

explain why this particular pavement is slippery right now, and to decide whether this can be prevented. The background store of knowledge includes general models of pavements, sprinklers, slipping, rain, and so on; these must be accessed and supplied with instance data to construct the specific Bayesian network structure. The store of background knowledge must utilize some representation that combines the expressive power of first-order logical languages (such as semantic networks) with the ability to handle uncertain information. Substantial progress has been made on constructing systems of this kind (Koller and Pfeffer, 1998), but as yet no overall cognitive architecture has been proposed.

Road Maps: Artificial Intelligence; Learning in Artificial Networks

Related Reading: Bayesian Methods and Neural Networks; Decision Support Systems and Expert Systems; Graphical Models: Probabilistic Inference

References

- Binder, J., Koller, D., Russell, S., and Kanazawa, K., 1997, Adaptive probabilistic networks with hidden variables, *Machine Learn.*, 29:213–244.
- Dean, T., and Kanazawa, K., 1989, A model for reasoning about persistence and causation, *Computat. Intell.*, 5:142–150.
- Doucet, A., de Freitas, J., and Gordon, N., 2001, *Sequential Monte Carlo Methods in Practice*, Berlin: Springer-Verlag.
- Friedman, N., 1998, The Bayesian structural EM algorithm, in *Uncertainty in Artificial Intelligence: Proceedings of the Fourteenth Conference*

- (G. F. Cooper and S. Moral, Eds.), San Mateo, CA: Morgan Kaufmann, pp. 129–138.
- Goldszmidt, M., and Pearl, J., 1996, Qualitative probabilities for default reasoning, belief revision, and causal modeling, *Artif. Intell.*, 84:57–112.
- Jensen, F. V., 1996, *An Introduction to Bayesian Networks*, New York: Springer-Verlag. ♦
- Jordan, M. I., Ed., 1999, *Learning in Graphical Models*, Cambridge, MA: MIT Press. ♦
- Jordan, M. I., and Sejnowski, T. J., Eds., 2001, *Graphical Models: Foundations of Neural Computation*, Cambridge, MA: MIT Press.
- Kim, J. H., and Pearl, J., 1983, A computational model for combined causal and diagnostic reasoning in inference systems, in *Proceedings of the Eighth International Joint Conference on Artificial Intelligence (IJCAI-83)*, San Mateo, CA: Morgan Kaufmann, pp. 190–193.
- Koller, D., and Pfeffer, A., 1998, Probabilistic frame-based systems, in *Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI-98)*, Menlo Park, CA: AAAI Press, pp. 580–587.
- Lauritzen, S. L., and Spiegelhalter, D. J., 1988, Local computations with probabilities on graphical structures and their application to expert systems (with discussion), *J. R. Statist. Soc.*, series B, 50:157–224.
- Pearl, J., 1988, *Probabilistic Reasoning in Intelligent Systems*, San Mateo, CA: Morgan Kaufmann. ♦
- Pearl, J., 2000, *Causality: Models, Reasoning, and Inference*, New York: Cambridge University Press. ♦
- Shafer, G., and Pearl, J., Eds., 1990, *Readings in Uncertain Reasoning*, San Mateo, CA: Morgan Kaufmann.
- Tenenbaum, J. B., and Griffiths, T. L., 2001, Structure learning in human causal induction, in *Advances in Neural Information Processing Systems 13*, Cambridge, MA: MIT Press.

Biologically Inspired Robotics

Noel E. Sharkey

Introduction

At the beginning of the twenty-first century, living organisms have still not been successfully replicated by machines. Computers are much faster at number crunching than humans and can even beat the greatest at chess, and other machines can perform routine physical tasks faster than us and with a precision that we cannot approach. However, animals exhibit such remarkable capacities for flexible adaptation to novel circumstances that roboticists can only gaze in wonder. It is thus an important goal of modern robotics to learn from the way organisms are constructed biologically, and how this creates adaptive behaviors.

Biologically inspired robotics, also known as biomimetic robotics or biorobotics, refers to robotics research in which the life sciences, including biology, psychology, ethology, neuroscience, and evolutionary theory, play a key role in motivating the research. It is necessarily broad because the field is just beginning to emerge as a unified discipline, and so it still has fuzzy boundaries. Biorobotics research ranges from modeling animal sensors in hardware for guiding robots in target environments to investigating the interaction between neural learning and evolution in a variety of robot tasks. There are, however, common themes that will be explored here.

In the following sections, some of the major issues in biorobotics research and the aims of this approach are examined. First we briefly consider the historical roots of the core ideas. The seminal work of Grey Walter (1953) sets the scene and introduces some of the key elements of biologically inspired robotics. The re-introduction and development of the ideas in the 1980s occurred with Braitenberg's *synthetic psychology* and Brook's *behavior-based robotics*. In summarizing the breadth of the current work, we attempt a threefold classification of biologically inspired robotics.

The Roots of Biologically Inspired Robotics

The roots of biologically inspired robotics date back to the early twentieth century, when Hammond constructed a heliotrope based on the biologist Loeb's tropism theory of animal behavior. Loeb proposed that animals are attracted and repelled by stimuli in the environment in a way similar to the phototropic responses of plants. Although Hammond's heliotrope did not model an animal, its mechanized movement toward light was sufficient to satisfy Loeb that his theory has physical plausibility (cf. Sharkey and Ziemke in Ziemke and Sharkey, 1998, pp. 361–392, for an account).

There were a number of robot learning studies during the first half of the twentieth century, before the birth of artificial intelligence (AI). However, the prototypical biorobotics work was conducted by Grey Walter (1953). He went far beyond Hammond in testing the mechanistic plausibility of animal tropism. His aim was to create a self-sustaining artificial life form that could adapt. This required the development of a robot that could seek out a source to recharge its batteries on demand.

Grey Walter used electromechanical robots equipped with two input "receptors": a photo-electric cell for sensitivity to light, and an electrical contact as a touch receptor. The controller, between sensors and motors, was a small artificial nervous system built from miniature valves, relays, condensers, batteries, and small electric motors—no computer. There was a hutch where a robot could drive in to have the battery automatically recharged.

Behavior resulted from the interaction of the internal states of the robot (battery level) and the intensity of light sources, as well as other environmental factors such as obstacles. When the battery levels were high, the robot was repelled by the bright light of the hutch and attracted by the moderate light in the room, where it

“explored.” With low battery levels, the robot was attracted to the bright light of the hutch for an automatic recharge. In this way Grey Walter demonstrated that mechanical tropism could work as a means of exploration and maintaining energy.

Grey Walter (1953) also investigated adaptation and showed how a simple learning mechanism could extend the behavior of a robot using the conditioned reflex analog (CORA) with a microphone for auditory input.

Biorobotics more or less died when Grey Walter moved on to other research in the 1950s. With the rise of AI and computing, the focus was on providing robots with human-inspired perception and cognition. The new robots had a series of modules, such as visual processing, planning, and reasoning, through which sensory information passed serially. Typically, a decision-making module controlled the output to the actuators. This was in contrast to the more direct control approach of Grey Walter, in which the only mediation between sensing and moving was provided by an artificial neural net consisting of two hardware neurons. Another difference was that whereas AI robotics focused on human cognition, Grey Walter focused on the question of how seemingly complex animal-like behavior could arise from simple mechanisms such as tropisms and reflexes.

Today the term *taxis* is used instead of tropism to refer to the movement of an animal directed by a stimulus, either negatively or positively. Examples of such stimulus-directed activity include chemotaxis (chemical taxis), geotaxis (gravity), phototaxis (light), and phonotaxis (auditory). Although Grey Walter worked only on individual taxes, biologists at the time (e.g., Fraenkel and Gunn; cf. Sharkey and Ziemke in Ziemke and Sharkey, 1998, pp. 361–392) proposed that the behavior of many organisms could be explained by a combination of taxes working together and in opposition. They cited Fraenkel’s study of the coastal slug, *Littorina neritoides*. *Littorina* combines positive and negative phototaxis with negative geotaxis to feed and survive. Subsequently, combinations of taxes have been used as powerful explanations of many animal behaviors, from bacteria feeding to insect pheromone trailing to fish breeding and feeding.

These ideas began to emerge in a new wave of robotics during the 1980s as a result of two major influences. First, the neuroanatomist Valentino Braitenberg showed how a number of complex behaviors could emerge from a combination of very simple neural networks encoding different taxes (Braitenberg, 1984). Second, Rodney Brooks’s development of subsumption architecture allowed autonomous control by a combination of taxes, and drove home the effectiveness of behavior-based robotics. His major papers from this period are reprinted in Brooks (1999). In this style of robotics, each behavior-producing module, such as *avoid obstacles* or *move toward light*, is encoded as a separate program module such that each is directly under the control of environmental circumstances rather than a central controller. For example, when there is light on the sensors, the *move toward light* module will be active until the light is occluded by an obstacle, at which point the *avoid obstacles* module takes over.

Current Directions in Biorobotics

A large emerging body of research in robotics is making the connection between sensing and moving simple, and the relationship between robot and world tightly coupled. It was the dramatic increase in robotics research, riding on the back of the new behavior-based approach, that enabled biologically inspired robotics to flourish. Since the behavior-based approach grew directly from ideas in the life sciences, it was only natural that once the tools and techniques of the approach had been developed, they would be turned back to work on the source of inspiration.

In this article, biorobotics is divided into three main classes. Although these classes are mutually supportive and their paths often cross, the distinctions between them are nonetheless useful.

- The *generalized* approach follows from the lineage of ideas that inspired Grey Walter to use robots to investigate and extend general mechanistic theories of animal behavior and adaptation. This includes research using neural network adaptation through learning and/or evolutionary methods (see REACTIVE ROBOTIC SYSTEMS).
- The *specific* approach uses methods from the generalized approach to investigate specific species or organisms. The research can range from studies of the physical plausibility of a simple neural explanation for some target behavior pattern to the physical modeling of a particular animal or some of its senses. Models can be evaluated by observing the target behavior of the robot interacting with the environment through sensing and moving. One of the main goals of specific biorobotics is to develop new methods for scientific modeling.
- The *theoretical* division is a mixed bag that provides an examination of the implications of the research for a number of disciplines. The issues range widely, from discussions of robot embodiment to the nature of life. Although all biorobotics has a theoretical component, the theoretical approach is distinct in not requiring empirical work.

The idea was to include only work that at least touched base with the life sciences with respect to the type of controllers and the method of adaptation used. Each of the classes is dealt with in more depth in the following three subsections.

Generalized Biorobotics

The research impetus is to use broad notions derived from the life sciences for robot control. Many of these notions are in the form of implicit assumptions, such as deriving complex behavior from the simplest possible mechanisms or using the ideas of taxis or tropism for automated control. In this sense, Grey Walter’s research was prototypical generalized biorobotics. His work on classical conditioning with the CORA architecture also foresaw the modern focus on adaptive techniques in robotics. The biological currency in the generalized biorobotics community mostly consists of abstract models of neural network learning, animal learning, or evolutionary processes, or a mixture. In the next two subsections, the main trends will be discussed.

Evolution. Evolutionary methods have been used for many applications since the 1950s when the first Genetic Algorithm (GA; see EVOLUTION OF ARTIFICIAL NEURAL NETWORKS) was developed by Friedman for his master’s thesis on evolving control circuits for autonomous robots. These methods are particularly useful for constraining search in very large search spaces. However, from the perspective of biorobotics, the most important reason for employing evolutionary methods is that they are abstractly related to the Darwinian principle of natural selection and may be seen as analogous to real evolutionary theory; i.e., there is a fitness function to decide how fit a particular program is in the context of the problem it is to solve, and there are mutation and crossover to operate on the computer equivalent of gene strings. Given the intended relationship between the behavior of biorobots and natural biological behavior, the development of an *evolutionary robotics* is a very important step.

A fairly typical example of evolutionary methods for single robots is Nolfi’s garbage collector (in Sharkey, 1997, pp. 187–198). The connection weights were evolved to control a miniature robot equipped with distance sensors and a gripper. The task was to

“clean” an arena by picking up objects and dropping them off outside. To do this, the robot had to move around the arena, avoid obstacles, locate an object, pick it up, move toward the walls, and release the object outside the arena. After 1,000 generations, robot controllers were evolved that performed the cleaning task to a high degree of accuracy.

Most evolutionary robotics research relies on using a fixed neural network architecture on which the weights are evolved. Another interesting approach is to let the evolutionary method decide on the type of connectivity between the units in the net; i.e., the pattern of connectivity is “genetically” represented (see Husbands et al. in Ziemke and Sharkey, 1998, pp. 185–210).

An important research area in biorobotics is concerned with how the environment and other species co-evolve with a given organism, resulting in an *evolutionary arms race*. This issue has been taken up in the simple form of evolving two competing robot controllers at the same time. For example, Floreano and Nolfi (1997) co-evolved the controllers for predator and prey behavior for two different “species” as part of each other’s environment. One of the main problems was that in one generation the predators would win but in the next generation the prey would win because a counterstrategy was evolved. This instability has been overcome by introducing neural network learning during the lifetime of the individuals. In this way the predators were able to adapt to the new evolved strategies of the prey.

The approach of combining the two adaptive techniques of evolutionary methods and neural network learning is proving to be a very effective adaptation technique that has a naturalistic flavor. Much of the research on combining has focused on how learning can help guide evolution—the *Baldwin effect*. The idea is that if the genotype of an individual is close to an optimal combination of genes, learning can allow that individual to increase its suitability for its environment, thereby increasing its probability of survival and reproduction. This could lead to a larger “basin” of fitness around optimal genotypes, channeling evolution toward optimal solutions (see Nolfi and Floreano, 1999, for a review).

Learning. One of the most widely used learning techniques in biorobotics is reinforcement or reward learning (RL) (see, e.g., Krose, 1995). RL has been studied psychologically since the beginning of the twentieth century. An advantage of RL techniques in robotics is that the learner needs only occasional reinforcement. RL is therefore unlike supervised learning, which requires a trainer to provide the learner with an exact target action in every time step, suited for use in unknown environments or tasks (but see Sharkey, 1998, on the use of innate controllers for training supervised learning).

More recently there has been a move toward using the *operant conditioning* techniques developed in the 1940s for studies of animal learning. Operant conditioning involves the shaping of pre-given behaviors. In particular, animals can be trained to produce an experimenter-required behavior when they are rewarded for successive approximations to that behavior. For example, to begin training a rat to press a bar for food, rewards are given for any reaching movement. Then successive approximations to the goal are rewarded until the target behavior is observed. In robotics, this has also been called behavior editing by Dorigo and Colombetti (1998), who have conducted most of the experimental work on this technique. An extension of this work to include incremental shaping is discussed by Urzelai et al. (in Ziemke and Sharkey, 1998, pp. 341–360).

In a realistic approach, Saksida et al. (in Sharkey, 1997, pp. 231–249) successfully used operant conditioning to modify the interaction between behaviors that had been preprogrammed into a robot. This departure from using reinforcement learning as a trial-and-error approach to modify existing behaviors is a step toward

real animal training. Furthermore, unlike most RL work, the training was conducted by a human trainer rather than a programmed reinforcer. Initially, the robot has three categories of objects: a bright orange jacket, green and pink plastic dog toys, and blue plastic recycling bins. One of its innate behaviors was to approach the plastic dog toys and pick them up. Successful (fast) shaping was shown for a number of new behaviors, including *Follow the Trainer*, *Recycling*, and *Playing Fetch*.

Specific Biorobotics

One of the attractions of robotics is that there is strong potential for testing the relationship between a model and some hypothesized behavioral consequences in the physical world. When Hammond built his heliotrope in the early twentieth century to test Loebian theory, it was essentially the physical plausibility of the theory that was under scrutiny. This was generalized biorobotics in that the hypotheses were about all animals. One of the goals of specific biorobotics is to extend such physical testing to test specific hypotheses about specific species. The motivation is that mathematical specification and computer simulation provide only a weak test of a model in that the inputs are typically chosen by the researcher and the outputs are designed to be interpretable as data points or graphs. The central idea of specific biorobotics is to test the model by situating the robot in a physical environment that provides the main features of the world of the target species.

There are a number of dangers with this approach, and a number of wrinkles will have to be ironed out before such modeling reaches maturity as a test methodology in biology and psychology. For example, with complex neural networks such as brains, it is not always possible to isolate a mechanism and test its behavioral consequences. Although robotics can offer a window on the possible behaviors resulting from particular models, a model cannot generally be used directly as a robot controller; a number of “gaps” between the sensors, the model, and motor output have to be filled in. This can be advantageous in forcing the theorist to extend the theoretical mechanisms, but care must be taken to ensure that mechanisms outside the theoretical framework do not play a causal role in the robot behavior.

Robotic modeling of living systems has taken a number of different forms, from behavioral modeling (see, e.g., Webb in Gauthier, 1996, pp. 117–134, on cricket phonotaxis; Grasso et al. in Chang and Gadiano, 2000, pp. 115–131 on lobster chemotaxis) to neuroscientific modeling (e.g., Burgess et al. in Ziemke and Sharkey, 1998, pp. 291–300, and Recce et al. in Sharkey, 1997, pp. 393–406, on the rat hippocampus; van der Smagt in Ziemke and Sharkey, 1998, pp. 301–320, on the human cerebellum for arm control) to modeling animal sensing (e.g., Lambrinos et al. in Chang and Gadiano, 2000, pp. 39–64, on ant solar compass sensing; Blanchard et al. in Chang and Gadiano, 2000, pp. 17–38, on locust sensing of approach; Rucci in Chang and Gadiano, 2000, pp. 181–193, on localization of auditory and visual structures in the barn owl) to biomechanics (e.g., Delcomyn and Nelson in Chang and Gadiano, 2000, pp. 5–15, and Quinn and Ritzman in Ziemke and Sharkey, 1998, pp. 239–254, on hexapod walking in the cockroach).

One of the most successful attempts at behavioral modeling has been the work of Webb and her associates (e.g., Webb in Gauthier, 1996, on mate selection in the female cricket). A wheeled robot was used to physically model a female cricket locating a conspecific male by following its calls. The robot was equipped with an auditory system capable of selectively localizing the sound of a male cricket stridulating (rubbing its wings together rapidly to produce a sound that attracts potential mates).

A similar approach has been taken by Lambrinos et al. (in Chang and Gadiano, 2000, pp. 39–64) for modeling the sensors of the

desert ant *Cataglyphus*, which maintains its heading across a largely featureless desert using polarized light sensing. Lambrinos et al. built special-purpose polarized light sensors based on what is known about the neural mechanisms of polarization that the honey bee *Apis mellifera*, the field cricket *Gryllus campestris*, and the desert ant *Cataglyphus bicolor* use to determine the position of the sun. The sensors were mounted on a wheeled robot and used to test different models of how *Cataglyphus* maintains its heading with polarized light. The research has been successfully conducted on a mobile robot in the ant's natural habitat with a homing performance similar to that of the ant.

As in Webb's work, the "ant robot" was used to model only a small part of the whole process of finding the direction to the nest. It did not, for example, accommodate the movement of the sun across the sky during the day (although this information was used to make corrections to the data). The sun moves relative to Earth at an average of 15° per hour (this figure varies greatly according to the time of day). In the early part of the twentieth century, this fact was used to show that ants both memorized the position of the sun and compensated for its movement. When the ants are imprisoned in a dark box for 2½ hours and released, they deviate from their original bearing by approximately the same number of degrees as the sun moved during their imprisonment. These findings reveal that *Cataglyphus* keeps track of the azimuth during the day and uses this information in maintaining a course.

Another important aspect of robotics used for modeling concerns legged locomotion. This leads to a two-way interaction between model testing and engineering. A number of researchers have turned to insect locomotion as a way to find a type of gait for a legged robot. Quinn and Ritzmann (in Ziemke and Sharkey, 1998, pp. 239–254) have designed and built a hexapod robot based on detailed neurobiological and kinematic observations of the locomotion of the death's head cockroach, *Blaberus discoidalis*. As a result, the robot's kinematics are remarkably similar to those of the real cockroach, and issues addressed in controlling the artificial cockroach have actually led to new understanding of its natural counterpart.

Moving onto the mammalian nervous system, Burgess, Donnett, and O'Keefe (in Ziemke and Sharkey, 1998, pp. 291–300) used a miniature mobile robot equipped with a camera to test a neuronal model of how internal and external sensory information contribute to the firing of place cells in the rat hippocampus, and how these cells contribute to rat navigation behavior. They tested hypotheses based on their earlier neurophysiological work on the rat hippocampus using single-cell recording techniques. The robot experiments showed that the information provided by the robot's on-board video, odometry, and proximity sensors was sufficient to allow reasonably accurate return to an unmarked goal location. Similar robot modeling work has also been carried out by Recce et al. (in Sharkey, 1997, pp. 393–406) using the hippocampus as a method of absolute localization.

Research in specific biorobotics is gathering momentum as robot and sensing technology continues to improve. There are still many modeling issues to be worked out in conjunction with biology. The next step would be to get the morphology of robots to more accurately model the bodies and movement of the target species and to work continuously toward the goal of modeling whole animals, rather than installing patches to cover the missing bits. Like computational modeling, great care must be taken to ensure that the patches do not have a causal role in the target behavior.

Theoretical Biorobotics

Theoretical biorobotics is the most abstract level of biologically inspired robotics. Essentially, theoretical biorobotics is an all-

encompassing category for work that does not involve implementation on a robot but rather addresses metaquestions about robotics. Although wide-ranging, the main theoretical focus of biologically inspired robotics concerns biological and psychological issues. Many of these issues draw on detailed philosophical reasoning. Here we will confine ourselves to setting out some of the main points and referencing more detailed works in the literature.

A strong impetus for the new wave in biologically inspired robotics was the way it differed from traditional AI. Rodney Brooks, one of the prime movers in the mid-1980s, was concerned with the inadequacy of the prevailing methods used in AI for robotics (see, e.g., Brooks, 1999). Based mainly on the cognitivist conception of human intelligence, the sensory input to robots went through a number of strategic stages such as perception, planning, and reasoning before each move. All of the information was presented to a central controller, which decided how to act. This slowed performance to a single small move about every 15 minutes.

Rejecting cognitivism, theoretical biorobotics views intelligence as *embodied* in the machine and in its interactions with the world in which it is situated. Extreme cognitivists hold that mind is essentially a computer program, a language of thought, that could be run on any machine capable of running it. Mind is simply linked to the machine running it and the external world through transducers. Extreme bioroboticians might claim that mind is inseparable from the individual machine and the more encompassing environmental machine of which the individual machine is a part. That is, the robot is *situated* in the world and is an *embodied* or *physically grounded* intelligence.

Varela, Thompson, and Rosch (1991) provide an insightful discussion of the details of embodiment in robots and its relation to life and mind. These authors are primarily interested in how living systems are embodied and how they are situated in their interactions with the world. Their purpose is to urge cognitive science to reject the vacuity of ungrounded thought. However, Sharkey and Ziemke (in Ziemke and Sharkey, 1998), while going along with some of the account by Varela et al. of living systems, argue for a weak embodiment in robotics, i.e., that robots can be used to *model* embodiment without themselves being embodied (see PHILOSOPHICAL ISSUES IN BRAIN THEORY AND CONNECTIONISM).

Another idea that has received considerable attention is that of *emergence* or *emergent behavior*. This is the notion that we can get something for nothing (or very little). One analogy is that from a collection of many molecules of water a cloud emerges that is greater than the sum of the parts. Perhaps a better example is the emergence of collective behavior in insects when each insect carries out very simple behaviors. For example, it is argued that the extraordinary structures that termites build in the desert emerge from very simple behaviors. Clark (1997) provides an in-depth discussion of emergent behavior and describes the two rules required by the termites: "If not carrying anything and you bump into a wood chip, pick it up"; and "If carrying a wood chip and you bump into another one, put it down." The resultant piling behavior emerges from the interplay between simple rules and the constraints of the environment.

The idea, then, is that coherent behavior emerges from a collection of simple taxes working together at the same time. This was an outright rejection of the notion of a central controller for action that was prevalent in AI. The idea in AI was to provide the robot with a model of the world, whereas one of the favorite slogans of the new roboticists is "the world is its own model." Nonetheless, Sharkey and Ziemke (2001) caution that even the taxes are emergent in the sense that they are in the eye of the beholder; i.e., they are distal descriptions of behavior.

One of the healthiest signs in the field is that some mainstream biologists and psychologists have begun to write about the relationship between specific biological findings and robotics. Navi-

gation, for example, is an important topic in both biology and robotics, and biologists (e.g., Collett in Ziemke and Sharkey, 1998, pp. 255–270; Etienne in Ziemke and Sharkey, 1998, pp. 271–290; Franz and Mallot in Chang and Gaudio, 2000, pp. 133–153) have discussed the relation between different aspects of navigation from an insect and mammalian perspective. Moreover, psychologists are beginning to take robot studies using animal learning techniques seriously enough to write detailed discussions of the relationship between the natural and the metallic (e.g., Savage in Ziemke and Sharkey, 1998, pp. 321–340).

Conclusions

The field of biologically inspired robotics has been classified into the three separate subfields of generalized, specific, and theoretical. General and theoretical biorobotics has a long but patchy history that is now a considerable and growing field. Specific biorobotics has gradually emerged from the other two and is fast making headway toward the goal of accurately modeling specific animal species. With the ever-increasing improvements in materials, sensors, and computing equipment, we can look forward to many exciting new developments over the coming decade and the transfer of the findings into engineering.

Road Map: Robotics and Control Theory

Related Reading: Arm and Hand Movement Control; Neuroethology, Computational; Potential Fields and Neural Networks; Reactive Robotic Systems

References

- Braitenberg, V., 1984, *Vehicles: Experiments in Synthetic Psychology*, Cambridge, MA: MIT Press. ♦
- Brooks, R., 1999, *Cambrian Intelligence: The Early History of the New AI*, Cambridge, MA: MIT Press.
- Chang, C., and Gaudio, P., Eds., 2000, *Biomimetic Robotics*, *Robot. Auton. Syst.*, 31(1–2):1–218 (special issue).
- Clark, A., 1997, *Being There: Putting Brain, Body and World Together Again*, Cambridge, MA: MIT Press.
- Dorigo, M., and Colombetti, M., 1998, *Robot Shaping: An Experiment in Behavior Engineering*, Cambridge, MA: MIT Press.
- Floreano, D., and Nolfi, S., 1997, Adaptive behaviour in competing co-evolving species, in *Proceedings of the Fourth European Conference on Artificial Life* (P. Husbands and I. Harvey, Eds.), Cambridge, MA: MIT Press.
- Gaussier, P., Ed., 1996, *Moving the Frontiers Between Robotics and Biology*, *Robot. Auton. Syst.*, 16:107–362 (special issue).
- Grey Walter, W., 1953, *The Living Brain*, New York: Norton. ♦
- Krose, B., Ed., 1995, Special issue on reinforcement learning and robotics. *Robot. Auton. Syst.*, 15:233–340.
- Nolfi, S., and Floreano, D., 1999, Learning and evolution, *Auton. Robots*, 7:89–113.
- Sharkey, N., Ed., 1997, *Robot Learning: The New Wave*, *Robot. Auton. Syst.*, 22(3–4):135–274 (special issue). ♦
- Sharkey, N., 1998, Learning from innate behaviors: A quantitative evaluation of neural network controllers, *Auton. Robots*, 5:317–334.
- Sharkey, N., and Ziemke, T., 2001, Mechanistic vs. phenomenal embodiment: Can robot embodiment lead to strong AI? *Cognit. Syst. Res.*, 2:251–262.
- Varela, F., Thompson, E., and Rosch, E., 1991, *The Embodied Mind: Cognitive Science and Human Experience*, Cambridge, MA: MIT Press.
- Ziemke, T., and Sharkey, N., Eds., 1998, *Biorobotics*, *Connect. Sci.*, 10(3–4):161–360 (special issue).