ECOLOGICAL PSYCHOLOGY, 4(4), 247-270

Ecological Perspectives on the New Artificial Intelligence

An essay review of *Intelligence as Adaptive Behavior: An Experiment in Computational Neuroethology*. Randall D. Beer. New York: Academic, 1990, xxiii + 213 pp., $29.95, and *Minimalist Mobile Robotics: A Colony-Style Architecture for an Artificial Creature*. Jonathan H. Connell. New York: Academic, 1990, xvii + 175 pp., $29.95.

Judith A. Effken and Robert E. Shaw
Center for the Ecological Study of Perception and Action
University of Connecticut

> *All great progress takes place when two sciences come together, and when their resemblance proclaims itself, despite the apparent disparity of their substance. (Poincaré, cited in de Latil, 1957, p. 4)*

Intelligence as adaptive behavior? Direct measurement of variables from the environment? Using the world as its own representation? Topics not unfamiliar to ecological psychology, but what are they doing in a series of books on AI? The Perspectives in Artificial Intelligence series, published by Academic Press, features work by several designers who are taking a nontraditional approach to AI and robotics that we find most engaging from an ecological psychology point of view. Two volumes: *Minimalist Mobile Robotics: A Colony-Style Architecture for an Artificial Creature*, by Jonathan H. Connell, and *Intelligence as Adaptive Behavior: An Experiment in Computational Neuroethology*, by Randall D. Beer, are able to capture the spirit of the new AI and are the focus of this review.

Requests for reprints should be sent to Judith A. Effken, Center for the Ecological Study of Perception and Action, Department of Psychology, Box U-20, 406 Babbidge Road, Storrs, CT 06269-1020.

Beer and Connell challenge some of the key assumptions of traditional AI, namely that: (a) most intelligent behavior can be modeled as deliberative reasoning or problem solving, (b) human reasoning must be portrayed as computation over symbolic representations, and (c) what we learn from modeling problem solving in restricted domains will eventually form the basis for our understanding of intelligent behavior in general.

Beer argues that the promise of an AI built on these assumptions has not been—and is not likely to be—realized. The intelligence exemplified in even the foremost AI programs remains inflexible and confined to specific domains. Beer believes that artificial intelligence resembles an idiot savant that performs expertly in a certain domain but that is often totally inept outside that domain. Whereas domain-specific intelligence has been captured in a variety of programs (from artificial block worlds to real-world applications of geology or medicine), endowing a program with general intelligence—or common sense—has been more elusive. In the authors' view, conventional AI approaches are bankrupt. It no longer makes sense for designers to respond to each critic's view of what is missing in an "intelligent" machine by tweaking parameters or adding another rule. Instead, they argue that the source of the problem is much deeper; it resides in AI's underlying assumptions. As we see shortly, each volume presents a strong challenge to one or more of these core beliefs.

As the first salvo in their attack on traditional AI, both authors urge that designers rethink their current model of intelligence as deliberative reasoning, arguing instead that a more appropriate model for intelligence would be adaptive behavior, the kind of practical "getting along in the world" that has been the bread and butter of ecological psychology and that reflects our heritage of pragmatic realism (e.g., Gibson, 1967). Note that the term *adaptive* is being used in the sense of contributing to the animal's short-term survival, rather than in the sense of "learning" or long-term structural change. The emphasis is on how an animal engaged in an ongoing interaction with its environment continuously adjusts its behavior to changing internal (propriospecific) and external (exterospecific) circumstances in such a way as to realize its objectives.

Connell has set his sights mainly on assumption b. Traditionally, roboticists and cognitive psychologists alike (e.g., Chatila & Laumond, 1985; Leiblich & Arbib, 1982; Moravec & Elfes, 1985) have claimed that an internal model must be cognitively constructed as a world map wherein the observer's current and desired locations are depicted and updated to keep track of where the observer is relative to the layout of the environment and the desired goal. The ecological approach has shunned this strategy, asserting that information for layout also comprises information for the observer's relationship to that layout. Indeed, one of Gibson's key insights was that there is information in the optic array to indicate both the layout of the environment and motion of the observer in that environment.

> A transformation of the whole optic array specifies a movement of the observer. It is propriospecific. Invariants under transformation of the whole optic array specify invariants of the layout of the environment. They are exterospecific. The separation of the information about the world from the information about the observer's movement in the world, the isolation of the invariants from the motions, is something that a visual system does in its dual role of being both exteroceptive and proprioceptive. (Gibson, 1966, p. 201)

Building on Gibson's ideas, Shaw and Mingolla (1982) criticized the traditional world graph (to use the term of Leiblich & Arbib, 1982) approach as one-sided. That is, an internal model depicts the situation from a single perspective, either that of the organism (as a vector field specifying lines of force) or that of the environment (as a gradient or one-form field specifying equipotential curves conjugate to the lines of force). Shaw and Mingolla argued that both perspectives are needed and can be instantiated in such a way that they function collectively in what Shaw and colleagues (Shaw & McIntyre, 1974; Shaw & Turvey, 1981) called a "coalition." For example, if the organism's perspective is instantiated as connected vectors in a world graph, as in the Leiblich and Arbib approach, then the environmental perspective can be instantiated as contiguous forms in a world map that is mathematically dual to the graph. The effect is analogous to—and no more mystical than—the duality relationship of affordances and effectivities. In each case, the component terms are reciprocal, and neither can be described adequately except in the context of the other.[1] To put it more concretely, the size of an object that will be perceived by an actor as graspable (affording or supporting grasping) is determined by the size of the grasping effector or tool of the actor. Taken together, the graspability of the object and the "grasper-ability" of the actor's physiological/biomechanical subsystems provide a clear demonstration of the inherent duality of affordances and effectivities.

But (and this is Connell's intuition as well as the argument of Shaw and Mingolla) these maps or graphs need not be stored as internal models, as cognitivists would have us believe. They can be created and revised as a robot moves about its world detecting new information. Coordination of Connell's robot is "through the world." That is, the environmental consequence of an action taken by the robot is perceived and used to modify and coordinate behavior directly without intervening mental models or plans.

Gibson is not cited by Connell (although he is cited in the editor's preface). Still, Connell's ideas are reminiscent of Gibson's idea of an active perceiver. In Gibson's formulation, perception is not a passive affair but an active endeavor to

---

[1] The duality of maps and graphs is well known in differential geometry and can be found in descriptions of the duality of vector fields as tangent bundles (e.g., Burke, 1985).

discern relevant information about a particular objective. Given a goal, the observer maneuvers himself or herself so as to better discriminate pertinent information. The information thus gained enables the observer to modulate more precisely further detection and control activities. Detection and control reciprocally fine-tune each other in a perceiving–acting cycle that persists until the desired goal is attained (Figure 1).

In an effective demonstration of the feasibility of this approach, Connell's robot navigates around its world without the aid of a world map or any overall representation of the situation. Instead, the robot depends on information updates from its ongoing perceiving–acting cycle to modify its current goal path. What Connell's robot loses in accuracy (the robot may not reach its target in the most direct route), it makes up in flexibility and extensibility. From a practical point of view, as some ecological psychologists have argued, optimality goals are probably unrealistic as real-world solutions:

> Nature is pragmatic rather than idealistic, demanding only that the goal-path generated by this cycling intention be *tolerably suboptimal* so that life-supporting needs are met and life-threatening situations avoided. Psychological development and health may require more, namely, that information be gathered in the process, or that stress be reduced, affections satisfied, and tastes pleased. (Kugler, Shaw, Vicente, & Kinsella-Shaw, 1991, p. 414)

In any event, flexibility and extensibility seem highly attractive characteristics for an adaptive system.

Beer views flexibility as mandatory in a world where circumstances are continually changing. It is impossible to develop plans to meet any and all

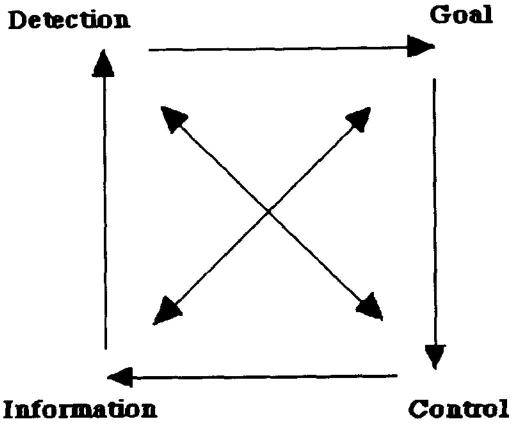

FIGURE 1 Perceiving–acting cycle. From "Reciprocities of Intentional Systems" (p. 589), by R. E. Shaw, P. N. Kugler, and J. Kinsella-Shaw, in *Perception and Control of Self-Motion* edited by R. Warren and A. H. Wertheim, 1990, Hillsdale, NJ: Lawrence Erlbaum Associates, Inc. Copyright 1990 by Lawrence Erlbaum Associates, Inc. Adapted by permission.

contingencies. Beer reminds us of Gladwin's (1964) comparison of Trukese and European navigation methods:

> Once the European navigator has developed his operating plan and has available the appropriate technical resources, the implementation and monitoring of his navigation can be accomplished with a minimum of thought. He has simply to perform almost mechanically the steps dictated by his training and by his initial plan synthesis. . . . [On the other hand, the] total process [of Trukese navigation] goes forward without any planning, unless the intention to proceed to a particular island can be considered a plan. (Gladwin, 1964, p. 175)

It seems that whereas European navigators begin with a plan based on universal principles and follow that plan explicitly, the Trukese begin merely with an overall objective. Pursuing that objective, Trukese navigators react to conditions as they arise, using sensory information provided by the wind, clouds, and water. Although Trukese navigators can point to their goal at any time during the voyage (if asked), they cannot delineate a specific course. In contrast, Europeans may not know where their objective is relative to themselves at any given time, but they know their course and can compute their location on that course (Berreman, 1966). Suchman (1987) asserted that we all act like Trukese—even though we may talk like Europeans! Action is situated; it is contingent on the way the situation actually unfolds. Only when pressed to justify our actions post hoc do we invoke plans, retrospectively reconstructing in careful detail a plan that was probably no more than ad hoc behavior when the activity was carried out.

In ecological theory, situated action has been referred to as the problem of context-conditioned variability (e.g., Turvey, Shaw, & Mace, 1978). For anatomical, mechanical, and physiological reasons, an action as seemingly simple as tossing a ball may be realized in a myriad of ways. Even when the outcome of the action seems to be the same, that is, the ball hits the target with each and every toss, there is no reason to assume that, at the level of muscles and nerves, the act was the same. Each toss will differ depending, at the very least, on the initial position of the arm relative to the body, the position of the body relative to the surface on which it is resting and relative to the target, the characteristics of the ball and the target, and initial neuronal states at the time the "toss" intention was initiated.

Beer contends that the contextual richness of a situation—even a situation as simple as tossing a ball—is lost when complex concepts are abstracted by program designers from observed events and reduced to simple properties that can be encoded and manipulated in an artificial intelligence application. First, the abstractions have no reality of their own but exist only in the mind of the program designer. Second, because the knowledge structures embedded in the program are specific to the domain in which they were created, they have no "meaning" in another domain. Finally, by freezing the abstractions in code, something essentially context sensitive is transformed into something context independent, resulting in diminished flexibility.

Fortunately, neither Beer nor Connell is content to simply criticize traditional approaches. Each offers an alternative of his own: Beer explores encoding intelligence as adaptive behavior and Connell demonstrates goal-based robot "behavior without representation."

INTELLIGENCE AS ADAPTIVE BEHAVIOR: BEER'S ARTIFICIAL COCKROACH

Beer locates control for action in the perceiving–acting cycle of an animal moving through an environment. His cockroach design regards perceptual sensitivity to available goal-specific information as of primary importance. Although conceding that mediating mental states may play a contributory role, Beer is convinced they are neither necessary nor sufficient for the explication of goal-directed behaviors, a point of view with which those of us working in ecological psychology are certainly in agreement.

Beer's emphasis on what ecological psychology portrays as animal–environment duality, characterized by both mutual compatibility and reciprocal dependency (e.g., Shaw, Kugler, & Kinsella-Shaw, 1990), is most directly derived from Maturana and Varela's (1980, 1987) theory of structural congruence in animal–environment relationships. Beer's specific interpretation of animal–environment mutuality has also been influenced by Agre's (1988) "situated activity," Schöner and Kelso's (1988) "dynamic patterns," and approaches to autonomous robot control systems taken recently by Brooks (1986, 1989) and Connell (1987). Beer believes that an animal develops the cognitive structure it does because of the history of structural changes it undergoes, as well as the history of structural changes of its evolutionary ancestors. In ecological psychology, Shaw and Bransford (1977) argued, similarly, that evolution tunes perceptual systems to be most sensitive to the invariants (i.e., aspects that remain unchanged under various transformations) of the greatest adaptive significance to the animal. Gibson (1966) referred to such evolutionary adaptation as "genetic preattunement" to distinguish it from attunement due to learning (i.e., the education of attention).

Having accepted the centrality of animal–environment mutuality, Beer changes his level of analysis from modeling specific cognitive processes in a limited domain to modeling simple whole animals in their environment (what he calls "computational neuroethology"). Echoing Dennett (1978), Beer submits that modeling neural control of behavior in simpler whole animals is more likely to produce insights into the nature of the dynamics required for adaptive behavior in higher animals than modeling subsystems of higher animals in a specific domain and attempting to transfer that knowledge to a different domain.

Beer also acknowledges a debt to Braitenberg's synthetic approach to psychology. Braitenberg's "law of uphill analysis and downhill invention" dictates that it will be easier to create little machines (or, in Beer's case, to model animals) that have simple behaviors—even if the behavior ultimately achieved goes beyond what was originally planned—than to analyze how they accomplish that behavior. It is much more difficult to start outside the "black box" and guess at internal structure from observing behavior. A "problem that taxes the minds of psychologists when they deal with real animals or humans, that of inborn concepts [finds] many solutions when . . . attacked . . . from the downhill, synthetic direction" (Braitenberg, 1984, p. 49).

Historically, the synthetic approach owes a debt to the electronic tortoises built by Grey Walter in the late 1940s and 1950s (de Latil, 1957). Battery-powered, Elmer (Electro-Mechanical Robot) and his "sister" Elsie (Electro-Light-Sensitive-Internal-External) moved freely about their environment, seeking and "feeding" on light, which they subsequently transformed into electric current to recharge their accumulators. When satiated (their batteries charged), their behavior changed. Rather than searching for bright light on which to feed, they instead sought out soft light in which to "rest." But searching for soft light consumed energy so that their motors gradually ran down; they became "hungry" and again went off in search of bright light on which to "feed." When hungry and, therefore, attracted by very bright light, they would head for the source of that light and there plug themselves into a power source to "eat." Created with differing sensitivity to three levels of illumination, the tortoises exhibited remarkably complex behavior patterns because of the complexity of the light patterns in the room. It became impossible for an observer to predict exactly what behavior the tortoises would demonstrate.

In a similar synthetic approach, Beer develops an artificial cockroach that he embeds in a simulated environment replete with natural obstacles. To obtain food when hungry, the insect must search for and find food sources, sometimes negotiating barriers between it and a food source. To create the appropriate actions, Beer incorporates mechanisms that mimic reflexes, taxes, fixed-action patterns, and motivational control previously described by ethologists. When observed in action, Beer's insect exhibits what appears to be goal-oriented, flexible, and stable (repeatable) behavior.

The design of the insect's nervous system is specialized for the body in which it is embedded and the niche in which it is located (what ecological psychology would characterize as affordance–effectivity compatibility) that is reminiscent of Runeson's (1977) special purpose "smart" devices. To provide you with the flavor of Beer's approach, we offer a sample of how the insect gets around and obtains "nourishment" in its world.

Basic rhythmic movements of the insect's legs are produced by a central pattern generator. Stance and swing motor neurons determine the force with which the leg is swung backward or forward, respectively. When a foot is positioned on the ground, those forces propel the insect's body forward or backward. To time transitions in leg movement, the central pattern generator uses propriospecific information about where the legs are at the current time. A backward angle sensor initiates a swing when a leg is all the way back. A forward angle sensor inhibits a current burst that stops the swing and initiates a stance when a leg is all the way forward. Between the starts and stops, the swing action is treated as a ballistic movement. But leg movement by itself is not sufficient for successful navigation.

The six individual leg controllers must be coordinated to ensure that the bug does not fall over when it begins to move. To accomplish this, Beer uses a mutual inhibition routine: Whenever the pacemaker for one leg fires, it inhibits the pacemakers of adjacent legs.

Observed at different speeds, the gaits of the locomoting artificial insect resemble gaits of actual insects at the same relative speeds. At high speeds, the insect moves in a tripod gait (with front and back leg on one side and middle leg on the contralateral side of the insect on the ground at the same time). At slow speeds, step patterns vary, depending on initial conditions. In real insects, metachronal waves (legs swinging in order from the back to the front of the insect) are commonly observed when insects walk slowly. To achieve this phenomenon, Beer lowers the burst frequency of rear pacers and induces phase-locking. Now, as the insect walks faster, metachronal waves (Figure 2) increasingly overlap as stance phases shorten till a phase transition occurs and the tripod gait appears. Beer argues that these gaits cannot be represented statically in a controller, even in a distributed sense, but instead must be actively constructed from dynamic interactions between various controller components and the insect's body and its environment—what ecological psychologists recognize as expropriospecific information (Lee, 1976).

Although Beer has been successful in achieving the typical locomotion patterns of the artificial cockroach, his insect is not yet able to compensate for a missing leg by changing its gait, as live insects can. From an ecological point of

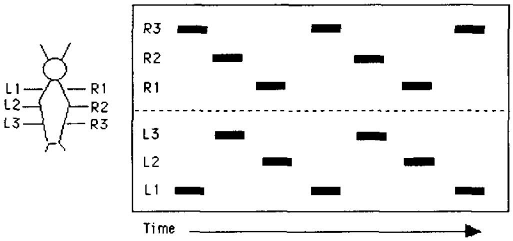

FIGURE 2 Abstract depiction of metachronal waves in one gait shown by Beer's artificial cockroach. Right and left leg steps are indicated. Movement progresses from back to front on each side of the insect. From "Insect Walking" by D. M. Wilson, 1966, *Annual Review of Entomology*, *11*, p. 106. Copyright 1966 by Annual Reviews, Inc. Adapted by permission.

view, the ability of insects to compensate for a lost leg is generally understood as an adaptive gain of a system whose responses are not predefined or "menu-driven" but self-assembled each time a response is required so as to allow the organism to achieve a given goal under a variety of disparate environmental conditions or physiological states. It appears that expropriospecific information about the contact of a limb (or absence of contact) with the surface provides the prospective control required to produce a change in gait that can support the animal's weight and maintain its balance during locomotion. If reflex detectors receive information that a given leg has not taken up its share of the animal's load in an ongoing gait, protraction of the foreleg is delayed. The initial gait is converted into a gait that can maintain the desired goal–locomotion (Wilson, 1980). In this view, as in our previous ball-tossing example, different neurons, muscle groups, or even limbs may be recruited to achieve the final goal.

In Beer's approach, the insect has a fixed repertoire of behaviors that–although varied with increased pace–are not assembled based on need. Motor units are assigned to a specific task or subtask. One motor unit cannot be substituted for another. One way Beer's insect might achieve a stable locomotion pattern when injuries destroy legs or connections among legs would be to have redundant structures that could be recruited in an emergency. But this requires that dominant nodes must be informed about the damage so they can identify the appropriate reserves to recruit. Turvey and colleagues (Turvey, Shaw, & Mace, 1978) argued that this is simply one more form of the context-conditioned variability problem that all action theorists must address–roboticists and ecological psychologists alike.

By optimizing control through relative autonomy, Beer has incorporated a partial solution to what Greene (1972) characterized as a problem in executive control. Given a general, central command, a subsystem can independently carry out its particular subtask unaided and unimpeded by other subsystems. What is lacking in this approach is a way to maintain functional integrity (i.e., some level of function that can attain the goal) through goal-directedness when injuries to effectors or changes in environmental conditions render impossible the usual procedures. Turvey, Shaw, & Mace (1978) suggested that to achieve such integrity, the control system must be designed initially with the entire environment in mind, not just a narrow niche. Instead of a hierarchical model, two heterarchies (organism and environment) that allow bi-directional communication must be linked. This might seem excessively complicated, but because the environment is just as organized as the organism and is specific to the organism (Gibson, 1979/1986; Henderson, 1913/1970), this approach should not increase the degrees of freedom of the problem to be solved. A key interpretation of this fundamental idea is that the minimal unit that can maintain functional integrity can only be a coalition of duals, as neither the organism nor the environmental components can be sufficiently constrained without the other (Shaw, 1987; Shaw & Turvey, 1981). Note that a critical distinction is made between simply adding contextual constraints to a given system and capitalizing on the reciprocity of two systems that serve as contextual constraints for each other.[2] In the ecological view, "a coalition is not a system plus context, but the minimal system that carries its own context" (Turvey, Shaw, & Mace, 1978, p. 592)—one interpretation of what it means for action to be situated.

But let us return to our description of the capabilities of Beer's insect. Whereas the locomotion controller requires an external dose of current to activate it, exploration is something the insect seemingly does spontaneously. The insect wanders, exploring its environment, not because of any exploratory bent but because a wired-in circuit causes periodic turns. This random behavior simply provides a basis on which more sophisticated, intentionally constrained behavior might be built. In this regard, the ecological view emphasizes that alert, awake animals actively seek, rather than passively wait for, stimulation from the sights, sounds, tastes, touches, and smells of their environments (Gibson, 1966). Consider Gibson's example: Harlow's infant monkeys actively sought maternal warmth and cuddling (what Harlow, 1958, called "contact comfort"), but when deprived of this, they found satisfaction in clinging—even to Harlow's mother surrogates. In Gibson's view, perceiving is an inherently satisfying, active activity, not a random behavior. Ultimately, it will be necessary to find ways to expand the insect's exploratory repertoire to include nonrandom, actively seeking behaviors as well as random wandering and reacting to objects encountered. For instance, in the current model, when the insect contacts an object with its antennae, it will turn away or back up, turn, then go forward again, depending on whether it contacts the object head on or at an angle. The challenge for robotics is to build true intentional behaviors based on these Trukese navigation functions. How this might be done is beautifully illustrated by Beers in the following example.

Feeding is an example of goal-oriented (autonomously motivated) behavior. To determine when it needs food, the ongoing energy state of the insect is monitored by sensor. A low energy state is interpreted as "hunger." To get a hungry insect to a food patch, Beer uses information from chemical sensors in the antennae to orient the insect toward a food patch. When the insect's energy level falls, so does the value of its energy sensor. As this value decreases, the normally inhibited feeding arousal neuron is gradually activated. At a predetermined feeding arousal activity level, the "search" command neuron is activated, connections between odor strength neurons and turn neurons are enabled, and the insect begins to orient to food. Once food has been located, consummatory behavior is triggered.

When the insect contacts a food source, rhythmic biting movements begin

---

[2] How contextual, or boundary, constraints might be assimilated into the intrinsic dynamics of an actor has been given a mathematical interpretation in terms of Volterra functionals by Shaw, Kadar, Sim, & Repperger (1992).

and continue until enough food has been "ingested" to raise its energy measure above the "get food" level. In a kind of hysteresis effect, the insect may take in more food than dictated by the internal energy requirements that initiated the behavior. Beer regards this characteristic as highly adaptive, as by "overeating" from time to time, the animal does not have to feed or search for food constantly—a gift from Mother Nature to at least most of us!

The insect's response to food is the result of complex interaction dynamics between an internal positive feedback loop and a negative feedback loop closed through the external environment. Beer disagrees with those who assume that if one understands the complexity of the environment, what seemed to be complexity in the animal will turn out to be simplicity (e.g., Barker, 1968; Simon, 1969). Beer alleges, instead, that complex dynamics exist on both sides. Behavior is due to the interaction between two complex systems. Ecological psychologists (e.g., Michaels & Carello, 1981; Turvey, Shaw, & Mace, 1978) argued in similar fashion that if we are ever to comprehend behavior, we must discover how biological and environmental systems constrain each other so as to achieve an ecological balance in relative complexity. If the environmental description minimizes what the organismic description maximizes—and vice versa, with respect to available information and action control work—then an extension of the powerful mini-max duality principle from mathematical programming theory to intentional dynamical control theory can provide a way to quantify the most parsimonious descriptions of the ecosystem (Shaw, Kadar, Sim, & Repperger, 1992). On the other hand, an imbalance in structural complexity of either the organism or the environment implies a lack of coevolutionary fit of the two—and a violation of the organism-environment duality principle (Shaw & McIntyre, 1974; Turvey & Carello, 1981; Turvey et al., 1978). To the extent that the cognitive constructivist view of the organism ignores this ecological balance principle, its models of the actor/perceiver will be nonparsimonious.

In Beer's insect, a degree of rudimentary choice is achieved by imposing a loose behavioral hierarchy. Generally, higher order behaviors suppress lower order behaviors, but not always. For example, although feeding normally suppresses edge-following, if an obstacle blocks the path to food, edge-following will override feeding in an attempt to get around the obstacle. Thus behaviors, although ordered to some extent by the hierarchy, are also somewhat context dependent. As we have previously noted, some ecological psychologists (e.g., Turvey, Shaw, & Mace, 1978) argued that a hierarchical model—or even a heterarchical model—will be insufficient to encompass the mutual law-based constraints implied by animal-environment mutuality. Only a coalition-styled "ecosystem" that is closed by a duality operation specifying the context of the mutual constraints will do (Shaw & Turvey, 1981).

From an ecological psychology perspective, Beer has begun to explore the ramifications of animal-environment mutuality and the interaction of complex systems. This is surely a welcome move. In the future we, as ecological psychologists, would hope to see additions to this promissory program along the following lines: Goal-directed behavior (intentionality) is manifest by his insect, especially when hungry, but otherwise intention is not running the show. To compensate for this lack, Beer's current tactic is to impose random wandering, a hierarchy of choice options, and somewhat arbitrary pattern suppression or replacement mechanisms that determine behavioral "fly or run" decisions. Designers might consider enhancing the model by incorporating the idea of affordances as possible goals to be selected from an affordance-rich environment by an observer (robot) with need-ordered effectivities. In this view, the observer's effectivities serve to order the process by which particular affordances are selected to be realized. Thus, the following might serve as a preliminary basis for a design strategy.

Iberall and McCulloch (1969) proposed a view of intentions as cycling through modes. For a concrete example of cycling modes of intentional behavior, consider the following illustration provided by Kugler and colleagues (Kugler & Turvey, 1987; Kugler et al., 1991).

In an extraordinary example of coordination, African termites periodically collaborate to build nests of monumental proportions that often survive for hundreds of years. In this cooperative effort, termites work together as independent agents. Their flight patterns are controlled locally by the pheromone (chemical) gradients from their own excretions that have been deposited at the building site. The global patterning or structure of the pheromone field due to the location of excreta results in a cycle of insect behaviors (modes) that are qualitatively different. Initially, excretions are deposited randomly. But given a sufficiently large number of insects and enough time, insects are likely to encounter a location where enough pheromone has been created to be detected. Once the insect encounters gradients produced by the pheromone field, further flight patterns are determined by those gradients. The insects track the gradient to its source and there deposit their waste. Once their waste has been deposited, the insects lose their affinity for pheromone, no longer responding to the pheromone field. As more waste is deposited, the gradient becomes more pronounced. What had been a random process of waste deposition becomes highly organized, resulting in large pillars of waste that serve as major attractors for the depositing of further waste. As the pillars grow, their attractor qualities compete, interacting so that deposits become more prevalent on the side of a pillar adjacent to a competing pillar. Over time, the asymmetrical deposits result in an arch connecting the two pillars. The arch itself acts as a powerful attractor that generates waste deposits in the shape of a flattened surface (dome) atop the arch. With time, the gradient is diffused sufficiently so that waste deposition occurs randomly over the entire dome surface. The entire process (random depositing → tower construction → arch construction → dome construction → random depositing, etc.) can be characterized as modes of behavior that—rather than being "governed by a force field" (as is commonly described in the physical sciences)—are "guided by an information field" (Kugler et al., 1991). The generally conservative nature of the information (pheromone) sources/control (flight paths) field guarantees satisfaction of the ecological balance (mini–max duality principle) and, hence, the most parsimonious account of the insect swarm–environment ecosystem. This can be contrasted with a cognitive constructivist view that imputes unrealistic (and empirically unaccounted for) communication. This view requires that each termite have, minimally, a cognitive plan of construction for the finished architecture and a constantly updated internal model of which termite is doing what, where, and when. Thus, each termite is expected to function as a European navigator. In contrast, the ecological approach requires only that each insect be a Trukese navigator in the nest-building trade. It will be instructive to keep this example of a termite colony in mind as we consider Connell's approach to an analogous problem in robotics design.

CONNELL'S MINIMALIST ROBOTICS

Connell's approach is similar in many respects to that of Beer, but his focus is on solving specific problems encountered by robot designers. The usual robotics approach to robot navigation and object retrieval has been to (a) locate a target, (b) construct a model of the environment, (c) use a pathfinding algorithm to plan the trajectory from the starting point to the target, and (d) map the trajectory back onto the servo system that controls joint movement. The final movement can then be treated as a ballistic act.

Connell maintains that prevailing problems in robotics stem from the centralized nature of world models and path planners. Because of this centralization, designers must channel all sensory inputs into a highly condensed form so they can be used to plan what is essentially a ballistic trajectory. Similar arguments have been made by Beer (as previously discussed) and by Selverston (1980), in his critique of central pattern generator models. In each case, the author argues against reducing complex designer-abstracted concepts into simple encodable properties. Connell relies heavily on vertical decomposition of larger tasks and parallel processing of the resultant simple tasks, assuming that if multiple parallel paths are used, the perceptual burden can be distributed, thereby simplifying the design of the sensors and motor units that must be used. Instead of a centralized sequential program, Connell's robot is controlled by a number of independent "agents" (following Minsky, 1986). The modular approach allows for graceful degradation and easy extensibility but requires that agents remain independent of each other (and this design choice sets limits on the robot's ultimate performance). Whereas each path includes perception, modeling, planning, and execution functions, components of that path are essentially special purpose devices that pay attention to—and expend resources on—only their own tasks. Control of behavior is based solely on local information about the robot's current sensory data and goals; no mental or world models are involved.

Connell's robot perceives and acts locally. In contrast, the termites described by Kugler and colleagues perceive globally and act locally in a mutual, cross-scales dynamical interaction that results in triggering qualitatively different modes of behavior. Because Connell's robot is confined to local perceiving and acting only, we can anticipate that the addition of global intentional constraints might be a reasonable and necessary focus for future development. But, for now, let us return to a description of the current strengths of Connell's robot.

To justify his move away from mental models, Connell cites ethological studies that show that animals typically do not have extensive models of the objects or environment with which they interact. To borrow his example, baby seagulls respond as well to a simple design (a circle with a narrow triangle affixed laterally and a red spot on the triangle) as to a parent. The critical element seems to be the red spot, probably because, in their natural environment (niche), seagulls are the only creatures that show this characteristic (a clear case of animal-environment mutuality).

Like Trukese navigators, all subtasks in Connell's robot need only local information to be completed; but also like Trukese navigators, these agents must monitor the plan along the way. The robot has no need for an internal world model, which is fortuitous because, as Connell asserts, in a distributed system such as this, there is no good place to store a complete world model anyway! It could be argued that a particular agent should be charged with carrying the world model; but because agents are independent, there is no way they can communicate. Thus, the cognitive constructivist solution will not work. Alternatively, each agent might carry its own world model (see Pattee, 1979), but this would be exceedingly expensive; and, even if the cost were not prohibitive, there would still be no guarantee that the copies would be identical. So, Connell wisely elects to use the world as its own representation—in effect, eliminating the need to solve the competing representations problem altogether. Connell has presented us with a "Gibsonian robot" that detects and measures interesting quantities directly from the world (in much the same way African termites detect pheromone gradients). Because it more closely approximates the principle of organism-environment duality, this "ecological" design is more likely to achieve parsimony through balance than traditional cognitivist designs.

Although the robot does not build plans or construct representations, Connell maintains that it does use them. Partial representations and sketchy plans are embedded in the control system, but these "representations" are intrinsic parts of the creature's design and do not change over time. For example, when the hand gropes for a can, it is implicitly using a representation of a table that is partly contained in the way the finger sensors were installed, partly in standard directions of arm movement and partly in three independent control agents. One agent (Descend) "knows" that you contact tables by going down, another agent (Bounce) "knows" that the hand cannot be in contact with a table if it wants to move, and a third agent (Surface) "knows" that a table can be followed by extending the arm outward. Another agent (Over) responds to things that fit its "ungraspable object model." Keyed to the robot's own size, this is an intrinsic metric of affordance-effectivity fit ecological psychologists can heartily endorse. This "genetic preattunement" is embodied in the robot by Connell as a collection of what he terms "instincts" (in the sense that they are not consciously deduced). The robot's "plan" is actually derived from the sequence in which the environment allows behaviors to be triggered (another example of animal-environment mutuality). As we soon see, this is sufficient to enable the robot to carry out a very specific task.

Connell's robot navigates through a cluttered environment to find and retrieve soda cans (Figure 3). To locate and pick up a can on a table, the robot raises its hand to the top of its workspace, extends the hand slightly, then brings the hand down to find the surface of the table. If the fingers touch the surface, the hand then bounces along the surface in search of a can. When the hand approaches a soda can, local proximity sensors align it with the can and guide grasping. Then the robot lifts the can from the table and brings it back to a position next to the body. Two underlying principles guide the robot's behavior:

1. The robot's trajectory is guided by environmental constraints. As the hand moves, changing properties of the environment are used to select and activate appropriate behaviors. Hence, the robot can navigate cluttered areas without a detailed world model.
2. A system composed solely of independent local agents can show globally directed behavior. The hand is controlled cooperatively by a set of six independent behaviors. These independent behaviors interact through suppressor nodes; that is, higher behavior can suppress a lower behavior or replace it with its own output. The arm is divided into a number of levels

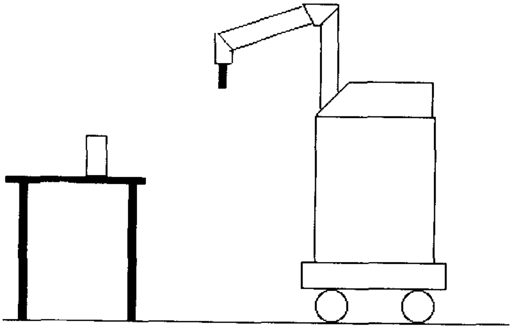

FIGURE 3 Connell's robot in its world. From *Minimalist Mobile Robotics: A Colony-Style Architecture for an Artificial Creature* (p. 45) by J. H. Connell, 1990, Boston: Academic. Copyright 1990 by Academic. Adapted by permission.

of competence – a hierarchy of subbehaviors. The designer can expand the robot's performance by building on previous levels, that is, by reducing the new behavior to a situation the robot knows how to manage so that preexisting routines can be used to carry out the activity. Consequently, various levels do not depend on each other's structure but on their functionality.

Detection and Object Retrieval

Connell's robot can detect objects within 30 degrees of the direction of travel and up to 2 meters away. In a simplifying move, the robot maps only the edges of the arm's effective workspace into visual coordinates. To retrieve a can, the robot moves so that the image of the object is within the image of its workspace and then releases its arm. Unknowingly (we assume), Connell has encoded Gibson's (1979/1986) rule for manipulating something graspable: "To manipulate something graspable, magnify the patch to such a degree that the object is within reach" (p. 233). If the robot approaches a large object originally perceived as graspable and its size becomes larger than the workspace on which it is superimposed, the robot loses interest in it as a graspable object and moves on to more promising candidates.

Connell contends that the complexity of the hand–eye coordination problem can be reduced by relying on the behavior of the arm controller. As argued earlier and here exemplified by Connell, the problem becomes simpler because of the introduction of mutual constraints exerted by dual control and detection components of the perceiving–acting cycle on each other (see Turvey, 1990).

Goal-Directed Navigation

In robot navigation, Connell's main argument is that a complete path plan need not be a data structure. Instead, one can use a simple decision strategy to choose between path segments and a local navigation procedure to detect relevant information about each segment. Appropriate behavior (in the sense of tactical navigation) comes as a consequence of two competing rules (intentions): (a) Avoid contact with obstacles and go around obstructions, and (b) remain within sensor range of some part of the environment. As the only way to stay near an object while moving is to follow it, the robot moves along the coast of its world. A useful trait emerges: The robot follows walls at a set distance.

All information used by the robot is task (or goal) oriented; perception is directly coupled to action. The robot concentrates only on that part of the environment directly relevant to current activities. No deep-seated understanding of obstacles is needed to maneuver around them.

Certainly, Connell's robot, in the version we described, is not what might ultimately be needed or desired, but most certainly it has provided us with a foundation on which future enhancements might be made that can alleviate its present limitations:

1. Currently, the robot can get caught in local minima or loops. To overcome this, Connell adds an agent that monitors some sensory variable measuring the robot's overall progress. If the current set of behaviors does not show progress, the agent turns off that set and turns on another. For example, if the hand has stopped moving, the current set of behaviors is deemed ineffective; so the agent HOME returns the hand to the parked position.
2. Connell allows the robot to execute paths within specified tolerance regions, rather than requiring it to achieve optimal paths. As we pointed out in our discussion of Beer's insect, this is also the approach taken by a number of ecological psychologists and should not be taken as a weakness in his model. By allowing design choices to fall within a tolerance region, Connell loses the advantage of a complete, specific mapping that would afford more direct movement toward the can, but gains in simplicity and flexibility as well as realism. If the robot could learn, then it could potentially improve its performance over time, but so far no such learning has been incorporated in this model, an obvious direction for future development in both Connell's and Beer's work.
3. The robot discovers certain sensory data related to its environment but does not detect the invariants (what we, as ecological psychologists, would consider the vital basis for information) in its world. Consequently, the designer has to build in knowledge like "cans are upright" and "cans are normally found on tables" as rules. But these rules lack flexibility. Empty soda cans may well be overturned or lying on the floor or under a chair—at least in our laboratory! From an ecological point of view, it would appear that designers could gain more flexibility by exploiting Gibson's affordance concept.

For example, in Connell's can-collecting task, the affordance structure of a soda can can be described as an object that is able to hold liquid, stand upright, be grasped, be lifted, and be drunk from by humans. To facilitate the robot's recognition of soda cans, the designer would need to incorporate information about the *structural* and *transformational* invariants that specify the encounters (events) allowed by objects with this affordance structure—assuming a robot that does not learn (Pittenger & Shaw, 1975; Shaw, Flascher, & Mace, in press; Warren & Shaw, 1981). Structural invariants are those properties that remain unchanged under transformation. For example, a vessel that can hold liquid will have a surface structure that is hollow and remains hollow even if the object is turned, tossed, or heated. Transformational invariants describe those styles of change that are perceived as the same change across a number of structures (or objects). For example, balls, soda cans, and infants all can be lifted (see Kim, 1992). As Gibson (1966) pointed out, "When the constant properties of constant objects are perceived (the shape, size, color, texture, composition, motion, animation, and position relative to other objects), the observer can go on to detect their *affordances*" (p. 285), that is, the behaviors or acts that objects permit (Michaels & Carello, 1981).

One of the noteworthy characteristics of invariants is that they remain unchanged under changes of perspective; hence, a robot built to recognize effectivity-based affordances would not need to have knowledge built in that cans are usually upright and normally on tables. Nor would the robot require countless perspectives of a given figure to be preprogrammed so that pattern matching can be performed. The design tradeoff would lie in (a) determining the invariant structure that specifies the affordance—an endeavor that lies at the heart of the ecological psychology program—and (b) building a robot with the capability to detect (and ultimately, discover) the invariants in its world. Connell speculates that one might be able to bootstrap learning on known behaviors (i.e., door recognition on wall hugging), but it seems that an ability to detect the invariants in its world will be a critical preliminary for learning.

4. Connell's robot is unable to plan an alternate route if the usual path is blocked. (Note the similarity to Beer's insect here.) Nor can it adjust its behaviors by redesigning its effectivities to compensate when a given behavior is thwarted. Thus, it would seem that one opportunity for further development might be to explore ways in which the robot might assemble and disassemble the relevant effectivities.

5. Arbitration is purely competitive; there is no compromise. Nor do priorities change over time. Connell conjectures that one might add graded responses or, abandoning suppressors and inhibitors altogether, use a potential field as Arkin (1987) has done. But Connell points out that Arkin's approach is world-centered, not robot-centered. Whereas we agree that a robot-centered view is needed, we would argue that neither world-centered nor robot-centered views will be sufficient by themselves. An animal–environment mutuality assumption dictates that it will be necessary to incorporate both views to obtain the full picture (see Shaw & Kinsella-Shaw, 1988). Moreover, it can be shown that moving between world- and robot-centered views is simply a coordinate transformation that can be effected by a self-adjoint system with an intentional operator (Shaw & Mingolla, 1982). A mathematical formulation for such a self-adjoint system—including the social operator required to move between the two coordinate systems—has been given by Shaw, Kadar, Sim, and Repperger (1992).

Another decision-making alternative Connell considers is theta-space summation. Under this approach, behaviors are rated in all possible directions. In this approach, commonly known as cost functional, minimal cost or its dual—maximal value—is computed. For an ecological (more specifically, intentional dynamics) version of this approach, we would suggest that Connell consider endowing his robot with a secondary preference for direction or assigning it several "manner parameters" (Shaw & Kinsella-Shaw, 1988) and propensities for each (i.e., although it is better to overshoot the corner of a fence than undershoot, it is better to undershoot the approach to a cliff than overshoot). When manner parameter preferences are added to target (goal) preferences, the robot effectively has a priority setting procedure for determining a goal path. Now, not only does the robot "know" where the goal is, but also how to reach it (Shaw & Kinsella-Shaw, 1988).

In the intentional dynamics view, intentionally selecting a goal defines the workspace (or context) in which further perceiving and acting will take place, effectively setting the boundary conditions on the range of affordances that might be individuated. Further, selecting a goal makes available specific characteristics of the goal (i.e., information about the target parameters of time-to-contact, direction-to-contact, distance-to-contact) that are dual to manner parameters (impulse-to-contact, torque-to-contact, and work-to-contact) describing how the goal is to be reached. This duality suggests the conditions under which one might move directly from perceiving a desired objective in the environment to self-assembling the required mechanisms to reach the goal in a perceiving–acting cycle.

If a robot were truly to be driven by target and manner goal parameters, it would effectively instantiate Weir's (1984) view of goal-directed systems as systems directed by goals, rather than as systems self-directed toward goals. Such goal-directed behavior is nonanalytic, bifurcatory behavior requiring "germ-determinacy" (one-to-many mappings) rather than "state-determinacy" (one-to-one mappings) for explanation (Shaw, 1987). That is, a goal is not the designated final state to be reached deterministically by a system but a unique way for a nondeterministic system to reach the final state over one of several potential paths (mathematically, a "germ" of paths or one-to-many mapping to be described later). Goals can exercise control over behavior because changes in the parameters of the goal (e.g., distance-to-the-target) are perceptually projected directly onto corresponding changes in the actor's control variables (e.g., work required to reach the target; Koditschek, 1989; Shaw & Kinsella-Shaw, 1988; Kugler et al., 1991).

How an organism's behavior could be determined by a goal without having an effect precede its cause (the teleological problem) has been a source of concern for philosophers and psychologists alike. To avoid this problem, some theorists have adopted a state-determinacy approach, claiming that goal states can be reduced to initial conditions plus efficient conditions (laws). This approach effectively reduces intention to determining how and for what purpose a system was designed (Kugler et al., 1991). Furthermore, state determinacy implies that an analytical solution can be employed in which one can project past or future behavior from a given segment of observed behavior. But this solution does not work. Mid-course corrections in goal-directed actions show up as choice-points (or bifurcations) in action paths that are not analytically continuous with either past or future states, even though the ultimate paths are patterned in anticipation of future goal states. To explain anticipatory control requires a means of allowing the goal to somehow determine mid-course corrections. One possible explanation has been called *path-determinacy*. Path-determinacy endeavors to characterize final causation lawfully by designing goal paths moving backward from the system's final state to its initial state according to a minimum principle. Mathematically, the method compares potential path integrals (curves) and selects the minimal curve. Although a similar comparison is made in state-determinacy approaches, there the comparison is made between points. Despite these differences, the two approaches turn out to be mathematically equivalent. Moreover, neither method has been able to explain how a goal can modify behavior (Kugler et al., 1991).

In contrast, Weir (1984) proposed that a goal path can be characterized as a bundle of virtual paths that may coincide to a point of bifurcation then branch into an assortment of separate paths, each a possible different realization of the goal. This bundle is known as a "germ" in mathematics because, at the bifurcation point, the mapping is one-to-many. In Weir's formulation, "the germ is the formal cause of a goal-directed path, that is, a dynamical principle which expresses nondeterministically the distinctive shapes that paths may assume" (Kugler et al., 1991, p. 411). Intentionally selecting a goal thus constrains behaviors to the set of possible paths that might be followed to achieve the goal, then sustains perceptual activities to detect more detailed information about the goal and determines the energetics for action required to reach the goal. This idea has been made mathematically explicit by Kugler and Shaw (1989); Kugler and Turvey (1987); Shaw and Kinsella-Shaw (1988); and Shaw, Kugler, and Kinsella-Shaw (1990). The key idea (made clear in the termite example) is that organisms respond, not only to local contingencies, but also to global goal constraints that can determine local actions, such as which path is selected (or how an action is modified) at a choice-point.

Connell claims (and we heartily agree) that his robot is an existence proof that a moderately complicated system can be built that is decentralized and that does not contain explicit world models. Moreover, as Simon (1969) pointed out, this simulation has shown us directly what occurs when a large number of variables interact: New behaviors emerge—behaviors that are appropriate for the prevailing environmental conditions.

CONCLUSION

Surely the work described in the two volumes represents a new departure for AI research, one that is moving in a direction more compatible with ecological psychology. It seems likely that the two efforts could do a great deal to facilitate each other's work. Both groups are working on similar problems—albeit with different tools. We contend that the analytic approach of ecological psychology can provide a complementary counterpoint to the synthetic approach taken by Beer and Connell and vice versa. For its part, ecological psychology offers AI an alternative theoretical framework for perception and action that seems to be congruous with the direction these designers have chosen. To some of us, the abstract model of intentional dynamics intuitively developed and given a generic mathematical description by Shaw and colleagues (e.g., Shaw et al., 1992; Shaw et al., 1990; Shaw & Repperger, 1990) seems made to order for a robotics testbed because it explicitly models how a perceiving–acting cycle is goal-driven in an environment.

Ecological theory is being applied successfully to computer interface design (e.g., Effken & Shaw's hemodynamic monitoring display designs, 1991; Gaver's ARKola, 1991; Vicente and Rasmussen's EID, 1990). It seems likely that the theory can be extended to AI and robotics with equal success.

The fact that a group of AI designers have recognized organism-environment duality as a basic design constraint suggests that the next logical step might be for this group of AI designers and ecological psychologists to join ranks to: (a) further explore the practical possibilities of affordance–effectivity compatibilities in goal-directed (intentional) systems and (b) develop systems that can detect, use, and learn from the invariant information available in the world. What we as ecological psychologists have learned about information detection, affordance–effectivity compatibilities, and navigation in cluttered environments should serve as valuable resources for this work, as should initial forays into learning theory (Johnston & Pietrewicz, 1985; Shaw & Alley, 1985) and intentional learning (Shaw et al., 1992) from an ecological perspective. On the other hand, the approaches that these roboticists are taking in developing constructive, working solutions to our shared problems will surely help us as much as we might help them.

Connell's demonstration of Gibson's rule for manipulating something graspable as a useful design strategy suggests that Gibson's other rules for the visual control of locomotion and manipulation might be equally useful. Gibson (1979/1986) has given us rules for standing up, starting, stopping, steering, approaching, entering enclosures, keeping safe distances, reaching, aiming, and throwing. Moreover, important empirical insights about information for time-to-contact when braking to reach a target abound (Kim, Turvey, & Carello, in press; Lee, 1976; Todd, 1982). But, it is clear from a review of the literature that adequate empirical testing of any of these rules remains to be done. Future AI demonstrations in the spirit of Connell's robot could do much to validate the utility of Gibson's rules—particularly if supported by the requisite empirical studies.

The synthetic approaches of AI and robotics designers offer ecological psychologists simulation and design possibilities that could have great value in the validation and extension of our theoretical frameworks. There is already an intersection of interest groups on complex systems, robotics, and ecological psychology from Boston University, the Massachusetts Institute of Technology, the University of Connecticut, Haskins Laboratories, and Yale University that frequently share theoretical insights and research data. It seems to us that another special interest group might be formed—composed of designers such as Beer, Brooks, Connell, and their colleagues together with ecological psychologists—for the purpose of discovering how perception and action might be integrated in goal-directed systems. For our part, we will be following with great interest and enthusiasm the line of research these designers and their colleagues have begun and, at the same time, hope for the opportunity to open such dialogues.

ACKNOWLEDGMENTS

We thank Claudia Carello for her insightful comments on an earlier version of this article.

REFERENCES

Agre, P. E. (1988). *The dynamic structure of everyday life* (Tech. Rep. No. AI-TR-1085). Cambridge, MA: Massachusetts Institute of Technology, Artificial Intelligence Laboratory.

Arkin, R. A. (1987). Motor schema based navigation for a mobile robot: An approach to programming by behavior. *Proceedings of the 1987 IEEE International Conference on Robotics and Automation*, 264–271.

Barker, R. G. (1968). *Ecological psychology*. Stanford, CA: Stanford University Press.

Beer, R. D. (1990). *Intelligence as adaptive behavior: An experiment in computational neuroethology*. New York: Academic.

Berreman, G. (1966). Anemic and emetic analyses in social anthropology. *American Anthropologist*, 68(2), 346–354.

Braitenberg, V. (1984). *Vehicles: Experiments in synthetic psychology*. Cambridge, MA: MIT Press.

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

Brooks, R. A. (1989). A robot that walks: Emergent behaviors from a carefully evolved network. *Neural Computation, 1*, 253–262.

Burke, W. L. (1985). *Applied differential geometry*. Cambridge, England: Cambridge University Press.

Chatila, R., & Laumond, J. (1985). Position referencing and consistent world modeling for mobile robots. *Proceedings of the 1985 IEEE Conference on Robotics and Automation*, 138–145.

Connell, J. H. (1987). Creature design with the subsumption architecture. *Proceedings of AAAI-87*, 1124–1126.

Connell, J. H. (1990). *Minimalist mobile robotics: A colony-style architecture for an artificial creature*. New York: Academic.

Dennett, D. C. (1978). Intentional systems. In D. C. Dennett (Ed.), *Brainstorms: Philosophic essays on mind and psychology* (pp. 1–22). Cambridge, MA: MIT Press.

Effken, J. A., & Shaw, R. E. (1991). Coordination in an intensive care setting. In P. J. Beek, R. J. Bootsma, & P. C. W. van Wieringen (Eds.), *Studies in perception and action: Posters presented at the VIth International Conference on Event Perception and Action* (pp. 289–292). Amsterdam: Rodopi.

Gaver, W. W. (1991). Sound support for collaboration. *Proceedings of ECSCW'91.*

Gibson, J. J. (1966). *The senses considered as perceptual systems.* Boston: Houghton Mifflin.

Gibson, J. J. (1967). New reasons for realism. *Synthese, 17,* 162–172.

Gibson, J. J. (1986). *An ecological approach to visual perception.* Hillsdale, NJ: Lawrence Erlbaum Associates, Inc. (Original work published 1979)

Gladwin, T. (1964). Culture and logical process. In W. Goodenough (Ed.), *Explorations in cultural anthropology: Essays presented to George Peter Murdick* (pp. 167–177). New York: McGraw-Hill.

Greene, P. H. (1972). Problems of organization of motor systems. In R. Rosen & F. M. Snell (Eds.), *Progress in theoretical biology* (Vol. 2, pp. 304–338). New York: Academic.

Harlow, H. H. (1958). The nature of love. *American Psychologist, 13,* 673–685.

Henderson, L. J. (1970). *The fitness of the environment.* Gloucester, MA: Peter Smith. (Original work published 1913)

Iberall, A. S., & McCulloch, W. S. (1969). The organizing principle of complex living systems. *Journal of Basic Engineering, Transactions of the ASME, 91*(2), 290–294.

Johnston, T. D., & Pietrewicz, A. T. (1985). *Issues in the ecological study of learning.* Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Kim, N-G. (1992). *Perceiving persistence under change and over structure: An investigation of the event perception hypothesis.* Unpublished doctoral dissertation, University of Connecticut, Storrs.

Kim, N-G., Turvey, M. T., & Carello, C. (in press). Optical information about the severity of upcoming contacts. *Journal of Experimental Psychology: Human Perception and Performance.*

Koditschek, D. E. (1989). Robot planning and control via potential functions. In O. Khatib, J. J. Craig, & T. Lozano-Perez (Eds.), *The robotics review I* (pp. 349–368). Cambridge, MA: MIT Press.

Kugler, P. N., & Shaw, R. E. (1989). Symmetry and symmetry-breaking in thermodynamic and epistemic engines: A coupling of first and second laws. In H. Haken & M. Stadler (Eds.), *Synergetics of cognition* (pp. 296–331). New York: Springer-Verlag.

Kugler, P. N., Shaw, R. E., Vicente, K. J., & Kinsella-Shaw, J. (1991). The role of attractors in the self-organization of intentional systems. In R. R. Hoffman & D. S. Palermo (Eds.), *Cognition and the symbolic processes: Applied and ecological perspectives* (pp. 387–431). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Kugler, P. N., & Turvey, M. T. (1987). *Information, natural law, and the self-assembly of rhythmic movement.* Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Latil, P., de. (1957). *Thinking by machine: A study of cybernetics.* Boston: Houghton Mifflin.

Lee, D. N. (1976). A theory of visual control of braking based on information about time to collision. *Perception, 5,* 437–459.

Leiblich, I., & Arbib, M. A. (1982). Multiple representations of space underlying behavior. *The Behavioral and Brain Sciences, 5,* 627–659.

Maturana, H. R., & Varela, F. J. (1980). *Autopoiesis and cognition* (Boston Studies in the Philosophy of Science, Volume 42). Boston: D. Reidel.

Maturana, H. R., & Varela, F. J. (1987). *The tree of knowledge.* Boston: Shambhala.

Michaels, C. F., & Carello, C. (1981). *Direct perception.* Englewood Cliffs, NJ: Prentice-Hall.

Minsky, M. (1986). *The society of mind.* New York: Simon & Schuster.

Moravec, H. P., & Elfes, A. (1985). High resolution maps from wide angle sonar. *Proceedings of the 1985 IEEE International Conference on Robotics and Automation,* 1838–1845.

Pattee, H. H. (1979). The complementarity principle and the origin of macromolecular information. *Biosystems, 11,* 217–226.

Pittenger, J., & Shaw, R. E. (1975). Aging faces as visco-elastic events: Implications for a theory of non-rigid shape perception. *Journal of Experimental Psychology: Human Performance and Perception, 1,* 374–382.

Runeson, S. (1977). On the possibility of "smart" perceptual mechanisms. *Scandinavian Journal of Psychology, 18,* 172–179.

Selverston, A. I. (1980). Are central pattern generators understandable? *The Behavioral and Brain Sciences, 3,* 535–571.

Schöner, G., & Kelso, J. A. S. (1988). Dynamic pattern generation in behavioral and neural systems. *Science, 239,* 1513–1520.

Shaw, R. E. (1987). Behavior with a purpose. *Contemporary Psychology, 32,* 243–245.

Shaw, R. E., & Alley, T. R. (1985). How to draw learning curves: Their use and justification. In T. D. Johnston & A. T. Pietrewicz (Eds.), *Issues in the ecological study of learning* (pp. 275–443). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Shaw, R. E., & Bransford, J. (1977). Introduction: Psychological approaches to the problem of knowledge. In R. E. Shaw & J. Bransford (Eds.), *Perceiving, acting and knowing: Toward an ecological psychology* (pp. 1–39). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Shaw, R. E., Flascher, O., & Mace, W. (in press). Dimensions of event perception. In W. Prinz & B. Bridgeman (Eds.), *Handbook of perception.* Berlin: Springer-Verlag.

Shaw, R. E., Kadar, E., Sim, M., & Repperger, D. W. (1992). The intentional spring: A strategy for modelling systems that learn to perform intentional acts. *Journal of Motor Behavior, 24,* 3–28.

Shaw, R. E., & Kinsella-Shaw, J. (1988). Ecological mechanics: A physical geometry for intentional constraints. *Human Movement Science, 7,* 155–200.

Shaw, R. E., Kugler, P. N., & Kinsella-Shaw, J. (1990). Reciprocities of intentional systems. In R. Warren & A. H. Wertheim (Eds.), *Perception and control of self-motion* (pp. 579–619). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Shaw, R. E., & McIntyre, M. (1974). Algoristic foundations to cognitive psychology. In W. Weimer & D. Palermo (Eds.), *Cognition and the symbolic processes* (pp. 305–361). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Shaw, R. E., & Mingolla, E. (1982). Ecologizing world graphs. *The Behavioral and Brain Sciences, 5,* 648–650.

Shaw, R. E., & Repperger, D. (1990). An adjoint systems approach to learning and transfer of training: 1 Introduction. *PAW Review: A Technical Report of the Perceiving–Acting Workshop in Ecological Psychology, IV*(2), 33–43.

Shaw, R. E., & Turvey, M. T. (1981). Coalitions as models for ecosystems: A realist perspective on perceptual organization. In M. Kubovy & J. Pomerantz (Eds.), *Perceptual organization* (pp. 343–415). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Simon, H. (1969). *The sciences of the artificial.* Cambridge, MA: MIT Press.

Suchman, L. A. (1987). *Plans and situated actions.* Cambridge, England: Cambridge University Press.

Todd, J. T. (1982). Visual information about rigid and nonrigid motion: A geometric analysis. *Journal of Experimental Psychology: Human Perception and Performance, 8,* 238–252.

Turvey, M. T. (1990). Coordination. *American Psychologist, 8,* 938–953.

Turvey, M. T., & Carello, C. (1981). Cognition: The view from ecological realism. *Cognition, 10,* 313–321.

Turvey, M. T., Shaw, R. E., & Mace, W. (1978). Issues in the theory of action: Degrees of freedom, coordinative structures and coalitions. In J. Requin (Ed.), *Attention and performance VII* (pp. 557–595). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Turvey, M. T., Shaw, R. E., Reed, E. S., & Mace, W. M. (1981). Ecological laws of perceiving and acting: In reply to Fodor and Pylyshyn. *Cognition, 9,* 237–304.

Vicente, K. J., & Rasmussen, J. (1990). The ecology of human–machine systems II: Mediating "direct perception" in complex work domains. *Ecological Psychology, 2,* 207–250.

Warren, W., & Shaw, R. (1981). Psychophysics and ecometrics. *Behavioral and Brain Sciences, 4,* 209–210.

Weir, M. (1984). *Goal-directed behavior.* New York: Gordon and Breach.

Wilson, D. M. (1980). Insect walking. In C. R. Gallistel (Ed.), *The organization of action: A new synthesis* (pp. 115–152). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.