A Biological Perspective on Autonomous Agent Design

Randall D. Beer [1,2], Hillel J. Chiel [2], and Leon S. Sterling [1]

Departments of [1] Computer Engineering and Science, and [2] Biology and the Center for Automation and Intelligent Systems Research, Case Western Reserve University, Cleveland, OH 44106, USA

The inability of current "classical" AI systems to handle unconstrained interaction with the real world has recently lead to a search for new control architectures for autonomous agents. We argue that simpler natural animals already exhibit most of the properties required by an autonomous agent, and suggest that designers of autonomous agents should draw directly upon the neural basis of behavior in these animals. The relevant behavioral and neurobiological literature is briefly reviewed. An artificial nervous system for controlling the behavior of a simulated insect is then developed. The design of this artificial insect is based in part upon specific behaviors and neural circuits from several natural animals. The insect exhibits a number of characteristics which are remarkably reminiscent of natural animal behavior.

Keywords: Autonomous behavior; Artificial insect; Artificial nervous system; Heterogeneous neural networks; Computational neuroethology; Behavioral hierarchy; Motivated behavior; Insect locomotion; Artificial intelligence; Situated action.

Randall Beer received a B.S. in Computer Engineering (1985), and an M.S. (1985) and Ph.D. (1989) in Computer Science from Case Western Reserve University. He is currently an Assistant Professor in the Departments of Computer Engineering and Science, and Biology at Case Western Reserve University. He is interested in computational neuroscience, artificial intelligence, and symbolic programming languages.

North-Holland
Robotics and Autonomous Systems 6 (1990) 169–186

1. Introduction

The real world is complex, unpredictable, and dynamic. It is simply not possible for a designer to foresee all of the circumstances that might be faced by an agent in continuous, long-term interaction with such an environment. Any truly intelligent agent must therefore possess a considerable degree of autonomy. It must be capable of flexibly adapting its behavioral repertoire to the moment to moment contingencies which arise without explicitly being told what to do in each situation. How should the control architecture of such an autonomous agent be organized?

The classical AI answer to this question is largely drawn from introspection on conscious

Dr. Hillel Chiel graduated from Yale University in 1974 with a B.A. in English, and received a Ph.D. from M.I.T. in 1980 in Neural and Endocrine Regulation. As a postdoctoral fellow at the Center for Neurobiology and Behavior at Columbia University's College of Physicians and Surgeons, he studied the cellular basis of feeding behavior in the marine mollusc *Aplysia*. He then became a Consultant in Neurobiology in the Department of Molecular Biophysics Research at AT &T Bell Laboratories. He is currently an Assistant Professor in the Departments of Biology and Neuroscience at Case Western Reserve University. He is interested in experimental and modeling studies of behavior in simpler organisms.

Leon Sterling is Associate Professor in the Department of Computer Engineering and Science at Case Western Reserve University (CWRU), and co-Director of the Center for Automation and Intelligent Systems Research there. He received his B.Sc.(hons.) from the University of Melbourne in 1976 and his Ph.D. in Pure Mathematics from Australian National University in 1981. He spent three years in the Artificial Intelligence Department at the University of Edinburgh, and one year at the Weizmann Institute before joining CWRU in 1985. His interests include logic programming, expert systems and new foundations for artificial intelligence.

human reasoning. Essentially the same process by which we deliberately reason through, say, an anagram is hypothesized to underlie all intelligent behavior. Abstractly, this technique can be formulated as the appropriate manipulation of symbolic representations of the situation. The problem of crossing a room, for example, can be formulated as a search over symbolic descriptions of the possible paths through all of the intervening obstacles in much the same way that solving an anagram can. The actual manipulations involved in the former problem may be very much more complex than for an anagram, and the structure of the symbolic representations may be very intricate, but the idea is essentially the same. This notion finds its strongest and most explicit expression in the Physical Symbol System Hypothesis, which states that formal symbol manipulation is both a necessary and sufficient mechanism for general intelligent behavior [27].

In many ways, this methodology has served AI well. Though various issues were more difficult than originally anticipated, numerous fragments of intelligent behavior have now been generated in this manner. From the point of view of research on autonomous agents, however, all of these systems currently suffer from several rather glaring deficiencies: (1) they are incapable of flexibly coping with contingencies not explicitly foreseen by their designers, (2) their performance is extremely sensitive to the representational choices made by their designers, and brittle in the face of inevitable small deviations of the real world from these abstractions, and (3) their time complexity scales very poorly with problem size, becoming intractable for even simple real world tasks.

Of course, current shortcomings of the classical AI methodology do not necessarily compromise its fundamental soundness. The exploration of techniques for addressing these limitations is currently an active area of AI research. However, in light of the continuing difficulties encountered by these efforts, we must at least face the possibility that much of our intelligent behavior, particularly that which involves taking action in the real world, is really not at all like conscious reasoning. For this reason, a number of researchers have begun to explore alternative architectures for the control of autonomous agents (e.g. [8,25,1,17]).

Historically, AI has almost exclusively attempted to emulate *human* behavior. Our own approach to designing architectures for autonomous agents is grounded in the recognition that human beings are not the only natural agents which exhibit interesting autonomous behavior. Given our current level of understanding, people may not even be the best examples to study at this time. When even the most mundane contingency arises in our everyday interactions with the real world, we may draw upon an incredibly diverse collection of cognitive skills and a lifetime's worth of accumulated knowledge to cope with it. But human beings are simply too complex to model whole, and very little is known about the mechanisms underlying human cognition.

Therefore, our work has focused on the behavior of simpler natural animals, such as insects. While such animals cannot play chess or prove theorems, they are capable of autonomously adapting their limited behavioral repertoires to the moment to moment contingencies of the real world in ways that no current AI system can match. In order to tap this very rich source of potential insights for autonomous agent control structures, we have undertaken a careful study and simulation of the biological mechanisms underlying the autonomous behavior of simpler natural animals [4]. We call this endeavor *Computational Neuroethology*, since Ethology is the study of the behavior of animals in their natural environments [24], and Neuroethology is the study of the neural mechanisms underlying this behavior [9].

This paper is organized as follows. The next section describes some important principles of animal behavior. A brief introduction to some of the neural mechanisms which are known or hypothesized to underlie this behavior is then provided in Section 3. With this background, Section 4 describes an initial exploration into biologically-inspired control architectures for autonomous agents which we call the Artificial Insect Project. By drawing upon specific behavioral principles and their underlying neural circuits from several natural animals, we have designed an artificial nervous system for controlling the behavior of a simulated insect. Finally, Section 5 discusses the advantages and disadvantages of this approach to designing autonomous agent control architectures, and suggests some directions for future research.

2. Concepts in Animal Behavior

Before embarking upon a detailed consideration of the underlying biological mechanisms, it is instructive to briefly consider the problem of natural autonomous behavior from an ethological perspective. Unlike most artificial systems, natural animals obviously thrive very well in the real world. What are the common behavioral principles exhibited by animals engaged in the everyday business of their existence? Which of these principles might be useful in artificial autonomous agents?

Perhaps the most important principle, readily apparent to even a casual observer, is that all animal behavior is *adaptive* in the following sense: as an animal confronts its environment, its behavior is continuously adjusted to meet the everchanging internal and external conditions of the interaction. For example, a feeding insect will suddenly turn and run if it is attacked by a predator [31]. In addition, a running insect will continuously alter its gait to compensate for changes in terrain and load, and can even adjust for amputations of one or more legs [15].

Broadly speaking, animal behavior can be divided into a number of major classes. Perhaps the simplest form of animal behavior is a *reflex*, in which some fast, stereotyped response is triggered by a particular class of environmental stimuli. The defining characteristic of a reflex is that the intensity and duration of the response is entirely governed by the intensity and duration of the stimulus [10]. Reflexes allow an animal to quickly adjust its behavior to sudden environmental changes. Reflexes are commonly employed for such things as postural control, withdrawal from painful stimuli, and the adaptation of gait to uneven terrain.

Taxes or orientation responses are another simple class of behavior [9]. These behaviors involve the orientation of an animal toward or away from some environmental agent, such as light, gravity, or chemical signals. For example, female crickets exhibit positive phonotaxis during courtship, that is they orient to the calling song of a male [26].

Fixed-action patterns are a somewhat more complex form of behavior [24]. A fixed-action pattern is an extended, largely stereotyped response to a sensory stimulus. The triggering stimulus for a fixed-action pattern is generally more complex and specific than for reflexes. The response usually involves a complex temporal sequence of component acts. While such a pattern may be triggered by the occurrence of a specific sensory stimulus, its intensity and duration is not particularly stimulus-governed. In fact, once a fixed-action pattern has been triggered, it will usually run to completion even if the triggering stimulus is removed. An example of a fixed-action pattern is an escape response, in which some distinguishing characteristic of an imminent predator attack triggers a sequence of evasive maneuvers on the part of the prey (e.g. cockroaches escaping from toads; [31]). The fixed-action patterns of individual animals can also be interrelated in intricate ways, as is demonstrated by the elaborate courtship rituals between the male and female members of many animal species (e.g. guppies; [3]).

Despite the ubiquity of such responses as reflexes, taxes and fixed-action patterns, animal behavior is by no means solely reactive. Factors internal to an animal can also play an important role in the initiation, maintenance, or modulation of a given behavior. The sign or intensity of reflexes, for example, can change depending upon internal factors. The threshold for triggering most fixed-action patterns similarly varies with internal state.

Behaviors which show no simple or rigid dependence on external stimuli, but are instead governed primarily by the internal state of the animal, are known as *motivated behaviors*. In these behaviors, an animal's propensity to exhibit a given behavior such as feeding depends not only upon the presence of the appropriate environmental stimuli (i.e. food), but also upon internal *motivational* variables (i.e. hunger). Motivated behaviors are typically characterized by (1) grouping and sequencing of component behavior in time, (2) goal-directedness: the sequence of component behaviors generated can only be understood by reference to some goal, (3) spontaneity: the behavior can occur in the complete absence of any eliciting stimuli, (4) changes in responsiveness: the modulatory effect of the motivational state varies depending upon its level of arousal or satiation, (5) persistence: the behavior can greatly outlast any initiating stimulus, and (6) associative learning [21].

Any individual animal consists of a large collection of reflexes, taxes, and fixed-action patterns, many aspects of which are under at least some motivational control. As an animal confronts its environment with this diverse behavioral repertoire, it must properly coordinate its many possible actions into coherent behavior directed toward its long-term survival. Toward this end, the behavioral repertoire of a natural animal typically exhibits a certain organization. Some behaviors normally take precedence over others. Some behaviors are mutually exclusionary (i.e. any behaviors which utilize the same motor apparatus for incompatible actions). Switches between different behaviors depend both upon environmental conditions and internal state. These relationships are often described as rigid and strictly hierarchical, with cleanly delineated behaviors and simple all or nothing switching between them. In reality, the relationships may be nonhierarchical, the organization can change depending upon the behavioral context, and behaviors can partially overlap so that discrete switches between them are sometimes difficult to identify.

Though the number and variety of behavior clearly varies from species to species, all of the principles described above are exhibited in one form or another by all natural animals. This basic organization of behavior supports the ability of natural animals to flexibly cope with real world environments. In addition to this propensity for adaptive behavior, however, natural animals also exhibit various forms of *plasticity*. Aspects of their future behavior can be modified as a result of their past history of interactions with the environment. The time scale of these modifications may range from seconds to years.

Several simple forms of plasticity have been identified in natural animals [18]. In *habituation*, the magnitude of response to a given stimulus decreases with repeated exposure to the stimulus. For example, while a loud clap may initially produce a startle response in an animal, subsequent claps will produce a progressively weaker response. In some cases, the startle response may disappear altogether. *Dishabituation* is the sudden restoration of an habituated response following a particularly strong or noxious stimulus to the habituated sensory apparatus. An extremely loud clap, for example, might restore the habituated startle response. *Sensitization* involves an enhancement of a response to a wide variety of stimuli following the presentation of another strong stimulus. For example, a strong pinch might increase the sensitivity of the startle response to sound. These simple forms of plasticity allow an animal to adjust its responsiveness to its environment.

None of the above forms of plasticity depend upon a pairing of the strong stimulus with the weaker one. In *associative learning*, on the other hand, pairing between two stimuli is crucial. In one form of associative learning, called *classical conditioning*, repeated pairing of an initially neutral stimulus with one which normally elicits some response will eventually lead to a situation in which the neutral stimulus alone triggers the response. A common example of classical conditioning is when dogs salivate at the sound of a bell if the bell has been paired with the appearance of food in the past. In another form of associative learning, called *instrumental conditioning*, an animal's behavior is reinforced by events in its environment. For example, a rat will learn to avoid a particular food if prior ingestion of that food was followed by sickness [13]. These associational forms of plasticity allow an animal to take into account the causal relationships within its particular environment. However, it is important to realize that most animals cannot make arbitrary associations, but only those that are biologically relevant. For example, though a rat can easily learn to associate illness with a particular odor or taste, it by and large cannot learn to associate illness with auditory or visual stimuli.

Though there are several other forms of behavioral plasticity, we will mention only one more here. *Latent learning* is plasticity which does not involve particularly strong stimuli or obvious reward or punishment, as when an animal learns about its environment through exploration. For example, even ants can learn to run a maze simply by repeatedly being placed within it [33]. By these and many other forms of behavioral plasticity, animals fine-tune the behavioral repertoire with which they are genetically endowed to the exigencies of the particular environment in which they find themselves.

In this section, we have focused primarily on characterizing the behavior of simpler animals. This emphasis should not be misunderstood. Human beings are obviously not insects, and there are many aspects of human behavior of interest to

AI which clearly cannot be directly addressed through a study of simpler animals. We maintain, however, that there are many more which can. In particular, we strongly believe that the behavior of simpler animals has all of the ingredients which artificial autonomous agents require in order to flexibly cope with the real world: it is goal-oriented, adaptive, opportunistic, plastic, and robust. While the specifics of any given animal behavior are unlikely to be of direct use to an engineered agent, the general principles most certainly are.

Furthermore, it is important to stress that simpler animals are not simple. Even *C. elegans*, a millimeter long worm with only 302 nerve cells (it has less than 1000 cells in its entire body!) has been shown to be capable of associative learning [20]. In addition, several species of insects are known to possess elaborate social structures and to employ complex forms of communication [42]. Finally, and perhaps most importantly from our perspective, the neurobiological mechanisms underlying many of the abovementioned behavioral principles are beginning to be worked out in simpler animals. It is to these neural mechanisms which we now turn.

3. Neurobiological Basis of Animal Behavior

Consider the following problem: You must design the control system for a device which can autonomously accomplish some open-ended task (such as "stay out of trouble" or "keep this area clean") in a complex, dynamic, unpredictable, and, in many ways, openly hostile environment. You have considerable general information about the structure of this environment, but cannot assume that this information is complete in any sense. Your system must therefore be capable of flexibly applying whatever behavioral repertoire you choose to give it to the actual situations it encounters. At the same time, it must be capable of modifying aspects of that repertoire to better fit the particular environment in which it finds itself. This task is far easier than the one that evolution faces, because evolution cannot benefit from the knowledge of any conscious designer. The only information that it has about the environment is whether or not a given design succeeds in reproducing itself. On the other hand, because it has so little information to go on, its designs make the fewest possible assumptions, resulting in the most robust control systems in existence. Evolution's answer to this challenge is nervous systems.

One of the most important facts about nervous systems is that they are extremely heterogeneous. Individual nerve cells possess complex intrinsic dynamics which endow them with often unique response properties. The activity of a nerve cell at any point in time is a function not only of the activity of other nerve cells which synapse upon it, but also of its shape, the characteristics and distribution of its current channels, its chemical environment, and its internal biochemical state. Far from being unimportant biological details, most of these properties appear to be functional: nervous systems actually take advantage of them for controlling the behavior of animals [34,23].

Nervous systems are not only heterogeneous in their elements, but also in the interconnections between those elements. Nervous systems consist of a great many specific circuits which are organized into highly structured architectures. These architectures are constructed during the development of an animal, and have been designed over the course of evolution. They are responsible for the basic complement of behavior with which an animal is endowed [36].

How are nervous systems organized to support the behavioral principles discussed in the previous section? By a careful analysis of both the behavior and the underlying neural circuitry, this question is beginning to be answered in simpler animals. Underlying reflexes in all animals, for example, are essentially direct connections between the sensory neurons which recognize the sensory stimulus and the motor neurons responsible for the response. Because these reflex circuits typically consist of short, fast pathways involving no more than one or two synapses, they are capable of only rudimentary sensory analysis and stereotyped motor responses. However, these reflex circuits may be affected by other circuits which interact with the neurons involved in the reflex [38]. Similarly, taxes and other orientation responses appear to be controlled by circuits which compare information from sensory receptors on each side of an animal's body [32].

All behaviors more complex than reflexes and taxes require the generation of temporally extended patterns of motor activity (e.g. fixed-action patterns). What is the neural basis of such patterns of behavior? How are the sequencing and timing of the individual components controlled? These questions have been most fully addressed in the context of rhythmic patterns of behavior, such as swimming or walking. The neural circuits underlying rhythmic behaviors are called *central pattern generators* [12]. They can be divided into two general categories: those employing pacemaker cells and those employing network oscillators. Pacemaker cells are neurons which are capable of producing rhythmic bursts solely by virtue of their own intrinsic dynamics. Network oscillators, on the other hand, are networks of neurons which generate rhythmic patterns due to the synaptic interactions between their component neurons, none of which are capable of rhythmic activity in isolation. Often, central pattern generators involve networks of neurons with intrinsic bursting properties, so that the final pattern depends both upon the intrinsic dynamics of each cell as well as the interconnections between them (e.g. [34]). In addition, the details and phasing of the basic pattern produced by a central pattern generator can be greatly affected by sensory feedback, sometimes making it difficult to distinguish between a *central* pattern generator and a *peripheral* one [29].

What are the neural mechanisms by which the internal state of an animal affects its behavior? There is no simple answer to this question. Many internal processes exist which can effect the function of particular neural circuits over time: (1) as already described, individual nerve cells have intrinsic cellular dynamics which influence their operation; (2) reverberating pathways exist in which any activity in a circuit leads to increased activity within that same circuit via positive feedback loops; (3) the activity of one nerve cell can dynamically alter the interactions between several others because neurons form connections on different parts of other neurons; and (4) the activity of a given neural circuit can be greatly influenced by a variety of chemical means, such as hormones [35].

Many situations require that a decision be made whether or not to generate a specific behavior (e.g. some fixed-action pattern) in a given context. An early notion regarding the neural basis of such decisions was that of a *command neuron* [22]. In this view, sensory information converges on a single neuron which initiates the response by activating the appropriate pattern generation circuitry only when the proper sensory stimulus is present. A few examples of putative command neurons have been found. However, though neurons whose activation can elicit specific motor patterns certainly exist, the notion of a single neuron being uniquely responsible for a given behavior has proven to be an oversimplification. Even for simpler animals, it now appears that behavioral choice is a much more distributed process. In general, decisions regarding which behaviors to generate in a given environmental context appear to be achieved by consensus involving interactions between the circuits responsible for each behavioral pattern [2].

The cellular basis of several forms of behavioral plasticity have been extensively studied in some invertebrates. For example, in the marine mollusc *Aplysia*, studies of habituation have shown that the decrement in the animal's response to a repetitive stimulus is associated with synaptic depression: due to the nature of the biochemical processes responsible for synaptic transmission, the efficacy of a specific synapse between the sensory neuron and the motor neuron involved in the response decreases with repeated use. Similarly, sensitization is caused by an enhancement of synaptic transmission within the affected pathway. This enhancement is triggered by another neuron associated with the sensitizing stimulus which forms a synapse near the affected synapse and releases chemicals which increase its efficacy. Finally, associative conditioning has similarities to sensitization in that the efficacy of a neural pathway is enhanced. In *Aplysia*, it appears that the prior activity of the neuron that receives reinforcement allows the affected synapse to be further enhanced by the mechanisms previously described for sensitization [19].

As a solution to the problem posed at the beginning of this section, nervous systems have many advantages. The flexibility and complexity of the individual neurons allow them to be configured in a large variety of ways. Because of their complexity, individual nerve cells can be utilized to process large amounts of information in parallel. The architectures that are typical of nervous systems are also highly distributed, with different parts performing overlapping, but not identical functions. Finally, the intrinsic properties of neurons, and their complex interactions, endow them with a rich dynamics that may be capable of responding much more effectively to the rapidly changing exigencies of the real world than more discrete, centralized systems.

4. The Artificial Insect Project

In the preceding sections, we have argued that artificial autonomous agents should aspire to the behavioral capabilities of simpler natural animals. Given the current state of the art, a robot with "only" the behavioral repertoire of an insect would, we believe, be quite an impressive achievement. We have further argued that the control architectures we design for our autonomous agents could benefit from a knowledge of the neural mechanisms underlying natural animal behavior. To test these ideas, we have undertaken the construction of a simulated insect whose behavior is controlled by an artificial nervous system. The overall design of this insect is inspired by the principles of natural animal behavior described in the previous two sections, and its nervous system is based in part on specific neural circuits in several natural animals. At present, the simulated insect is capable of locomotion, wandering, edge-following, and feeding, as well as properly managing the interactions between these behaviors in order to survive within its environment for an extended period of time. In this section, we provide an overview of this *Artificial Insect Project*, which represents a first cut at designing a complete, biologically-inspired artificial agent. Full details can be found in [4].

4.1. Physical Models

As for a natural animal, the physical characteristics of an autonomous agent's body and environment have a significant impact on the design of its controller. The body model we have chosen for the artificial insect is shown in *Fig. 1*. Though this design is loosely based on the American Cockroach [7], it resembles the basic body plan of many insects. The antennae contain tactile and chemical sensors. The mouth can open and close, and also contains tactile and chemical sensors. The insect has an internal store of energy, as well as a simple metabolism in which energy is consumed at a fixed rate. If its energy level ever reaches zero, the insect is removed from its environment. When the insect's mouth closes over a patch of food, a fixed amount of energy is transferred from the food patch to the insect's internal energy store.

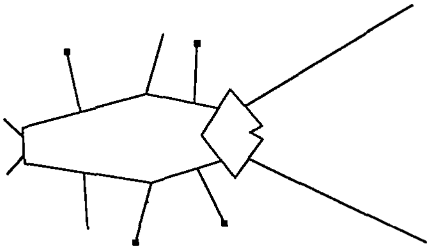

Fig. 1. Body model.

The artificial insect has six legs, each with a foot that may be either up or down. When its foot is up, a leg assumes a fixed length and any forces it applies cause it to swing. When its foot is down (denoted by a black square), a leg stretches between its foot and the body, and any forces it generates may result in movement of the body. Despite the fact that the insect is only two-dimensional, it can fall down. The insect becomes statically unstable whenever its center of mass lies outside of the polygon formed by the feet which are down. If this condition persists for longer than 40 msec, the insect is considered to have fallen down and the legs are no longer allowed to move the body.

The environment in which the artificial insect exists contains unmovable obstacles and food patches. When an insect encounters an obstacle, it bounces back along its direction of motion a small, fixed amount. Food patches are circular areas of the environment which contain energy. These patches emit an odor whose strength is proportional to the number of food units in the patch, which is in turn proportional to its area. As odors diffuse through the environment, their intensity falls off as the inverse square of the distance from the center of the food patch.

4.2. Neural Model

In order to utilize neurobiological principles for controlling the behavior of an artificial agent, we must choose a neural model which strikes the proper balance between biological reality and computational and conceptual tractability. We cannot possibly model an entire nervous system at the detailed biophysical level. On the other hand,

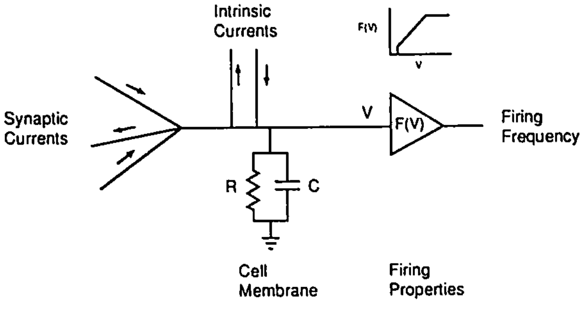

Fig. 2. Neural model.

as discussed in Section 3, certain biological characteristics appear to be fundamental to the way nervous systems control behavior. Our neural model is therefore intermediate in complexity between biological nerve cells and the formal neurons typically employed in artificial neural networks. The model is shown schematically in *Fig. 2*.

The output of a model neuron corresponds to the firing frequency of a nerve cell. In the model, this frequency is a nonlinear function of the neuron's potential. We have employed saturating linear threshold functions with an initial jump discontinuity to represent this relationship (see inset). Three parameters characterize this function: the threshold voltage at which the neuron begins to fire, the minimum firing frequency, and the gain. An RC circuit is used to capture the ability of nerve cells to temporally sum their inputs. Model neurons are interconnected by weighted synapses through which they can inject current into one another. These aspects of the model are similar, though not identical, to several neural models that have been previously explored in the field of artificial neural networks (e.g. [16]).

One of the most striking differences between real nerve cells and the formal neurons that are typically employed in artificial neural network research is their rich internal dynamics. Nerve cells are not simple functions, but dynamical systems which are capable of spontaneous activity and whose input/output characteristics change over time. Without modeling the detailed biophysical mechanisms responsible for these characteristics, we have nevertheless captured their net effect through the addition of intrinsic currents to our model. These currents may be both time and voltage dependent. For example, we have used a pair of intrinsic currents to design a model pacemaker neuron which is employed in several circuits within the artificial insect's nervous system.

Finally, our neural model supports compound synapses, in which the output of one neuron effects the connection between two others. Compound synapses come in two varieties. *Gating* synapses allow the activity of one neuron to enable or disable connections between two others. *Modulatory* synapses, on the other hand, allow a neuron to modify the strength of a connection between two others in a multiplicative fashion.

4.3. Locomotion

All of the other behaviors require some means for the artificial insect to traverse its environment. Therefore, the first behavior we sought to implement was locomotion [6]. In locomotion, each leg must swing rhythmically. However, because the insect can fall down, the controller must also properly coordinate the movements of the six individual legs in order to achieve successful locomotion.

The design of the neural circuit which controls locomotion in the artificial insect is largely based on the work of Pearson and his colleagues on the neural basis of locomotion in the American cockroach [28]. While a complete circuit has not yet been worked out, several principles of its operation have been identified [30]. (1) each leg is probably controlled by a separate central pattern generator, (2) reflexes involving leg position and load play an important role in shaping the output of the central pattern generators, and (3) the central pattern generators controlling different pairs of legs are probably coupled by some form of inhibition.

During walking, each leg rhythmically alternates between a swing phase and a stance phase. During the *swing phase*, the foot is up and the leg is swinging forward. During the *stance phase*, the foot is down and the leg is swinging back, propelling the body forward. These basic movements are produced by the pattern generator circuit shown in *Fig. 3*. There are six copies of this circuit, one for each leg, except that a single command neuron LC makes the same two connections on all six leg controllers. The design of this circuit is based upon Pearson's Flexor Burst-Generator Model [28].

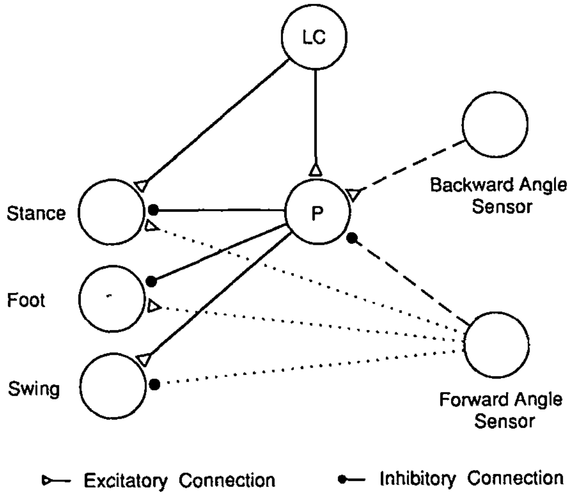

Fig. 3. Leg controller.

The alternating swing and stance movements of each leg are primarily produced by the central pattern generator shown in sold lines in *Fig. 3*. Each leg is controlled by three motor neurons. The *swing* and *stance* motor neurons determine the force with which the leg is swung forward or backward, respectively, while the *foot* motor neuron controls whether the foot is up or down. A stance phase is produced when the foot is down and the stance motor neuron is active. Periodically, a swing phase is initiated by a burst of activity produced by the pacemaker neuron P. This activity lifts the foot and swings the leg forward by inhibiting the foot and stance motor neurons and exciting the swing motor neuron. Another stance phase begins when the pacemaker burst terminates. The alternating swing/stance cycle required for walking can thus be produced by rhythmic bursting in P. Note that both the force applied by the leg during each stance phase and the time between bursts in P depend upon the steady level of excitation supplied by the locomotion command neuron LC.

Most central pattern generators require some sensory feedback to fine-tune the basic pattern. In our controller, this information is supplied by two sensors which signal when a leg has reached an extreme forward or backward angle (shown with dashed lines in *Fig. 3*). When a leg is all the way back, the *backward angle sensor* encourages it to swing by exciting the pacemaker. The *forward angle sensor*, on the other hand, encourages the leg to terminate the current swing by inhibiting the pacemaker. In addition, the direct connections from the forward angle sensor to the motor neurons (shown with dotted lines in *Fig. 3*), comprise a stance reflex which smooths the transition from swing to stance.

In order to generate statically stable gaits, the movements of each individual leg must be properly coordinated or the insect will fall down. This coordination is achieved by appropriate coupling of the central pattern generators controlling each leg. One useful rule of thumb is that adjacent legs should not swing at the same time. This constraint is implemented by mutual inhibitory connections between the pacemakers of adjacent legs, as shown in *Fig. 4*. For example, when the middle right leg is swinging, the front and back right legs and the middle left leg are discouraged from also swinging, but the other legs are unaffected.

While these constraints generate statically stable gaits at high speeds of walking, at lower speeds they are not sufficient to guarantee statically stable gaits. The slower gaits of many animals exhibit a stepping sequence known as the *metachronal wave*, in which a wave of swings progresses from the rear of the animal to the front. In insects, for example, the back leg swings, then the middle leg, then the front leg on each side of the body. This appears to be a particularly stable pattern of stepping. Metachronal waves were implemented in our model by increasing the leg angles of the two rear legs, which lowers the natural frequency of their pattern generators due to the sensory feedback [14]. Because of the inhibitory coupling between the pacemakers of adjacent legs, this results in entrainment between the pacemakers on each side of the body. In the stable phase relationship

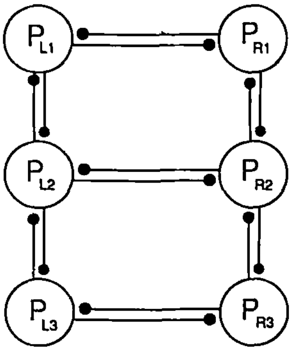

Fig. 4. Central coupling between pacemakers.

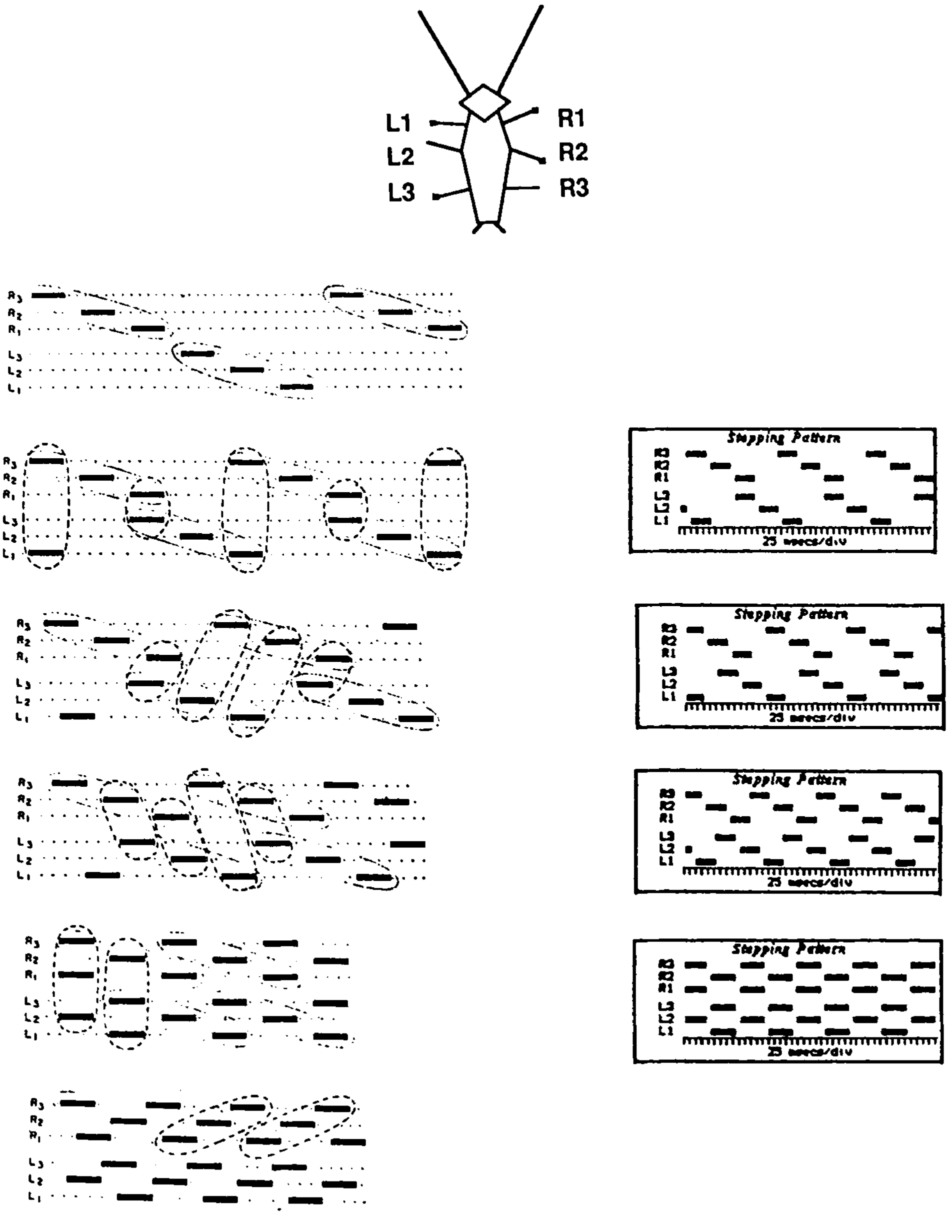

Fig. 5. A comparison of the gaits generated by the artificial insect (right) and natural insects (left; [41]). Leg labeling conventions are shown at top.

results from this entrainment, the swing of any given leg immediately follows the one behind it.

Using this locomotion controller, the insect exhibits a continuum of statically stable gaits as the firing frequency of the locomotion command neuron LC is varied. Gaits can be conveniently described by their *footfall patterns*. In this representation, a black bar is displayed during the swing phase of each leg. The space between bars represents the stance phase. Selected gaits exhibited by the artificial insect are shown at the right in *Fig. 5* as the firing frequency of LC is varied from lowest

(Top) to highest (Bottom). At low levels of LC activity, the metachronal waves on each side of the body are very apparent. However, they increasingly overlap as LC activity increases until the *tripod gait* appears at the fastest walking speed. In this gait, the front and back legs on each side of the insect swing and stance together with the middle leg on the opposite side. This sequence of gaits bears a striking resemblance to those that have been described by Wilson [41] for natural insects (*Fig. 5; Left*).

This sequence of gaits emerges from the interaction between the dynamics of the neural circuitry responsible for locomotion and the body and environment in which it is embedded. In order to better understand the operation of this neural network, we undertook a series of lesion studies, in which the response of the controller to the removal of various elements was examined [11]. We briefly summarize some of our results here.

In general, we found the locomotion controller to be remarkably robust to removal of any single element or connection. The lesion of one component often led to compensatory effects in other components, although the overall robustness of the controller to further perturbations always decreased with any lesion. Indeed, we found the interaction between the central and peripheral components of the controller to be quite unexpectedly subtle. For example, lesioning connections from the backward angle sensors to the pacemakers in the rear legs completely abolished the metachronal wave in slower speed gaits, but left the tripod gait virtually unaffected. In contrast, higher speed gaits were more sensitive to lesions of the central inhibitory connections between pacemakers than were lower speed gaits. This suggests that the higher speed gaits are primarily generated centrally, while lower speed gaits are more dependent upon sensory information.

While the command neuron would appear to be important to the generation of all gaits, it is not, in fact, essential. For example, the full range of normal gaits were exhibited even after complete removal of LC if the insect was pushed along by an external force. In addition, removing all of the connections from LC to the six pacemakers had no effect whatsoever on a walking insect.

When these same six connections were removed before an insect was allowed to establish a normal gait, the reason for these remarkable results became clear. A normal insect usually establishes a stable gait within a single step. However, an insect with this lesion required well over a dozen steps to achieve the proper coordination. On the other hand, once a normal gait was established, the lesioned insect was indistinguishable from a normal one. It appears that LC normally acts to set the burst frequency of the pacemakers close to the final value required for a given speed of walking, and then the sensory information simply fine-tunes them. However, in the complete absence of this central information, the sensory feedback alone is still sufficient to establish normal gaits, though this process is cruder and takes longer than normal.

4.4. Wandering

Once an insect is capable of locomotion, it can begin to wander through its environment if it also has an ability to turn. In straight-line locomotion, the legs apply forces which translate the insect's body. Turning was implemented by also allowing the legs to apply lateral forces to the body, thereby rotating it. Since the neural circuitry responsible for wandering behavior in insects has not yet been worked out, we designed a simple neural network which is capable of generating the necessary motor patterns [4]. In this controller, two pacemaker-like neurons whose burst and interburst characteristics vary randomly are used to excite the motor neurons controlling the lateral extension of the front legs at random intervals and for random periods of time.

4.5. Edge-Following

An animal must have some means for coping with any obstacles it encounters as it wanders through its environment. One strategy that is commonly employed by insects is edge-following ([7], p. 373). During edge-following, an insect maintains a nearly parallel orientation between its body and the edge of the obstacle it is following. If the angle between the insect and the edge is too small, the insect must turn toward the edge so as to increase this angle. If this angle is too large, the insect must decrease it by turning away from the edge. In addition, a momentary loss of contact with the edge should not terminate the behavior.

Of course, if the insect is unable to reestablish contact within a certain period of time, then its attempts to do so should cease. Thus, edge-following exhibits behavioral hysteresis: once triggered, it persists for a short period of time even after the sensory stimulus which initially triggered it has been removed. Unfortunately, the neural circuitry controlling edge-following behavior in insects is currently unknown. We therefore designed a neural network which is capable of generating edge-following behavior with the characteristics described above [4]. There are two copies of this controller, one for each antenna.

4.6. Feeding

As previously discussed, motivated behaviors are among the most interesting and complex behaviors exhibited by simpler animals. Feeding is a prototypical motivated behavior in which attainment of the goal object (food) is clearly crucial to an animal's survival. In this case, the relevant motivational state is hunger. When an animal is hungry, it will exhibit a sequence of *appetitive* behaviors which serve to identify and properly orient the animal to food. Once food is found, *consummatory* behaviors are generated to ingest it. On the other hand, a satiated animal may ignore, or even avoid, sensory stimuli which suggest the presence of food [21].

An animal's interest in feeding (its *feeding arousal*) may be a function of more than just its energy requirements. Other factors, such as the exposure of an animal to the taste, odor, or tactile sensations of food, can significantly increase its feeding arousal. This relationship between feeding and arousal, in which the very act of feeding further enhances an animal's interest in feeding, leads to a form of behavioral hysteresis. Once food is encountered, an animal may feed well beyond the internal energy requirements which initiated the behavior. In many animals, this hysteresis is thought to play a role in the patterning of feeding behavior into discrete meals rather than continuous grazing [37]. At some point, of course, the ingested food must be capable of overriding the arousing effects of consummation, or the animal would feed indefinitely.

Because the artificial insect possesses a simple metabolism and a limited energy store, it too requires some form of feeding behavior in order to survive for any extended period of time. Based in part upon neuroethological data on the feeding behavior of the marine mollusc *Aplysia*, we have designed neural controllers for feeding in this insect. An appetitive controller is responsible for finding food when the insect is in need of energy by following the odor which diffuses from a food patch. Once a food patch has been found, a separate consummatory controller is responsible for the actual ingestion. This controller also implements the arousal and satiation characteristics described above.

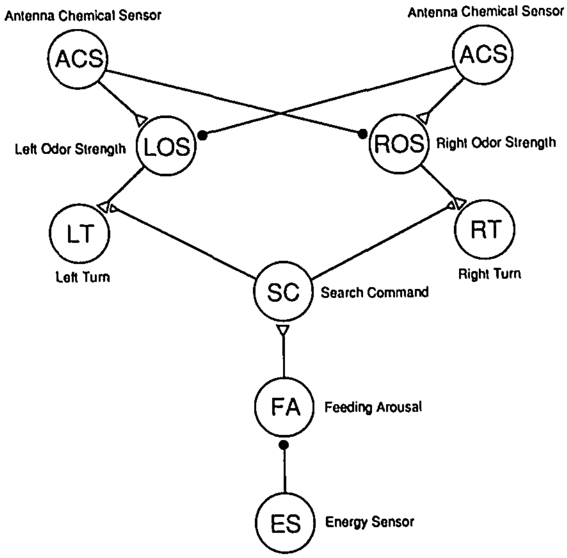

Fig. 6. Appetitive controller.

The appetitive component of feeding behavior in the artificial insect is an example of a taxis. The appetitive controller is shown in *Fig. 6*. Its design follows the general outlines of several proposed neural circuits controlling taxes in various animals, but it is not directly based upon any specific circuit. This controller consists of two components. The first is responsible for orienting the insect to a food patch by following its odor. These odor signals detected by the chemical sensors in each antenna (ACS) are compared (by LOS and ROS) and the difference between them is used to generate a turn toward the stronger side by exciting the corresponding turn interneuron (LT or RT) by an amount proportional to the strength of the odor gradient. These turn interneurons connect to motor neurons controlling the lateral extension of the front legs.

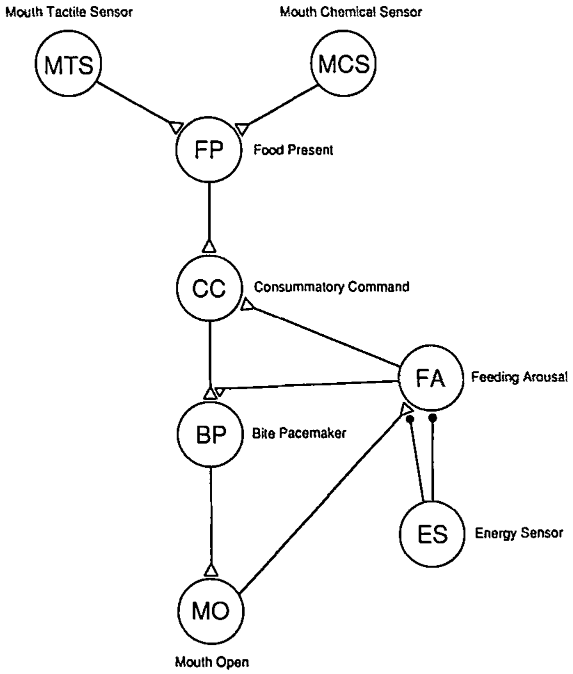

Fig. 7. Consummatory controller.

The second component of the appetitive controller is responsible for controlling when the insect orients to food. Though the odor gradient is continuously being sensed, the connections from the odor strength neurons (LOS and ROS) to the turn neurons (LT and RT) are normally disabled, preventing this information from actually turning the insect. As the insect's energy level falls, however, so does the activity of its energy sensor (ES). This gradually releases the spontaneously active feeding arousal neuron (FA) from inhibition. When the insect becomes sufficiently aroused to fire the search command neuron (SC), the connections between the odor strength neurons and the turn neurons are enabled by gating synapses from SC, and the insect begins to orient to food.

The consummatory component of feeding behavior is a fixed-action pattern which is triggered by the presence of food. The consummatory controller is shown in *Fig. 7*. When chemical (MTS) and tactile (MCS) sensors in the mouth signal that food is present and the insect is sufficiently aroused to feeding, the consummatory command neuron CC fires. Both tactile and chemical signals are required to prevent attempts to ingest nonfood patches and to prevent consummation from beginning before the food is actually reached (due to the diffusion of odors). Once CC fires, it triggers the bite pacemaker neuron (BP) to generate the rhythmic bursts necessary for driving the motor neuron (MO) which opens the mouth. Because the threshold of the consummatory command neuron (CC) is lower than that of the search command neuron (SC), an insect which is not sufficiently aroused to actively search for food may nevertheless consume food that is presented directly to its mouth.

The motor neuron controlling the mouth also makes an excitatory connection onto the feeding arousal neuron, which in turn makes an excitatory modulatory connection onto the connection between the command neuron and the bite pacemaker. The net effect of these excitatory connections is a positive feedback loop: biting movements excite FA, which causes BP to burst more frequently, thereby generating more frequent biting movements which further excite FA until its firing frequency saturates. This positive feedback loop is inspired by work on the neural basis of feeding arousal maintenance in *Aplysia* [39].

As the insect consumes food, its energy level begins to rise. This increasing activity both inhibits FA directly, and decreases the gain of the positive feedback loop via an inhibitory modulatory synapse onto the connection between MO and FA. At some point, this effect will overcome the positive feedback and activity in FA will cease, causing the insect to stop feeding. This neural mechanism is based upon a similar one hypothesized to underlie the satiation of feeding in *Aplysia* [40]. Thus, complex interactions between the energy sensor (ES) and the feeding arousal neuron (FA) implement the motivational state governing feeding behavior in this insect.

With these two neural controllers in place, the feeding behavior of the artificial insect exhibits four of the six characteristics of motivated behavior defined by Kupfermann (1974) and presented in Section 2:

Grouping and sequencing of behavior in time. A "hungry" artificial insect generates appetitive and consummatory behaviors with the proper sequence, timing, and intensity in order to obtain food.

Goal-Directedness. Regardless of its environmental situation, a hungry insect will generate movements which serve to obtain food. Thus, at certain times, the behavior of the artificial insect can be understood only by reference to an internal goal.

Changes in responsiveness due to a change in internal state. While a hungry insect will attempt to orient to and consume any nearby food, a satiated one will ignore it. In addition, once a hungry insect has consumed sufficient food, it will walk right over the food patch which initially attracted it.

Persistence. If the artificial insect is removed from food before it has fed to satiation, its feeding arousal will persist for some time.

One technique that has been applied to the study of arousal and satiation in natural animals is to examine the time interval between successive bites as an animal feeds under various conditions. In *Aplysia*, for example, the interbite interval progressively decreases as an animal begins to feed (showing a buildup of arousal) and increases as the animal satiates. In addition, the rate of rise and fall of arousal depends upon the initial degree of satiation [37].

It is interesting to compare these results to those obtained from similar experiments on the artificial insect. Food was directly presented to artificial insects with differing degrees of satiation, and the time interval between successive bites was recorded for the entire resulting consummatory response. Above approximately 80% satiation, insects could not be induced to bite. Below this level, however, insects began to consume the food. As these insects fed, the interbite interval decreased until some minimum was reached as feeding arousal built up (*Fig. 8*). The rate of arousal build-up was slowest for insects with the highest initial level of satiation. In fact, an insect which was already 75% satiated never achieved full arousal.

As the feeding insects neared satiation, the interbite interval again increased as arousal waned. It is interesting to note that, regardless of the initial level of satiation, all insects in which biting was triggered fed until their energy stores were approximately 99% full. The appropriate number of bites to accomplish this were generated in all cases. Feeding behavior in the artificial insect thus exhibits a number of very interesting characteristics which are quite reminiscent of natural animals. These issues are further explored in [5].

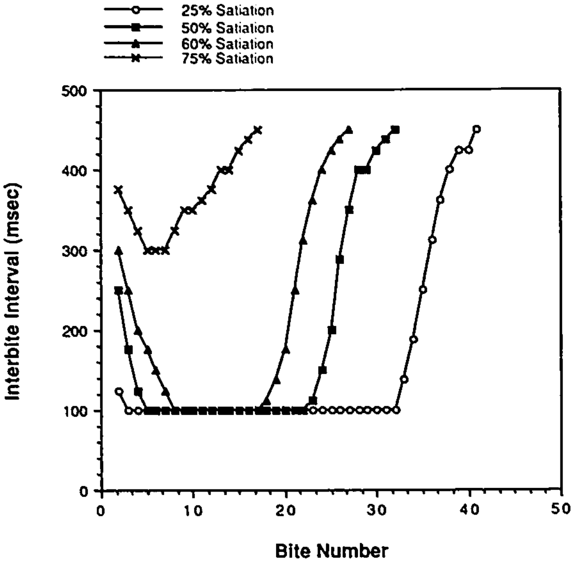

Fig. 8. Build-up of arousal and satiation.

4.7. Behavioral Choice

As described above, the artificial insect is capable of locomotion, wandering, edge-following, and feeding (which in turn consists of appetitive and consummatory component behaviors). Many of these behaviors are potentially incompatible because they share the same motor apparatus. For example, the wandering, edge-following, and appetitive controllers all utilize the lateral extensors of the front legs to turn the insect, often in opposite directions. In addition, while locomotion is crucial to these three behaviors, locomotion during the consummatory behavior would be disastrous. The artificial insect must therefore constantly decide what to do next given its current internal and external situation. How should its nervous system be organized so that the many individual neural controllers always generate globally coherent behavior?

Generally speaking, feeding should take precedence over edge-following, which in turn should take precedence over wandering. The artificial insect's behavioral repertoire can therefore be organized as shown in *Fig. 9*. Each major behavior is represented by an ellipse. Locomotion is not explicitly represented as a separate behavior, since it is implicitly utilized by most of the other behaviors. In addition, certain other important interactions, such as that between the edge-following

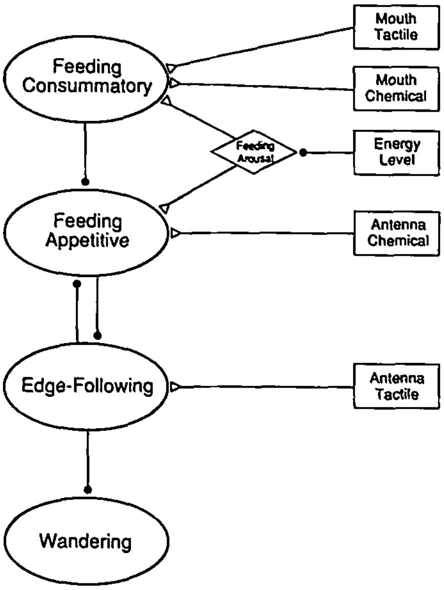

Fig. 9. Behavioral organization.

controllers on each side of the body (crucial in corners), are not explicitly represented in this diagram. Sensory stimuli which play a role in triggering a given behavior are shown as rectangular boxes, while the motivational state governing feeding is represented by a diamond. The interactions between these various components are illustrated by excitatory and inhibitory connections. This diagram roughly corresponds to that which an Ethologist might construct to describe the interrelationships between the various behaviors of a natural animal.

In general, whenever a higher order behavior is triggered, it suppresses lower order behaviors. Note, however, that the diagram in *Fig. 9* is not strictly hierarchical. While feeding normally takes precedence over edge-following, this precedence reverses if an obstacle blocks the insect's path to food. In this case, the insect follows the edge of the obstacle in the hopes of getting around it. The relationship between these two behaviors is therefore dependent upon the environmental context.

The excitatory and inhibitory connections in *Fig. 9* are meant only to illustrate the interactions between the artificial insect's various behaviors. How can the interactions in this diagram be implemented neurally? In some cases, the required interactions between two behaviors can be directly implemented by explicit connections between the corresponding command neurons. For example, edge-following behavior can suppress wandering via direct inhibitory connections from a key neuron in each of the edge-following controllers to a key neuron in the wandering controller.

The neural implementation of other interactions is more complex. The edge-following and appetitive controllers are particularly interesting in this regard. Neither of these controllers can simply suppress the other because situations exist in which either one should dominate. The neural implementation of this relationship therefore requires additional circuitry which modifies the interaction between these two controllers depending upon the environmental context. Thus, behavioral choice in the artificial insect is implemented in a distributed fashion: decisions are made by consensus among the various neural controllers rather than by a centralized decision module. A complete discussion of the neural circuitry which mediates these behavioral interactions can be found in [4].

The artificial insect's complete nervous system, which implements the behavioral repertoire illustrated in *Fig. 9*, contains a total of 78 model neurons and 156 model synapses. The capabilities of this nervous system are illustrated in *Fig. 10*, which shows the path followed by an artificial insect as it solves a simple but important problem in its environment. At (1), the insect is low on energy and immediately begins to locomote toward the food patch at the upper left (note that obstacles do not block the diffusion of odor). At (2), however, it collides with the intervening wall

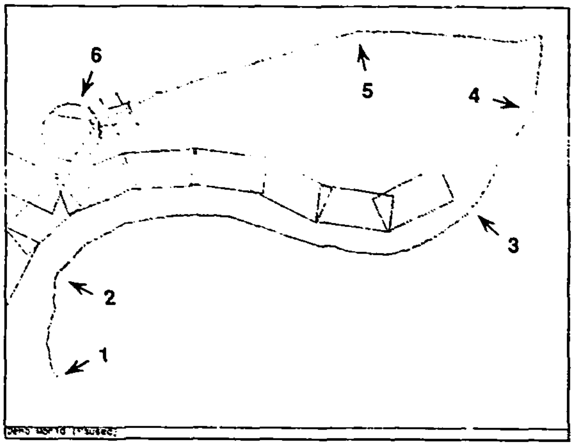

Fig. 10. An illustration of the artificial insect's behavioral repertoire.

and begins to follow its edge. When the insect loses contact with the wall at (3), it briefly tries to reestablish contact by turning back toward it. When no further contact is forthcoming, the insect begins to wander. Note that, due to the inverse square decay of odor intensity with distance, the insect's chemosensors can no longer detect the food patch at this point. After a short period of wandering, it collides with the right wall at (4) and begins to follow it, negotiating a corner in the process. As it continues to follow this edge, the insect once again comes within range of the odor at (5). It immediately leaves the wall it was following and heads toward the food patch, finally feeding successfully at (6).

5. Discussion

While it would certainly never be mistaken for a natural insect, the artificial insect described above nevertheless exhibits a number of characteristics which are strikingly reminiscent of the autonomous behavior of simpler natural animals. It is capable of locomotion, wandering, edge-following, and feeding. Its locomotion controller, which is directly based upon the neurobiological data for cockroach walking, can generate a continuum of statically stable gaits simply by varying the activity of a single neuron. These gaits are quite similar to those that have been described for natural insects. Lesion studies of this controller have demonstrated a remarkable robustness and subtlety of operation. The feeding behavior of the artificial insect similarly displays many of the characteristics associated with motivated behavior in natural animals, including a build-up of arousal and satiation as feeding progresses. Finally, the artificial insect is capable of flexibly organizing its behavioral repertoire in a variety of ways in order to survive within its simulated environment. We believe that this richness is a direct consequence of the biological details which we have incorporated into our model.

The artificial insect currently suffers from a number of limitations in its present form. While portions of its nervous system are based directly upon neurobiological data (e.g. locomotion and feeding), other portions are rather *ad hoc*. Even the biologically-inspired controllers had to be fine-tuned by trial and error (there are over 500 parameters in the insect's nervous system). We would have preferred a more principled approach to these issues, but the required neurobiological data was simply not available. Because we were interested in designing a complete autonomous agent rather than modeling only isolated pieces of behavior, we were forced to fill in many missing details. Only further study of natural nervous systems and considerably more design experience with artificial ones will increase the sophistication of our neural controller designs and deepen the principles upon which they are based.

The behavioral repertoire of the artificial insect is also currently rather limited. While it does exhibit, in one form or another, most of the behavioral characteristics described in Section 2, its repertoire is still quite impoverished compared to that of any natural animal. In addition to locomotion, wandering, edge-following, and feeding, insect behavior typically includes fleeing, fighting, nest building, foraging, grooming, mating, and communication. Some of these behaviors would certainly be useful to an artificial agent, and they are all interesting objects of study in their own right. Unfortunately, although there is a rich body of literature on the ethology of these behaviors, neural circuitry for many of them is not currently available.

Also conspicuously absent from the artificial insect is any form of plasticity. Plasticity is clearly crucial to an autonomous agent, and several neural mechanisms for it were mentioned in Section 3. However, we chose to focus first on nervous system design. Plasticity is a means by which evolutionarily good designs are fine-tuned to the particular environment of an individual animal. It is not a process for producing good designs in the first place from unstructured controllers. Now that we have designed a nervous system which is capable of endowing the artificial insect with the basic behavior essential to its survival, we can begin to explore the behavioral implications of introducing plasticity into specific portions of this nervous system.

In conclusion, we believe that the behavior of even simpler natural animals already exhibits most of the characteristics which we seek to instill in artificial autonomous agents. Animal behavior is goal-oriented, adaptive, opportunistic, plastic, and robust. All of these qualities are crucial for continuous, long-term interaction with the real world.

Furthermore, we feel that the current level of understanding of the neural basis of behavior in simpler animals is sufficiently mature that fruitful interactions between Neuroethology and AI are possible. The artificial insect that we have described represents only one example of such interaction. A great many others are possible.

Acknowledgements

We would like to thank Roy Ritzmann for his comments on an earlier draft of this paper. We are also grateful to Yoh-Han Pao and the Center for Automation and Intelligent Systems Research. RDB was supported in part by a grant from the Cleveland Advanced Manufacturing Program. HJC was supported by NSF grant BNS 88-10757.

References

[1] P. Agre and D. Chapman, Pengi: An implementation of a theory of activity. *Proceedings of the Sixth National Conference on AI* (AAAI 87) (Morgan Kaufmann Publishers, 1987) 268–272.

[2] J.S. Altman and J. Kien, New models for motor control, *Neural Computation* 1 (1989) 173–183.

[3] G.P. Baerends, R. Brouwer and H.T.J. Waterbolk, Ethological studies on *Lebistes reticulatus* (Peters), I. Analysis of the male courtship pattern, *Behaviour* 8 (1955) 249–334.

[4] R.D. Beer, *Intelligence as Adaptive Behavior: An Experiment in Computational Neuroethology*, Ph.D. Dissertation, Dept. of Computer Engineering and Science, Case Western Reserve University (1989). Also available as Technical Report TR 89-118, Center for Automation and Intelligent Systems Research.

[5] R.D. Beer, and H.J. Chiel, (forthcoming), Neural implementation of motivated behavior: Feeding in an artificial insect, to appear in D.S. Touretzky (Ed.), *Advances in Neural Information Processing Systems 2* (Morgan Kaufmann Publishers).

[6] R.D. Beer, H.J. Chiel and L.S. Sterling, Heterogeneous neural networks for adaptive behavior in dynamic environments, in D.S. Touretzky (Ed.), *Advances in Neural Information Processing Systems 1* (Morgan Kaufmann Publishers, San Mateo, 1989) 577–585.

[7] W.J. Bell and K.G. Adiyodi, (Eds.) *The American Cockroach*. (Chapman and Hall, New York, 1981).

[8] R.A. Brooks, A robust layered control system for a mobile robot, *IEEE Journal of Robotics and Automation* RA-2(1) (1986) 14–23.

[9] J.M. Camhi, *Neuroethology* (Sinauer Associates, Sunderland, MA, 1984).

[10] T.J. Carew, The control of reflex action. In E.R. Kandel and J.H. Schwartz (Eds.) *Principles of Neural Science* (Elsevier, New York, 1985) 457–468.

[11] H.J. Chiel and R.D. Beer, A lesion study of a heterogeneous neural network for hexapod locomotion, *Proceedings of the First International Joint Conference on Neural Networks* (IJCNN 89) (1989) 407–414.

[12] F. Delcomyn, Neural basis of rhythmic behavior in animals, *Science* 210 (1980) 492–498.

[13] J. Garcia, W.G. Hankins, K.W. and Rusiniak, Behavioral regulation of the milieu interne in man and rat, *Science* 185 (1974) 824–831.

[14] D. Graham, Simulation of a model for the coordination of leg movement in free walking insects, *Biological Cybernetics* 26 (1977) 187–198.

[15] D. Graham, Pattern and control of walking in insects, *Advances in Insect Physiology* 18 (1985) 31–140.

[16] J.J. Hopfield, Neurons with graded responses have collective computational properties like those of two-state neurons, *Proc. Natl. Acad. Sci.* 81 (1984) 3088–3092.

[17] L. Kaelbling, An architecture for intelligent reactive systems, *Proceedings of the 1986 Workshop on Reasoning about Actions and Plans*, (Morgan Kaufmann Publishers, 1987) 395–410.

[18] E.R. Kandel, *Cellular Basis of Behavior*. (W.H. Freeman, San Francisco, 1976).

[19] E.R. Kandel, Cellular mechanisms of learning and the biological basis of individuality, In E.R. Kandel and J.H. Schwartz (Eds.), *Principles of Neural Science* (Elsevier, New York, 1985) 816–833.

[20] N. Kumar, M. Williams, J. Culotti and D. van der Kooy, Evidence for associative learning in the nematode *C. elegans*, *Society for Neuroscience Abstracts* 15 (1989) 1141.

[21] I.J. Kupfermann, Feeding behavior in *Aplysia*: A simple system for the study of motivation, *Behavioral Biology* 10 (1974) 1–26.

[22] I.J. Kupfermann and K.R. Weiss, The command neuron concept, *Behavioral and Brain Sciences* 1 (1978) 3–39.

[23] R.R. Llinás, The intrinsic electrophysiological properties of mammalian neurons: Insights into central nervous system function, *Science* 242 (1988) 1654–1664.

[24] K.Z. Lorenz, *The Foundations of Ethology*. (Simon and Schuster, New York, 1981).

[25] P. Maes, The dynamics of action selection, *Proceedings of the Eleventh International Joint Conference on AI* (IJCAI 89) (Morgan Kaufmann Publishers, 1989) 991–997.

[26] R.K. Murphey and M.D. Zaretsky, Orientation to calling song by female crickets, *Scapsipedus marginatus(Gryllidae)*, *J. Exp. Biol.* 56 (1972) 335–352.

[27] A. Newell and H.A. Simon, Computer science as empirical inquiry: Symbols and search, *Comm. of the ACM* 19(3) (1976) 113–126.

[28] K.G. Pearson, The control of walking. *Scientific American* 235 (1976) 72–86.

[29] K.G. Pearson, Are there central pattern generators for walking and flight in insects? in W.J.P. Barnes and M.H. Gladden (Eds.), *Feedback and Motor Control in Invertebrates and Vertebrates* (Croom Helm, London, 1985) 307–315.

[30] K.G. Pearson, C.R. Fourtner and R.K. Wong, Nervous control of walking in the cockroach, in R.B. Stein, K.G. Pearson, R.S. Smith and J.B. Redford (Eds.), *Control of Posture and Locomotion* (Plenum Press, 1973) 495–514.

[31] R.E. Ritzmann, The cockroach escape response. in R.C. Eaton (Ed.), *Neural Mechanisms of Startle Behavior* (Plenum Press, New York, 1984) 93–131.

[32] C.H.F. Rowell, H. Reichert and J.P. Bacon, (1985). How Locusts fly straight, In W.J.P. Barnes and M.H. Gladden (Eds.), *Feedback and Motor Control in Invertebrates and Vertebrates* (Croom Helm, London 1985) 337–354.

[33] T.C. Schneirla, Basic problems in the nature of insect behavior, in K.D. Roeder (Ed.), *Insect Physiology* (Wiley, New York, 1953) 656–684.

[34] A.I. Selverston, A consideration of invertebrate central pattern generators as computational data bases, *Neural Networks* 1(2) (1988) 109–117.

[35] G.M. Shepherd, *Neurobiology*. (Oxford University Press, New York, 1988).

[36] N.C. Spitzer, (Ed.) *Neuronal Development* (Plenum Press, 1982).

[37] A.J. Susswein, K.R. Weiss and I. Kupfermann, The effects of food arousal on the latency of biting in *Aplysia*, *J. Comp. Physiol.* 123 (1978) 31–41.

[38] A.H.D. Watson and M. Burrows, The synaptic basis for integration of local reflexes in the Locust. In W.J.P. Barnes and M.H. Gladden (Eds.), *Feedback and Motor Control in Invertebrates and Vertebrates* (Croom Helm, London, 1985) 231–250.

[39] K.R. Weiss, H.J. Chiel, U. Koch and I. Kupfermann, Activity of an identified histaminergic neuron, and its possible role in arousal of feeding behavior in semi-intact *Aplysia*, *J. Neuroscience* 6(8) (1986) 2403–2415.

[40] K.R. Weiss, H.J. Chiel and I. Kupfermann, Sensory function and gating of histaminergic neuron C2 in *Aplysia*, *J. Neuroscience* 6(8) (1986) 2416–2426.

[41] D.M. Wilson, Insect walking, *Annual Review of Entomology* 11 (1966) 103–122.

[42] E.O. Wilson, *Insect Societies* (Harvard University Press, 1971).