simulation model *Rana computatrix*, the computational frog, in homage to W. Grey Walter's seminal *Machina Speculatrix* robots from the 1950's.

The *R. computatrix* models are faithful to the known biology, and there is an interplay between the experimental and theoretical work: an initial first approximation model was extended and refined in a number of stages, leading to a family of models.

Arbib's approach involves the definition of a number of functional *schemas*: schemas can be modelled by interacting layers of neuron-like elements, or by nets of intermediate-level units; the network models can be related to experimental data concerning neural circuitry, and the development process iterates (Arbib, 1987, p.411 ff.). Further details can be found in SCHEMA THEORY.

The primary focus in the *R. computatrix* models has been on how frogs and toads use vision to detect and catch prey, in environments that include obstacles and barriers. Arbib has developed a series of schema-based models that account for depth perception as interaction between accommodation and binocular clues, and at the lowest level the schemas are plausibly based on known details of the relevant neurological data.

One of the more striking results from this work, with reference to Marr's well-known theory of vision, is the indication that (in frogs and toads at least) there are different perceptual mechanisms for different visual stimuli. That is, the depths to prey and to barriers are extracted from the optic array by different processing channels, and are integrated in the sensorimotor pathways much later than Marr's theory might suggest. Arbib and Liaw (1995) went on to demonstrate how lessons learned from the *R. computatrix* project could inform the design of visually-guided robot systems.

In more recent work, Ijspeert and Arbib (2000) have reported on experiments where a sophisticated simulation of a 3-D multi-segmented bio-mechanical model of a salamander’s body is controlled by a complex neural network model. The network is composed of separate central pattern generators (CPGs – see LOCOMOTION, VERTEBRATE) for the body and the limbs, each of which may be activated and modulated by descending tonic inputs. Ijspeert and Arbib use this simulation system to explore the neural circuitry underlying the generation of visually steerable salamander locomotion behaviors in water and on land. One notable aspect of this work in relation to the earlier studies of anuran circuitry is that, while the gross morphology of the CPG circuits is decided by the experimenters, a genetic algorithm is used to determine: the fine details of the CPG circuits' internal connectivity and parameter values; the intersegmental coupling; and the coupling between the limb CPG and the body CPG. Thus, unlike the hands-on incremental modelling employed in the *R. computatrix* models, the salamander model is the product of a semi-automatic evolutionary design process.

Computational Cockroaches, and *Vehicles* Redux

Beer's (1990) book contains both methodological arguments for CNE, and also details of experimental work on his model of a computational cockroach, *Periplaneta computatrix*, which is a simulated hexapod agent embedded in an environment, inspired by neuroethological studies of the cockroach *Periplaneta americana*. The real cockroach uses chemotaxis as one of several strategies to locate food sources. If its path along an odour-gradient is blocked by an obstacle, then it performs stereotyped `edge-following' behavior. The artificial cockroach is controlled by a heterogeneous neural network which was inspired by biological data, and has been used to study issues in locomotion, guidance, and behavioral choice.

The primary external sensory input was simulated chemosensory information: patches of food in the environment gave off odour gradients detectable under an inverse square law relating distance to odour intensity. The neural nets also received mechanosensory input from proprioceptors in the limbs and tactile sensors which signal the presence of food under the mouth. The simulation model included elementary kinematics: if the artificial cockroach failed to adopt a stable position for a sufficient length of time, it fell down.

Results from the simulation sessions demonstrated behavior in the model that was highly similar to behavior in the real animal, and Beer subsequently performed “lesion” experiments by selectively deleting connections or units from the *P. computatrix* control network. Again, the results from the artificial system were in agreement with the biological data.

P. computatrix was inspired by biological data, but was not intended as a biological model. The various behaviors were generated by heterogeneous neural networks. The neuron model employed by Beer was more faithful to biology than many of the “formal neurons” used in conventional artificial neural network research: the units involved differential equations modelling membrane potentials, which gave his model neural assemblies a rich intrinsic dynamics. For further details, see LOCOMOTION, INVERTEBRATE.

The central focus in Beer’s (1990) work was on designing architectures composed from such neural units that could act as controllers for the various behaviors that *P. computatrix* should exhibit. Thus, there was no treatment of learning in the initial body of work on the cockroach. Subsequently, Beer reported on work that extended the original *P. computatrix* simulation model, testing it by allowing it to control walking in a real hexapod robot (Beer et al.1992).

In the robot implementation, the control network was still simulated (i.e., the units in the neural network were not realised physically) but the sensorimotor connections to the artificial neural network were interfaced to physical sensors and actuators by means of analogue-digital and digital-analogue converters. Beer et al. report that in all cases, the response of the physical robot was highly similar to that previously observed in simulation. The implementation did however reveal one problem in the controller which had not been examined in the simulation. This problem (involving disturbances in the crossbody phasing of the legs) was easily rectified, but nevertheless this demonstrates that simulation models cannot be trusted as *perfectly* replicating any physical implementation they may ultimately be intended for.

For a wider unified perspective on this work, see LOCOMOTION, INVERTEBRATE; VISUOMOTOR COORDINATION IN FROG AND TOAD; and LOCOMOTION, VERTEBRATE.

Subsequent to his work on *P. computatrix*, one line of research that Beer has pursued is in comparison radically simplified, divorced from any specific animal; and yet in its simplicity it reaches to the core of fundamental issues in cognitive science and adaptive behavior research. Rather than be constrained (and potentially confused) by biology, Beer (2002) developed a series of simple idealised embodied and embedded model agents, each of which is capable of "minimally cognitive" behaviours. Beer defines a minimally cognitive behavior as one that is just above the threshold for raising issues of genuine interest to cognitive science (see also COGNITIVE MODELLING: PSYCHOLOGY AND CONNECTIONISM.)

Beer's minimal agents exist in a two-dimensional world, but can only move along a bounded horizontal base-line. Various geometric shapes such as circles or diamonds drop from above, toward the agent's base-line. In each experiment, the intention is that the minimal agents use their sensors to detect the nature of whatever geometric shape or shapes is or are currently falling toward it, and thereby generate behavior "appropriate" to the current situation. The definition of appropriate behavior depends on the experiment, but may for example be as apparently trivial as "intercept circular objects and avoid diamond-shaped ones". To achieve this sensorimotor coordination, each minimal agent is equipped with a small continuous-time recurrent neural network (CTRNN) – see RECURRENT NETWORKS: SUPERVISED LEARNING.

The CTRNN for each minimal agent has a small number (e.g., seven) of fixed-orientation ray-casting "visual" proximity sensors (that each send a straight limited-length ray out at a particular angle to the agent's body and report on how far the ray travelled before it intercepted an object, if at all). Each sensor feeds onto a small number (e.g., five) of fully interconnected "interneurons", and all of these in turn feed onto a small number (e.g., two) of "output" neurons: one for moving to the left and one for moving to the right. Thus, a typical minimal agent may have fourteen units and perhaps seventy connection-weights in its CTRNN.

Any particular design for a CTRNN sensorimotor controller for one of Beer's minimal agents specifies the time-constant, bias, gain, and input weights for each neuron. Rather than design appropriate networks by hand, Beer employs a "hands-off" genetic algorithm to explore a very large space of possible network designs, evaluating each design on a measure of its observed behavior. To halve the size of the search space, Beer imposed a bilateral symmetry requirement. Other than this enforcement of symmetry, there is very little a priori commitment to any particular CTRNN solution. Over a reasonably small evolutionary experiment (e.g. 2000 generations with a population size of 100), minimal agents evolve that reliably score well on the experiment's evaluation function, and that also generalise well to situations not encountered in the evolutionary adaptation period.

So far, so simple. Yet, in a series of papers published since 1996, Beer and his colleagues have reported on the evolution of CTRNNs for sensorimotor control in minimally cognitive agents that have been evaluated on the basis of their ability to perform a variety of increasingly sophisticated behaviors. These behaviors include: orientating toward and reaching for objects; discrimination between objects; judging the passability of openings relative to the agent's own body size; discriminating between visible parts of the agent's body and other objects in the agent's environment; predicting and remembering the future location of falling objects so that they can later be intercepted "blind"; and switching attention between multiple objects as they fall. All with the same simple agent CTRNN architecture outlined above.

This array of cognitively interesting behaviors achieved by Beer's minimally cognitive agents prompts the question of what, precisely, is happening at the mechanistic level within the evolved CTRNNs to generate these behaviors. And at this point we return to the arguments and issues explored in the opening sections of this article. Beer presents concrete analyses of the CTRNNs of these agents, demonstrating a full understanding of their mechanistic activity from a *dynamical systems* perspective; and yet, as he points out, this analysis is of little or no use in attempts at elucidating an understanding from a *computational* (and hence *representational*) perspective: there is nothing readily identifiable in the CTRNNs that represents a circle or a diamond, or the action of intercepting or of avoiding. Rather, a full explanation of the behavior exhibited by one of Beer's minimally-cognitive agent's CTRNNs can only be given in the context of the dynamics of that agent's embodiment and of the environment that it is situated within. See (Beer, 2002; Cliff & Noble 1997) for further details.

DISCUSSION

Computational neuroethology studies neural mechanisms that generate adaptive behaviors, and hence requires that embodied agents are studied within the situated context of their environmental and behavioral niches.

From the above descriptions, some patterns emerge: the animal-specific CNE projects mentioned are dependent on the availability of fairly detailed neuroethological data. Such data invariably comes from invasive *in vivo* experimentation, and the neuroanatomy of "simpler" animals such as arthropods or simpler vertebrates is particularly amenable to such techniques: for arthropods in particular, certain neurons performing particular functions are readily locatable in different individual animals of the same species. While there are manifest obstacles preventing the collection of such data from more complex vertebrate subjects, research in these areas is making significant progress: see ACTION MONITORING AND FORWARD CONTROL OF MOVEMENTS; ARM AND HAND MOVEMENT CONTROL; EYE-HAND COORDINATION IN REACHING MOVEMENTS; MOTOR CORTEX – CODING AND DECODING OF DIRECTIONAL OPERATIONS; PURSUIT EYE MOVEMENTS; REACHING MOVEMENTS – IMPLICATIONS FOR CONNECTIONIST MODELS; SENSORIMOTOR LEARNING; and VESTIBULO-OCULAR REFLEX. Furthermore, by definition, any truly *general* principles underlying the neural generation of adaptive behaviors are those which are common to a number of species, so only cross-species studies will help identify general principles (Cliff, 1990, p.37).

Yet surely the most general principles of all are those that apply to all agents within a certain class of cognitive or behavioral niches, regardless of the hardware (or software) that those agents are implemented in. In this respect, Beer’s minimally cognitive agents are highly cogent. Until the representation-manipulating explanatory language that has traditionally been brought to bear on the supposed neural behavior-generating mechanisms of “complex” animals (including humans) can be demonstrated to be routinely applicable to “simpler” agents (including Beer’s *Vehicle*-like minimal agents), the rigour and limits of that explanatory language will remain in doubt.

REFERENCES

* Arbib, M.A., 1972, The Metaphorical Brain: an introduction to cybernetics as artificial intelligence and brain theory, New York: Wiley-Interscience.

Arbib, M.A., 1987, Levels of modelling of mechanisms of visually guided behavior. The Behavioral and Brain Sciences, 10:407--465.

Arbib, M.A., and Liaw, J., 1995, Sensorimotor transformations in the worlds of frogs and robots. Artificial Intelligence 72:53-79.

Arbib, M.A., 1997, From visual affordances in monkey parietal cortex to hippocampal-parietal interactions underlying rat navigation. Philosophical Transactions of the Royal Society of London Series B 352:1429-1476.

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

Beer, R.D., Chiel, H.J., Quinn, R.D., Espenschied, K., and Larsson, P., 1992, A distributed neural network architecture for hexapod robot locomotion. Neural Computation, 4(3):356--365.

Beer, R.D., 2002, The Dynamics of Active Categorical Perception in an Evolved Model Agent. The Behavioral and Brain Sciences, forthcoming. http://vorlon.cwru.edu/~beer/Papers/BBSpaper.pdf

* Braitenberg, V., 1984, Vehicles: Experiments in Synthetic Psychology, Cambridge MA: MIT~Press Bradford Books.

Cliff, D., 1990, Computational neuroethology: A provisional manifesto, in From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior (SAB90), (J.-A. Meyer and S.W. Wilson, Eds.), Cambridge MA: MIT Press Bradford Books, pp.29—39. http://citeseer.nj.nec.com/cliff91computational.html

* Cliff, D. and Noble, J., 1997, Knowledge-Based Vision and Simple Visual Machines. Philosophical Transactions of the Royal Society of London (Series B), 352(1358):1165-1175. ftp://ftp.cogs.susx.ac.uk/pub/reports/csrp/csrp459.ps.Z

Dean, J. 1998, Animats and what they can tell us. Trends in Cognitive Science, 2(2):60-67.

Ijspeert A.J., 2001, A connectionist central pattern generator for the aquatic and terrestrial gaits of a simulated salamander. Biological Cybernetics 84(5):331-348.
http://rana.usc.edu:8376/~ijspeert/PAPERS/BC.pdf

Ijspeert, A.J., and Arbib, M. A., 2000, Visual Tracking in simulated salamander locomotion, in From Animals to Animats 6, Proceedings of the Sixth International Conference of the Society for Adaptive Behavior (SAB2000), (J.-A. Meyer, A. Berthoz, D. Floreano, H. Roitblat, and S.W. Wilson, Eds), Cambridge MA: MIT Press, pp.88-97.
http://rana.usc.edu:8376/~ijspeert/PAPERS/SAB2000.pdf

* Simmons, P. and Young, D., 1999. Nerve Cells and Animal Behaviour. Cambridge UK: Cambridge University Press.

Webb, B., 2001, Can robots make good models of biological behavior? The Behavioral and Brain Sciences 24(6) in press. http://www.bbsonline.org/Preprints/Webb/