

Behavior-Based Robotics as a Tool for Synthesis of Artificial Behavior and Analysis of Natural Behavior

Maja J Matarić

Computer Science Department and Neuroscience Program

University of Southern California

941 West 37th Place, SAL 228

Los Angeles, CA 90089-0781

tel: (213) 740-4520, fax: (213) 740-7285

mataric@cs.usc.edu

(This paper appeared in Trends in Cognitive Science, Vol. 2, No. 3, March 1998, 82-87.)

Work in behavior-based systems focuses on functional modeling, that is, the synthesis of life-like and/or biologically-inspired behavior that is robust, repeatable, and adaptive. Inspiration from cognitive science, neuroscience, and biology drives the development of new methods and models in behavior-based robotics, and the results tie together several related fields, including Artificial Life, evolutionary computation, and multi-agent systems. Ideas from Artificial Intelligence and engineering continue to be actively explored and applied to behavior-based robots as their role in animal modeling and practical applications is being developed.

This paper overviews behavior-based robotics, a field of robotics that is guided by principles from nature and aims to develop methods for synthesizing artificial systems, ranging from physical robots to autonomous software agents, and use robotics to model and analyze natural systems, ranging from insects to humans. We define the key principles of behavior-based robotics, and overview a variety of examples of its practical applications and its models of natural systems. We describe the overlap between the two, but focus on biologically-inspired robotics work, giving details of several new areas of research.

Behavior-based robotics is a branch of robotics that bridges Artificial Intelligence, Engineering, and Cognitive Science. Its dual goals are 1) to develop methods for controlling artificial systems, ranging from physical robots to simulated ones and other autonomous software agents, and 2) to use robotics to model and better understand biological systems, typically animals ranging from insects to humans. This paper focuses largely on work toward satisfying the first goal, gives a brief review of the key approaches and types of systems that have been implemented with a strong biological inspiration.

In the field of robotics, “control architectures” are methodologies that supply structure and impose constraints on the way robots are controlled. The behavior-based approach is a methodology for designing robot architectures and controllers for endowing robots with intelligent behavior. The methodology is based on a biologically-inspired philosophy that favors parallel, decentralized architectures, and allows for some freedom of interpretation. The approach is general and fits well within other powerful frameworks such as schema theory [4].

In behavior-based systems, the robot controller consists of a collection of “behaviors”, each of which achieves and/or maintains a specific goal. For example, the “avoid-obstacles” behavior maintains the goal of preventing collisions with objects in the environment, and the “go-home” behavior achieves the goal of reaching some home region. Each behavior is a processing element or a procedure, also called a control law in the engineering field of Control Theory. that can be implemented either in software or hardware; each can take inputs from the robot’s sensors (e.g., camera, ultrasound, infra-red, tactile) and/or from other behaviors, and send outputs to the robot’s effectors (e.g, wheels, grippers, arms, speech) and/or to other behaviors in the system. Consequently, a behavior-based robot is controlled by a structured network of interacting behaviors.

System Organization

The organizational methodology of behavior-based systems differs from other robot control methods in its approach to modularity, that is, the way in which the system is organized and subdivided. The behavior-based philosophy mandates that the behaviors be relatively simple, incrementally added to the system, and not executed in a serial fashion. The systems are meant to be constructed in a bottom-up fashion resembling evolution in its incremental refinement as well as its utilitarian exploitation of existing modules.

Behaviors are activated in response to external and/or internal conditions, sensory inputs and internal state. The system as a whole activates entire subsets of behaviors so that parallelism can be exploited, both in speed of computation and in the resulting dynamics. The latter is a critical aspect of behavior-based control: as multiple behaviors or modules are active, dynamics of interaction arise both within the system itself (from the interaction among the behaviors) and within the environment (from the interaction of the behaviors with the external world). Inspired by biological organisms, designers of behavior-based systems exploit these dynamics to create (by hand or automatically through the use of learning, described below) repeatable, stable, and, ultimately, intelligent behavior without relying on top-down, centralized, and often even hierarchical control [3, 2, 13].

The Design of Behaviors

A methodological constraint of behavior-based systems is their use of state and representation: information is not centralized or centrally manipulated. Instead, various forms of distributed representations are used, ranging from static table structures and networks, to active, procedural processes, providing a rich medium for innovative interpretations.

Behaviors can be designed at a variety of levels of abstraction. In general, they are made to be higher than the robot’s atomic actions (i.e., typically above “go-forward-by-a-small-

increment”, “turn-by-a-small-angle”), and they extend in time and space. This effectively elevates the representational level of the system, which has been shown to facilitate higher-level cognition and learning [53, 59]. Some commonly implemented behaviors include: “go-home”, “find-object”, “get-recharged”, “avoid-collisions”, “pick-up-object”. More specialized behaviors include: “avoid-the-light”, “aggregate-with-group”, “find-mate”, “follow-edge”, etc.

The internal behavior structure of a system need not necessarily mirror its externally manifested behavior. For example, a robot that flocks with other robots may not have a specific internal “flocking” behavior; instead, its interaction with the environment and other robots may produce flocking. Behavior-based systems are typically designed so the effects of the behaviors interact in the environment rather than internally through the system, so as to take advantage of the richness of the interaction dynamics. These dynamics are sometimes called “emergent” because they result from the interactions and are not internally specified by the robot’s program [77].

Coordinating Multiple Behaviors

A key issue in behavior-based systems concerns the coordination of the multiple behaviors, thus making “arbitration”, i.e., deciding what behavior to execute at each point in time, one of the central challenges. For the sake of simplicity, most implemented systems use a built-in, fixed priority ordering of behaviors. More flexible solutions, which can be less computationally efficient and harder to analyze, have been suggested, commonly based on selecting a behavior by computing some function of the behavior activation levels, such as voting or activation spreading [48, 72].

A Historical Overview

Behavior-based systems were founded on the work in reactive robotics and in particular on the Subsumption Architecture [12], which achieves rapid real-time responses by embedding the robot’s controller into a collection of preprogrammed parallel condition-action rules, or reflexes, with minimal internal state (e.g., “if bumped, stop”) [15, 3]. In contrast to these so-called bottom-up systems, traditional AI deliberative planner-based systems are top-down, and require the robot to perform a serial sequence of processing sense-plan-act steps (e.g., “combine the sensory data into a model of the world, then use the planner to find a path in the model, then send each of the steps of the plan to the robot’s wheels”) [37, 66, 45]. Hybrid systems attempt a compromise between the “thinking” and “acting” extremes by employing a reactive system for low-level control and a planner for higher-level decision making [30, 36, 8, 73, 51, 22]. Hybrid systems tend to separate the control system into two or more communicating but largely independent parts. Behavior-based systems are an alternative to hybrid systems; they enable fast real-time responses through simple reactive behaviors that directly link sensors and effectors, but also provide for higher-level deliberation, by distributing the representation and computation over more sophisticated concurrent behavior processes. The power, elegance, and complexity of behavior-based systems all stem from the ways in which their constituent behaviors are defined and used. Some proponents of behavior-based systems claim that they better model cognition, while others employ them purely from

pragmatic motivations, including their ease of system development and the robustness of the results.

Learning and Adaptation

Learning has been called the hallmark of intelligence; thus achieving adaptive and learning capabilities in artificial systems is one of the greatest challenges of Artificial Intelligence. Learning is particularly difficult in robotics, because sensing and acting in the physical world involves a great deal of uncertainty due to incomplete and noisy information and dynamically changing environment conditions. It is often difficult for robots to correctly perceive (due to limited sensory technology) and act on (due to limited effectors) the variety of situations that arise in the physical world. Nonetheless, robot learning is an active branch of robotics, and is one of the variations and adaptations of standard machine learning techniques, and in particular reinforcement learning, that have been effectively applied to robots. In particular, behavior-based robots have been demonstrated learning to walk [49], navigate [54, 63, 74], communicate [81], divide tasks [71], behave socially [56], and even identify opponents and score goals in robot soccer [9].

Reinforcement learning is the most popular method for learning in mobile robotics [50, 42]. It refers to a set of problems (rather than methods), in which the robot must improve its behavior based on rewards or punishment from the environment. The reinforcement learning model is based on early conditioning work in psychology, and recently an increasing number of robot learning systems have employed related concepts from biological reinforcement learning, most notably shaping [57, 59, 27] and operant conditioning [35, 78]. Supervised learning methods using neural networks have also been used extensively [10, 74, 34]. Some of the most effective demonstrations of learning in mobile robots have been inspired by biological learning systems.

Demonstrations and Applications

Consistent with the dual goal of the field itself, behavior-based robotics has been used both in practical applications and in futuristic exploratory endeavors. Behavior-based robots have demonstrated various standard robotic capabilities, including obstacle avoidance, navigation, terrain mapping, following, chasing/pursuit, object manipulation, task division and cooperation, and learning maps, and walking. Application domains have included mobile robots, underwater vehicles, space robotics (most recently Sojourner, the robot that autonomously explored the surface of Mars), as well as robots capable of manipulation, grasping, walking, and running.

Interestingly, as the behavior-based approach is being explored for modeling natural systems, the resulting research is demonstrating methods with immediate practical applications. For example, models of large-scale group behavior, that have been developed with behavior-based systems, appear to be the methodology of choice for deploying robots in hazardous or inaccessible environments, including under water, in mine fields, under ground, and in space. These domains require a combination of individual independence and group cohesion capable of adapting to varying group sizes and organizations, as it does in natural societies. Conse-



Figure 1: A group of 13 flocking mobile robots. These robots also demonstrated wandering, following, aggregation, dispersion, chain-formation, and foraging.

quently, behavior-based approaches have presented popular options for addressing analysis of natural behavior and synthesis of practical artificial behavior.

Faithful simulation and robotic models have been developed of a variety of natural behaviors, ranging from reflexive behavior selection strategies [21], cricket phonotaxis for flight and mating behaviors [79], lobster odor location [38], fly [33] and hover-fly [20] vision, to insect navigation, trail formation, and path-finding [25, 26], the application of the schema theory to modeling navigation [5] and frog behavior [6], the use of evolutionary computation methods, modeled after natural selection, to develop individual robotic behaviors [68] as well as group behaviors [1] and many others. A bi-annual international conference, “Simulation of Adaptive Behavior” is devoted to the subject, and has been convening since 1990; its proceedings provide a more representative review of the various activities in this field [62, 61, 19, 47].

An Example of Biologically-Inspired Navigation

Our own work has used behavior-based systems to model navigation, map learning, and path finding mechanisms loosely modeled on the rat’s hippocampal place cells [53]. In this system, the behaviors served not only for general movement and obstacle avoidance, but also for landmark detection and representation. As new landmarks were discovered, behaviors became associated with them, and subsequently became activated whenever the robot returned to the same location, much like hippocampal place cells in rats [70, 28, 32, 69]. Unlike its neural counterpart, whose network topology has no obvious mapping to the

physical space it represents, our synthetic navigation system maintained a clear isomorphic mapping between the two. Consequently, the resulting robot-generated maps were easily readable by humans interacting with the robot. The landmark behaviors also served as predictors that allowed the robot to localize more precisely in its environment; an active landmark used context to activate its network neighbor in the robot’s direction of travel, thus “priming” it, and generating expectation. A lack of expectation indicated novel locations or incorrect localization.

An Example of Ethologically-Inspired Group Behavior

Our more recent work has focused on using behavior-based systems to model group behavior. Inspired by ethologically common natural behaviors, we have used groups of up to 13 robots to demonstrate aggregation, dispersion, following, flocking, foraging, task division, specialization, and dominance hierarchy formation (see Figure 1). The robotic implementations resemble the equivalent behaviors found in species ranging from ants, crabs, and chickens to chimps and humans [18, 55], but are not designed to be careful mechanistic models. Instead, they serve as demonstrations of possible mechanisms which push the state of the art in robotics, as well as allow us to postulate theories about their inspirations from nature.

Current Developments: Biologically-Inspired Imitation

Our most recent work with behavior-based systems extends the control spectrum, from planar mobile robots to articulated, anthropomorphic bodies. Again inspired by certain neuroscience theories of motor control (e.g., [11, 67], which demonstrate evidence of a finite set of additive force fields controlling the movement repertoire of frogs and rats), we are developing behaviors for the control of three-dimensional movement. As with our work on group behavior, we are using “basis behavior” primitives as a substrate for a broader repertoire of higher-level behaviors, obtained by sequencing and combining the basis set [58]. Our current basis set includes behaviors for movement to point, posture maintenance, and oscillatory movements, all based on theories of human motor control [31, 65, 7, 17], with the eventual goal of modeling *learning by imitation*.

Acquiring new skills by imitation is a well known robotics problem. It is usually classified as learning by demonstration, where the robot uses vision sensors to interpret the behavior of a human user and thus acquire a new task. Historically, assembly tasks have been learned by demonstration, without an effort to precisely model the observed behavior but focusing on achieving the demonstrated goals [40, 44, 46, 39]. More recently, imitation learning between two mobile robots has been demonstrated [24], as has skill learning between a human demonstrator and an articulated robot arm with human kinematics, such as learning to balance a pole [76] and play Kendama [64]. The latter was an instantiation of the bi-directional theory of motor control [43], another example of a robotic implementation of a neuroscience model. Modeling human skill acquisition, while a tremendously challenging task, is gaining popularity in robotics. It has recently been approached from a Piagetian perspective, using developmental stages in order to simplify the complex learning problem [16, 52, 80]. Computational modeling of motor control and learning is an active research

area outside the scope of this paper.

Inspired by data from neuroscience and psychophysics that provide evidence for combined perceptual and motor activation during movement observation and imagination [75, 23, 29, 41], we are developing a set of behaviors that not only produce movement, but also facilitate its perception. These behavior primitives simultaneously recognize and plan movements, thus combining perception and generation. Furthermore, the primitives facilitate prediction, in that they represent complete movements and when presented with an incomplete visual input, can complete it based on their own model of the movement. Practically, the system functions by continuously classifying the observed movements it into its known repertoire, thus enabling imitation.

This approach is an extension of our earlier hippocampus-inspired navigation work, described above, in which landmark behaviors primed and anticipated other landmarks based on their local topology. Similarly, in the current work partial movement matches recognize and prime complete behaviors. Inspired by developmental psychology work providing evidence for infant prediction of goals implied in observed incomplete and incorrect actions [60], our primitives infer movement goals by internally matching, predicting, and completing the observed movements. The ultimate goal of this work is dual, in a manner typical of behavior-based work: it is to 1) provide insight into the animal/human imitation process, and to 2) facilitate automated programming of new tasks and skills in robotic systems.

Outstanding Questions

To estimate the pace of development of behavior-based systems, see [14], which provides an in-depth overview of the foundations, goals, and the state of this interdisciplinary field in 1991. While major accomplishments have been made, many interesting and important research questions remain to be addressed. Among them:

- What organizational principles will be needed in order to scale up behavior-based systems to truly complex problems, such as those involving multiple types of very different tasks and goals, akin to animal behavioral repertoires?
- How well will behavior-based systems scale to increasingly more cognitive problems, such as those involving symbolic reasoning including natural language discourse?
- Can researchers in natural and artificial sciences overcome their stylistic, terminological, and methodological differences and collaborate more closely and efficiently in order to better utilize their complementary expertise?

Acknowledgments

This work is supported by the Office of Naval Research Grant N00014-95-1-0759, the National Science Foundation Career Grant IRI-9624237 and Infrastructure Grant CDA-9512448. The author is grateful to Paolo Gaudiano for his extensive insightful comments on an earlier draft of this paper.

References

- [1] Arvin Agah and George A. Bekey. Phylogenetic and ontogenetic learning in a colony of interacting robots. *Autonomous Robots*, 4:85–100, 1997.
- [2] Philip E. Agre. *The Dynamic Structure of Everyday Life*. PhD thesis, MIT, October 1988.
- [3] Philip E. Agre and David Chapman. Pengi: An implementation of a theory of activity. In *Proceedings, AAAI-87*, pages 268–272, Seattle, WA, 1987.
- [4] Michael Arbib. Perceptual structures and distributed motor control. In V. B. Brooks, editor, *Handbook of Physiology: Motor Control*, pages 809–813. The MIT Press, 1981.
- [5] Michael Arbib. From visual affordances in monkey parietal cortex to hippocampoparietal interactions underlying rat navigation. *Phil. Trans. Roy. Soc of London, B*, 1997.
- [6] Michael A. Arbib. Visuomotor coordination: Neural models and perceptual robotics. In *Visuomotor Coordination: Amphibians, Comparisons, Models and Robots*, pages 121–171. Plenum Press, 1989.
- [7] Michael A. Arbib, N. Schweighofer, and W. T. Thatch. Modeling the cerebellum: From adaptation to coordination, in motor control and sensory-motor integration: Issues and directions. In *Motor Control and Sensory-Motor Integration: Issues and Directions*, pages 11–36. North-Holland Elsevier Science, 1995.
- [8] Ronald C. Arkin. Towards the unification of navigational planning and reactive control. In *AAAI Spring Symposium on Robot Navigation*, pages 1–5, March 1989.
- [9] Minoru Asada, Eiji Uchibe, Shoichi Noda, Sukoya Tawaratsumida, and Koh Hosoda. Coordination of multiple behaviors acquired by a vision-based reinforcement learning. In *Proceedings, IEEE/RSJ/GI International Conference on Intelligent Robots and Systems*, Munich, Germany, September 1994.
- [10] A. A. Baloch and A. M. Waxman. Visual learning, adaptive expectations, and behavioral conditioning of the mobile robot mavin. *Neural Networks*, 4:271–302, 1991.
- [11] Emilio Bizzi, Ferdinando A. Mussa-Ivaldi, and Simon Giszter. Computations underlying the execution of movement: A biological perspective. *Science*, 253:287–291, 1991.
- [12] Rodney A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, RA-2:14–23, April 1986.
- [13] Rodney A. Brooks. Artificial life and real robots. In *Toward A Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, pages 3–10. MIT Press, 1991.
- [14] Rodney A. Brooks. Intelligence without reason. In *Proceedings, IJCAI-91*, pages 569–595, Sydney, Australia, 1991.

- [15] Rodney A. Brooks and Jonathan H. Connell. Asynchronous distributed control system for a mobile robot. In *SPIE's Cambridge Symposium on Optical and Optoelectronic Engineering*, pages 77–84, Cambridge, MA, 26–31 October 1986.
- [16] Rodney A. Brooks and Lynn Andrea Stein. Building brains for bodies. *Autonomous Robots*, 1:7–25, 1994.
- [17] Vernon Brooks, Frank Hilperath, Martin Brooks, Hans-Georg Ross, and Hans-Joachim Freund. Learning what and how in a human motor task. *Learning & Memory*, 2:225–242, 1995.
- [18] Ivan D. Chase and Sievert Rohwer. Two methods for quantifying the development of dominance hierarchies in large groups with application to harris' sparrows. *Animal Behavior*, 35:1113–1128, 1987.
- [19] D. Cliff, P. Husbands, J-A. Meyer, H. Roitblat, and eds. S. Wilson. *From Animals to Animats: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. The MIT Press, 1994.
- [20] Dave Cliff. The computational hoverfly; a study in computational neuroethology. In J-A. Meyer and S. Wilson, editors, *From Animals to Animats: International Conference on Simulation of Adaptive Behavior*, pages 87–96. MIT Press, 1990.
- [21] Jonathan H. Connell. *Minimalist Mobile Robotics: A Colony Architecture for an Artificial Creature*. Academic Press, Boston, 1990.
- [22] Jonathan H. Connell. Sss: A hybrid architecture applied to robot navigation. In *IEEE International Conference on Robotics and Automation*, pages 2719–2724, Nice, France, 1991.
- [23] Jeanne Decety. Do imagined and executed actions share the same neural substrate? *Cognitive Brain Research*, 3:87–93, 1996.
- [24] John Demiris and Gillian Hayes. Imitative learning mechanisms in robots and humans. In Volker Klingspor, editor, *Proceedings of the 5th European Workshop on Learning Robots*, pages 9–16, Bari, Italy, 1996.
- [25] Jean L. Deneubourg, S. Aron, Simon Goss, J. M. Pasteels, and G. Duernick. Random behaviour, amplification processes and number of participants: How they contribute to the foraging properties of ants. *Physica 22D*, pages 176–186, 1986.
- [26] Jean L. Deneubourg, Simon Goss, J. M. Pasteels, D. Fresneau, and J. P. Lachaud. Self-organization mechanisms in ant societies, ii: Learning in foraging and division of labor. In *From Individual to Collective Behavior in Social Insects*, volume 54, pages 177–196. 1987.
- [27] Marco Dorigo and Marco Colombetti. *Robot Shaping: An Experiment in Behavior Engineering*. MIT Press, Cambridge, MA, 1997.

- [28] Howard Eichenbaum and Neal J. Cohen. Representation in the hippocampus: What do hippocampal neurons code? *Trends in Neuroscience*, 11(6), 1988.
- [29] L. Fadiga, L. Fogassi, G. Pavesi, and G. Rizzolatti. Motor facilitation during action observation: a magnetic stimulation study. *J. Neurophysiol.*, 73(6):2608–2611, 1995.
- [30] R. James Firby. An investigation into reactive planning in complex domains. In *Proceedings, Sixth National Conference on Artificial Intelligence*, pages 202–206, Seattle, 1987.
- [31] Tamar Flash and Nevile Hogan. The coordination of the arm movements: an experimentally confirmed mathematical model. *Journal of Neuroscience*, 7:1688–1703, 1985.
- [32] Tom C. Foster, Carl A. Castro, and Bruce L. McNaughton. Spatial selectivity of rat hippocampal neurons: Dependence on preparedness for movement. *Science*, pages 1589–1592, June 1989.
- [33] N. Franceschini. Sampling of the visual environment by the compound eye of the fly. In A. Snyder and R. Menzel, editors, *Photoreceptor Optics*, pages 98–125. Springer, Berlin, 1975.
- [34] Paolo Gaudiano and Carolina Chang. Neural competitive maps for reactive and adaptive behavior. In *IEEE International Symposium on Computational Intelligence in Robotics and Automation*, pages 13–18, Piscataway, NJ, 1997. IEEE Press.
- [35] Paolo Gaudiano, Zalama Eduardo, Carolina Chang, and Juan Lopez Córonado. A model of operant conditioning for adaptive obstacle avoidance. In Pattie Maes, Maja J Matarić, J.-A. Meyer, Jordan Pollack, and Stewart Wilson, editors, *From Animals to Animats: International Conference on Simulation of Adaptive Behavior*, pages 373–381. MIT Press, 1996.
- [36] Michael P. Georgeff and Amy L. Lansky. Reactive reasoning and planning. In *Proceedings, Sixth National Conference on Artificial Intelligence*, pages 677–682, Seattle, 1987.
- [37] G. Giralt, Raja Chatila, and M. Vaisset. An integrated navigation and motion control system for autonomous multisensory mobile robots. In Michael Brady and R. Paul, editors, *First International Symposium on Robotics Research*, pages 191–214. MIT Press, Cambridge, MA, 1983.
- [38] F. Grasso, T. Consi, D. Moutain, and J. Atema. Locating odor sources in turbulence with a lobster-inspired robot. In Pattie Maes, Maja J Matarić, J.-A. Meyer, Jordan Pollack, and Stewart Wilson, editors, *From Animals to Animats: International Conference on Simulation of Adaptive Behavior*, pages 101–112. MIT Press, 1996.
- [39] Geir E. Hovland, Pavan Sikka, and Brenan J. McCarragher. Skill acquisition from human demonstration using a hidden markov model. In *Proceedings, IEEE International Conference on Robotics and Automation*, pages 2706–2711, Minneapolis, MN, 1996.

- [40] Katsushi Ikeuchi, Takashi Suehiro, Peter Tanguy, and Mark Wheeler. Assembly plan from observation. Technical report, Carnegie Mellon University Robotics Institute Annual Research Review, 1990.
- [41] Marc Jeannerod and Jean Decety. Mental motor imagery: a window into the representational stages of action. *Current Opinion in Neurobiology*, 5:727–732, 1995.
- [42] Leslie Pack Kaelbling, Michael L. Littman, and Andrew W. Moore. Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4:237–285, 1996.
- [43] Jurgen Konczak and Esther Thelen. Bi-directional theory approach to integration. In T. Inui and J. McClelland, editors, *Attention and Performance XVI*, pages 335–367. The MIT Press, Cambridge, Massachusetts, 1996.
- [44] Y. Kuniyoshi, M. Inaba, and H. Inoue. Learning by watching: extracting reusable task knowledge from visual observation of human performance. *IEEE Transactions on Robotics and Automation*, 10(6):799–822, 1994.
- [45] John E. Laird and Paul S. Rosenbloom. Integrating, execution, planning, and learning in soar for external environments. In *Proceedings, AAAI-90*, pages 1022–1029, 1990.
- [46] Christopher Lee and Yangsheng Xu. Online, interactive learning of gestures for human/robot interfaces. In *Proceedings, IEEE International Conference on Robotics and Automation*, volume 4, pages 2982–2987, Minneapolis, MN, 1996.
- [47] P. Maes, M. Matarić, J-A. Meyer, J. Pollack, and eds. S. Wilson. *From Animals to Animals: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*. The MIT Press, 1996.
- [48] Pattie Maes. The dynamics of action selection. In *IJCAI-89*, pages 991–997, Detroit, MI, 1989.
- [49] Pattie Maes and Rodney A. Brooks. Learning to coordinate behaviors. In *Proceedings, AAAI-90*, pages 796–802, Boston, MA, 1990.
- [50] Sridhar Mahadevan and Leslie P. Kaelbling. The national science foundation workshop on reinforcement learning. *AI Magazine*, 17(4):89–97, 1996.
- [51] Chris Malcolm and Tim Smithers. Symbol grounding via a hybrid architecture in an autonomous assembly system. In Pattie Maes, editor, *Designing Autonomous Agents*, pages 145–168. MIT Press, 1990.
- [52] Matthew Marjanovic, Brian Scassellati, and Matthew Williamson. Self-taught visually-guided pointing for a humanoid robot. In Pattie Maes, Maja Matarić, Jean-Arcady Meyer, Jordan Pollack, and Stewart Wilson, editors, *Fourth International Conference on Simulation of Adaptive Behavior*, pages 35–44, Cape Cod, MA, 1996. MIT Press.

- [53] Maja J. Matarić. Navigating with a rat brain: A neurobiologically-inspired model for robot spatial representation. In J-A. Meyer and S. Wilson, editors, *From Animals to Animats: International Conference on Simulation of Adaptive Behavior*, pages 169–175. MIT Press, 1990.
- [54] Maja J. Matarić. Integration of representation into goal-driven behavior-based robots. *IEEE Transactions on Robotics and Automation*, 8(3):304–312, June 1992.
- [55] Maja J. Matarić. Kin recognition, similarity, and group behavior. In *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society*, pages 705–710, Boulder, Colorado, 1993. Lawrence Erlbaum Associates.
- [56] Maja J. Matarić. Learning to behave socially. In D. Cliff, P. Husbands, J-A. Meyer, and S. Wilson, editors, *From Animals to Animats: International Conference on Simulation of Adaptive Behavior*, pages 453–462, 1994.
- [57] Maja J. Matarić. Reward functions for accelerated learning. In William W. Cohen and Haym Hirsh, editors, *Proceedings of the Eleventh International Conference on Machine Learning (ML-94)*, pages 181–189, New Brunswick, NJ, 1994. Morgan Kaufman Publishers, Inc.
- [58] Maja J. Matarić. Designing and understanding adaptive group behavior. *Adaptive Behavior*, 4(1):50–81, December 1995.
- [59] Maja J. Matarić. Reinforcement learning in the multi-robot domain. *Autonomous Robots*, 4(1):73–83, 1996.
- [60] Andrew N. Meltzoff and M. Keith Moore. Infants’ understanding of people and things: From body imitation to folk psychology. In Jose Luis Bermudez, Anthony Marcel, and Naomi Eilan, editors, *The Body and the Self*, pages 44–69. MIT Press/Bradford Books, 1995.
- [61] J-A. Meyer, H. Roitblat, and eds. S. Wilson. *From Animals to Animats: Proceedings of the Second International Conference on Simulation of Adaptive Behavior*. The MIT Press, 1993.
- [62] J-A. Meyer and eds. S. Wilson. *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*. The MIT Press, 1990.
- [63] José Del R. Millán. Learning reactive sequences from basic reflexes. In *Proceedings, Simulation of Adaptive Behavior SAB-94*, pages 266–274, Brighton, England, 1994. MIT Press.
- [64] Hiroyui Miyamoto, Stefan Schaal, Francesca Gandolfo, Hiroaki Gomi, Yasuharu Koike, Rieko Osu, Eri Nakano, Yasuhiro Wada, and Mitsuo Kawato. A kendama learning robot based on bi-directional theory. *Neural Networks*, 9(6), 1996.
- [65] Pietro Morasso. Spatial control of arm movements. *Experimental Brain Research*, 42:223–227, 1981.

- [66] Hans P. Moravec and Alberto Elfes. High resolution maps from wide angle sonar. In *IEEE International Conference on Robotics and Automation*, St. Louis, MO, March 1985.
- [67] Ferdinando A. Mussa-Ivaldi and Simon Giszter. Vector field approximation: a computational paradigm for motor control and learning. *Biological Cybernetics*, 67:491–500, 1992.
- [68] Stefano Nolfi, Dario Floreano, Orazio Miglino, and Francesco Mondada. Now to evolve autonomous robots: Different approaches in evolutionary robotics. In *Proceedings, Artificial Life IV*, pages 190–197. The MIT Press, 1994.
- [69] John O’Keefe. Computations the hippocampus might perform. In Lynn Nadel, Lynn A. Cooper, Peter Culicover, and R. Michael Harnish, editors, *Neural Connections, Mental Computation*, pages 225–284. MIT Press, Cambridge, MA, 1989.
- [70] John O’Keefe and Lynn Nadel. *The Hippocampus as a Cognitive Map*. Oxford University Press, Oxford, 1978.
- [71] Lynne E. Parker. An experiment in mobile robotic cooperation. In *Proceedings, Robotics for Challenging Environment*, Albuquerque, New Mexico, 1993.
- [72] D. Payton, D. Keirse, D. Kimble, J. Krozel, and K. Rosenblatt. Do whatever works: A robust approach to fault-tolerant autonomous control. *Journal of Applied Intelligence*, 3:226–249, 1992.
- [73] David Payton. Internalized plans: a representation for action resources. In Pattie Maes, editor, *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back*, pages 89–104. MIT Press, 1990.
- [74] Rolf Pfeifer and Paul Verschure. Distributed adaptive control: a paradigm for designing autonomous agents. In F. J. Varela and P. Bourguine, editors, *Toward a Practice of Autonomous Systems*, pages 21–30. MIT Press, 1992.
- [75] Giacomo Rizzolatti, Luciano Fadiga, Vittorio Gallese, and Leonardo Fogassi. Premotor cortex and the recognition of motor actions. *Cognitive Brain Research*, 3:131–141, 1996.
- [76] Stefan Schaal. Learning from demonstration. In M.C. Mozer, M. Jordan, and T. Petsche, editors, *Advances in Neural Information Processing Systems 9*, pages 1040–1046. The MIT Press, 1997.
- [77] Luc Steels. The artificial life roots of artificial intelligence. *Artificial Life*, 1(1):75–110, 1994.
- [78] David S. Touretzky and Lisa M. Saksida. Operant conditioning in skinnerbots. *Adaptive Behavior*, 5(3/4):219–247, 1997.
- [79] Barbara Webb. Robotic experiments in cricket phonotaxis. In *From Animals to Animals: International Conference on Simulation of Adaptive Behavior*. MIT Press, 1993.

- [80] Matthew Williamson. Postural primitives: Interactive behavior for a humanoid robot arm. In Pattie Maes, Maja Matarić, Jean-Arcady Meyer, Jordan Pollack, and Stewart Wilson, editors, *Fourth International Conference on Simulation of Adaptive Behavior*, pages 124–131, Cape Cod, MA, 1996. MIT Press.
- [81] Holly Yanco and Lynn Andrea Stein. An adaptive communication protocol for cooperating mobile robots. In J-A. Meyer, H. Roitblat, and S. Wilson, editors, *From Animals to Animats: International Conference on Simulation of Adaptive Behavior*, pages 478–485. MIT Press, 1993.