# Behavior-Based Primitives for Articulated Control


**Maja J. Matarić*** **Matthew Williamson** **John Demiris** **Aswath Mohan***

* Computer Science Department
University of Southern California
Los Angeles, CA 90089-0781
mataric@cs.usc.edu

MIT Artificial Intelligence Lab
545 Technology Square
Cambridge, MA 02139, USA
matt@ai.mit.edu

Department of Artificial Intelligence
University of Edinburgh
Edinburgh, EH1 2QL, UK
johnde@dai.ed.ac.uk



## Abstract

Our previous work has demonstrated how biologically-inspired behaviors can serve as an effective substrate for control, representation, and learning in mobile robots and multi-robot systems, in order to generate adaptive individual and group behavior. In this paper, we expand the behavior-based approach to the domain of manipulator control, and demonstrate first results in applying biological inspiration to a 20 degree-of-freedom humanoid dynamical torso simulation.


## 1. Introduction

Behavior-based systems take inspiration from biology, ethology, and neuroscience, in order to construct controllers for agents situated in noisy and dynamic environments (Matarić 1997). Behaviors impose a distributed, bottom-up approach to control, which eliminates bottlenecks and enables emergent functionality (Steels 1994). Aside from its pragmatic role in achieving distributed, real-time control in robotics, the behavior-based framework has also been used to model and implement simulations of biological systems. In this paper we present the use of behavior-based control on a novel domain, where its properties are used both to model its biological inspiration and to use its principles to achieve more robust control.

The domain we focus on is three-dimensional articulated movement. In contrast, the vast majority of behavior-based systems to date have been applied to planar navigation problems. We present a generalization of behavior-based control to the domain of manipulator control, which has been largely under the umbrella of traditional robotics and geometric reasoning (Paul 1981, Brady, Hollerbach, Johnson, Lozano-Perez & Mason 1982). In contrast, we take inspiration from neuroscience and describe an approach toward implementing a general articulated motion controller gbased on distributed behaviors. We demonstrate the first stages of our approach on a 20 degree-of-freedom humanoid torso simulation with dynamics. Given the complexity of the problem, we consider different control schemes, and describe a combination of approaches used so as to choose the most appropriate approach for a given movement task.

## 2. Motivation

Biological motion provides a challenging model for control of articulated robots and graphical agents. Our work, both in navigation and now in manipulation, is inspired by a specific principle derived from evidence in neuroscience, which proposes an elegant organization of the frog and rat motor systems. Giszter, Mussa-Ivaldi & Bizzi (1993), Mussa-Ivaldi, Giszter & Bizzi (1994), and related work on spinalized frogs and rats suggests the existence of force-field motor primitives that produce complete high-level motor behaviors such as reaching and wiping. This evidence has been demonstrated by stimulating sections of the spinal cord, and measuring the resultant forces in the foot, which converge to a single equilibrium point, i.e., come to rest in the same position. Additionally, Mussa-Ivaldi et al. (1994) show that stimulating two or more regions in the frog spinal cord results in either a linear superposition of the corresponding force fields or a "winner-take-all" response that suppresses the other force fields.

This data suggest an elegant organizational principle for motor control, in which entire behaviors are coded with low-level force-fields, and can be combined into higher-level, more complex behaviors. The idea of supplying an agent with a collection of *basis behaviors* or *primitives* to generate motion, and combining those into a general repertoire for complex motion, is very appealing. Our previous work (Matarić 1995, Matarić 1997), inspired by the same biological results, has already successfully applied the idea to control of planar mobile robots. This paper extends the notion to agents with more DOFs, not by implementing a careful model of the frog force-fields, but rather by using the biological evidence as inspiration.

## 3. Experimental Testbed

We used a dynamic anthropomorphic simulation environment for implementing and testing our ideas. Adonis is a rigid-body simulation of a human torso, without an articulated spine, and with static legs (Figure 1).[1] The simulation consists of eight rigid parts or links: the head,

[1] Adonis was developed by Jessica Hodgins at Georgia Tech, using SD/Fast and models purchased from Viewpoint labs.

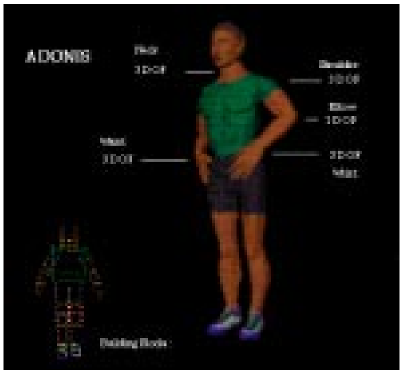


Figure 1 The 20 degree-of-freedom (DOF) physics-based torso simulation testbed. In the experiments presented here, 14 DOFs were used.

torso, left upper arm, right upper arm, left lower arm, right lower arm, left hand and right hand. The head is connected to the torso with a rotary three degree of freedom (DOF) joint, as is the torso to the waist. In the experiments described here, we used only the DOFs in the arms, consisting of three DOF rotary joints at the shoulders and wrists, and one DOF pin joints at the elbows. Thus, Adonis has a total of 20 DOFs of which we used 14. The testbed presents a very complex control problem and introduces novel challenges for the behavior-based framework. Much like a physical robot, the simulation software provides no built-in collision detection and avoidance capabilities, so the controller must resolve the many issues involved with self- and other collisions.

## 4. The Behavior Primitive Approach

Designing a basis behavior set is a difficult problem, especially in a high-dimensional space. We drew inspiration and constraints from the human motor psychophysics literature, and selected a small initial set of general primitives with the goal of exploring their effectiveness, and subsequently using learning techniques to tune their parameters and generate their combinations. The following primitives were used:

1. **move-to-point** Move the end-point to a Cartesian goal position, implemented using impedance control.
2. **get-posture** Achieve a posture by adjusting the angles of the arm to reach the desired angles.
3. **avoid** Avoid collisions, was implemented based on generalized potential fields.

The implementation of each is described in turn.

### *4.1 Move-to-Point by Impedance Control*

The general principle of impedance control (Hogan 1985) is to modulate the mechanical impedance of the end-point of an arm, by altering the torques at the arm's joints. Mechanical impedance for an object is defined as the relationship between an imposed disturbance and a generated force. For example, a compressed spring exerts a force proportional to the displacement. The impedance of such a system is constant and equal to the stiffness of the spring. For a more complicated mechanism like a robot arm, the mechanical impedance is determined by the control at the joint level. A mechanical arm can be made to appear as if a virtual spring and damper are connected to some equilibrium point; moving the point will drag the arm around, and the arm will automatically return to its equilibrium position if disturbed. Arranging the control of the arm in this way has advantages in terms of stability, especially when interacting with different environments (Colgate & Hogan 1989).

We implemented an impedance-based controller as follows. The force $F$ from the "virtual" spring and damper is:

$$F = K(x_e - x) - B\dot{x} \tag{1}$$

where $x$ is the 6-D vector defining the position and orientation of the end-point (hand) in space, $x_e$ is the 6-D vector of desired positions and orientations, and $K$ and $B$ are stiffness and damping matrices. This desired force is implemented by applying torques ($\tau$) at the joints, using the following simple relation (Craig 1989):

$$\tau = J^T F \tag{2}$$

The method has several advantages over position control techniques using inverse kinematics (Baker & Wampler II 1988). It is computationally simple, requiring only the forward kinematics and the Jacobian (Whitney 1982), and it is stable both when moving freely, and during contact with surfaces (Hogan 1985). It also fits well into the primitive-based approach; rather than use inverse kinematics for all joints, which would be very computationally intensive, a mixture of simpler control schemes can be used. For example, if **move-to-point** impedance control is used on the arms, and **get-posture** joint control on the body, the hands will remain in the same positions as the torso is moved.

To improve the robustness of the system, two additional terms were implemented. First, a term was added to the torque calculation to compensate for the effect of gravity on the limbs, so improving the tracking accuracy of the hand. Second, damping was applied at the shoulder to reduce redundant arm motions. With a small amount of damping at the shoulder joints, the arm moves smoothly without additional internal movements.

The addition of these terms makes the torque equation:

$$\tau = J^T F + G + B_s \tag{3}$$

where $G$ is the gravity compensation, and $B_s$ the added damping term.

The impedance controller was used to effectively implement the **move-to-point** behavior. The user simply specifies the desired end-point and orientation of the arm in Cartesian coordinates, and the impedance controller calculates the torques to move the arm to that position.

### 4.2 Get-Posture by Joint Angle Control

Many articulated movements involve achieving (and holding) postures, such as gestural communication, dance, etc., which are most easily defined in terms of joint angles. This capability is implemented in the **get-posture** behavior. PD servo control is used at the joints, and the desired angles are set manually by the user. The latter is a drawback of the method, however it allows the user to specify whole-arm position (rather than just the end-point in **move-to-point**). Although we did not attempt to implement the motor control that might be used by frogs, **get-posture** is a reasonable approximation to a frog spinal field. The shape of the frog's force fields remains mostly unchanged when the animal is deafferented (Loeb, Giszter, Borghesani & Bizzi 1993), indicating that the field is largely generated by the spring-like muscles in the leg, acting through the leg kinematics. In Adonis, the analogue of the muscle springs are the PD servos of **get-posture**, which also generate a restoring force field at the hand.

Note that **move-to-point** and **get-posture** behaviors are duals of each other, since every posture involves a specific Cartesian location of the end-point, and every movement to a point results in a set of angles for the wrist, elbow, and shoulder. However, task specifications typically vary between the two; some tasks are naturally presented in Cartesian space while others are best described in joint angles. The task we chose featured both types of specifications.

### 4.3 Avoid by Using Virtual Forces

Detecting and avoiding collisions is a crucial part of control, especially in systems where self-collisions are possible. To prevent such collisions, we applied repulsive forces from obstacles to each moving arm. The arms are constructed out of three point-mass building blocks, so the force is applied from the center of mass of the obstacle along a straight line to the centers of mass of the three blocks of the arm. Manz, Liscano & Green (1991) used repulsive force inversely proportional to the computed distance and directly proportional to the velocity of the approaching object. We adapted this approach, because our task required reaching targets that were frequently very close to the obstacle in 3D space, so the repulsive force had to be reduced after a certain proximity was acheived. This can be visualized as a force-field with a peak at a certain distance from the obstacle. Formally:

$$F = \frac{V\,S}{e^{D^2}} \tag{4}$$

where $F$ is the applied repulsive force, $V$ is the velocity, $D$ is the distance, and $S$ is the scaling factor. The force was computed for and applied to each component of the arm, as described above. Note that this is added to the internally generated forces of Adonis' dynamical simulation, and serves to deflect the arms' trajectories away from obstacles and other body parts (the head being the most relevant in our task).

This avoidance approach is successful when the initial distance between the body part and the obstacle is sufficiently large. When combined with other behaviours, it can generate movement trajectories around obstacles and towards goals. However, the approach is not sufficient for generating smooth trajectories around obstacles. In fact, while direct reaching in biological systems is known to typically proceed along a straight Cartesian path (Flash & Hogan 1985), little is know about trajectory generation around obstacles. Most theories involve the use of carefully chosen intermediate via points (Miyamoto, Schaal, Gandolfo, Gomi, Koike, Osu, Nakano, Wada & Kawato 1996) which produce a smooth path around the obstacle. We are working on a similar implementation of a generic "curved-trajectory-generator".

### 4.4 The Task and Results

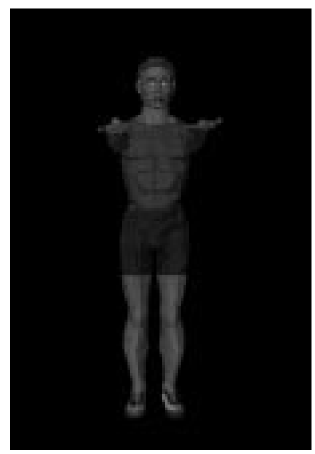

Figure 2 The start position for achieving sub-task five.

For the first demonstration of our approach, we chose a well-known motor task, the popular dance "Macarena", which involves a sequence of upper-body movements, whose order is well defined and success of achievement is easily evaluated. The goal behind our approach is to present a high-level, rather abstract description of the task, and have the motor behaviors execute the appropriate primitives to reach each of the sub-tasks. We used the following verbal specification of the Macarena to define the sub-tasks (sub-goals) for the motor primitives:

1. Extend right arm out, palm down
2. Extend left arm out, palm down
3. Rotate right hand (palm up)
4. Rotate left hand (palm up)
5. Touch right hand to top of left shoulder
6. Touch left hand to top of right shoulder
7. Touch right hand to back of the head
8. Touch left hand to back of the head

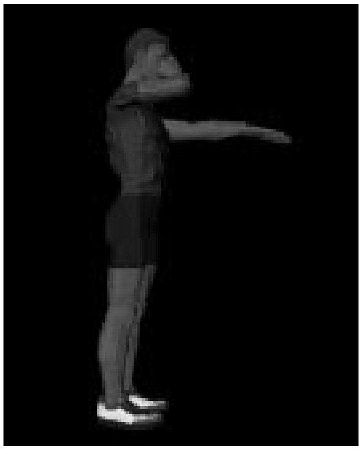

Figure 3 An example of an illegal solution based only on **get-posture** behavior.

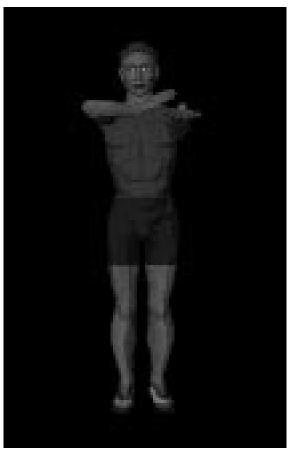

Figure 4 Successful achievement of subtask five by blending **get-posture** and **avoid** behaviors.

9. Touch right hand to the left side of the chest
10. Touch left hand to the right side of the chest
11. Put right hand to the right hip
12. Put left hand to the left hip

The Macarena task is sufficiently complex to test our approach; it involves the control of all of the 14 degrees of freedom in the two arms, and many of the movements involve the collision avoidance system. Furthermore, the task involves both Cartesian and joint-space specifications. The first four sub-tasks are best described in terms of postures (e.g., "sticking the arm out" does not aim at a Cartesian end position), and the remaining eight are best envisioned in Cartesian space.

A direct combination or blending of **avoid** with the other behaviors (much like in our previous work on *flocking* (Matarić 1992)) produces legal and correct (i.e., goal-achieving) movements. However, the resulting trajectory is not optimal and may be quite unnatural. For example, the achievement of sub-goal 5 requires the arm to move to the shoulder, which involves contacting the body but not colliding with it, while avoiding any contact with the head. A simple conservative addition of the two behaviors results in a trajectory that moves the arm first away from the head and toward the right shoulder, and then back around the face to the goal on the left shoulder. A small amount of parameter tuning results in a smoother trajectory. Unfortunately, the behavior tuning is specific to the shoulder-head configuration and does not generalize to other self-collisions. There are two ways of addressing this issue: assuming that self-collisions are critical enough to merit special-purpose fine-tuned avoidance strategies (as may be the case in biological movement (Arbib, Schweighofer & Thatch 1995)), or developing a higher-level smoothing behavior, as done in 2D path planning.

Figure 2 above demonstrates the state of the system at the beginning of executing subtask 5; Figure 3 demonstrates the illegal solution, without the use of collision-avoidance; Figure 4 shows the same transition, with blending of the **avoid** behavior to generate a legal movement. Movies of Adonis performing the Macarena, are available at http://www-robotics.usc.edu/~agents/macarena.html.

One of the benefits of the behavior decomposition of control systems is not only that there are different ways of structuring a given system (i.e., different collections of behaviors can solve the problem), but also that once a behavior decomposition is achieved, the specific behavior controllers can themselves vary, depending on the available sensors and effectors. To validate this point, we experimented with two different types of **move-to-point** behaviors: impedance control described above, and force-field control. The latter approach, described in detail in Matarić, Zordan & Mason (1998), applied forces to the end point of the arm to move it to the desired Cartesian location. The approach was limited since it imposed no constraints on the other DOFs of the arm, causing the resulting motion to resemble a rag doll; for example in sub-tasks one and two, the arm would reach the goal, but the elbow would swing about. Consequently, we opted for impedance control as a more general approach to generating movement in Cartesian space.

The current implementation of the Macarena does not address the question of sequencing the behaviors; we assume the existence of a task-level plan. Our current research is addressing the production of such a plan automatically, directly from a continuous stream of movement data. Our preliminary experiments with human movement data have found salient cues based on the velocity profiles of the movements and zero-velocity breaks at the sub-goals. The issue of sequencing is addressed to the extent that **get-posture** and **move-to-point** are mutually exclusive in the Macarena task, and our implementation awaits termination of the current behavior before staring the next. Our task does not require concurrent execution of two behaviors on the different arms, but many more complex tasks that would require this capability. The behavior-based framework we have implemented lends itself to scaling up to such coordination; it is not difficult to imagine how additional behaviors that constrain and guide coordinated two-arm movement can be added to the current system, and this is an active direction of our work.

## 5. Related Work

Manipulator control is a mature and active area of robotics, and a comprehensive review is well beyond the scope of this paper. Instead, we review only the most related approaches, both in behavior-based control in general, and in manipulation in specific. Manipulator work largely focuses on the difficult problems of inverse kinematics and dynamics (Paul 1981, Brady et al. 1982), and learning techniques are often employed to overcome the high dimensionality of the problem (Atkeson 1989, Schaal & Atkeson 1994). By using generic behaviors, we hope to develop "stereotyped" parametric solutions to common motor control problems, such as reaching to a point and achieving postures, so that a motor system can perform less run-time computation (i.e., avoid solving the find-path problem (Lozano-Pérez 1982) for each new movement task) without losing most of its generality. Previous work by Williamson (1996) is inspired by the same motivation; it demonstrates the control of a 6-DOF physical robot arm based on primitive postures. The control system uses four primitives (three reaching and one rest posture) combined using linear interpolation. This scheme was also used by Marjanovic, Scassellati & Williamson (1996) to learn to reach toward a visual target. Their work and ours are both based on the same biological inspiration, but their system uses postural primitives while we use movement behaviors instead.

Behavior-based control has been used extensively in navigation; (Arkin 1987) introduced motor schemas, which are akin to behaviors, for mobile robot control, and related methods for combining them. The behaviors we discuss are based on the same principles as the generic **homing** behavior used in Matarić (1995), which moves a mobile robot toward a specified 2D Cartesian goal. Our go-to-point behavior is a more complex 3D generalization. A simple special-purpose behavior-based arm controller was implemented by Connell (1988), for controlling a 3 DOF arm mounted on a mobile robot and used for soda can collection.

Sensorimotor primitives have been successfully applied to assembly tasks. Malcolm & Smithers (1990) used a hybrid system to generate a sequence of grasps in order to perform an assembly task with an ADEPT manipulator. Unlike theirs, our current system does not yet address the sequence generation process. Furthermore, while their execution module generates a specific path for each arm movement, our underlying behaviors are designed to be more general purpose. Morrow, Nelson & Khosla (1995) use primitives for translation, rotation, guarded moves, and grasping sequenced by a planner, and Nelson, Jägersand & Fuentes (1995) extend it to model "virtual tools" for grasping and assembly. The emphasis in this work is building skill libraries that can be robustly used by a planner. In contrast, our work is aimed at computationally simpler gross motion primitives.

## 6. Extensions and Conclusion

A key direction of our continuing work is in the application of the above primitives to the perceptual side of manipulation, in particular within an imitation learning framework (Kuniyoshi, Inaba & Inoue 1994, Demiris, Rougeaux, Hayes, Berthouze & Kuniyoshi 1997). Again inspired by biological data, which indicate that perceptual and motor systems may share some representational substrates, we have been working on adapting our motor primitives into sensori-motor behaviors, capable of not only generating but also recognizing typical movements. We give the agent a set of innate motor behaviors, and allow it to learn through imitation by observing other agents' movement, and attempting to map those onto its own set of primitives, and internally (and externally, if needed) imitate them, thereby generating a learning signal.

The goal of the work described in this paper has been to take inspiration from biological motor primitives and apply it to the high-dimensional manipulator control problem. Based on our previous work with basis behaviors, we designed a set of motor primitives and used the to produce a non-trivial sequential motor task, the Macarena, using 14 DOFs of a dynamic torso simulation. We demonstrate how behavior-based techniques of behavior sequencing and blending can be directly applied to this new and more complex control domain, and continue to refine our approach to begin to tackle more sophisticated control problems.

## Acknowledgments

This work was supported by the NSF Career Grant IRI-9624237 to M. Matarić. The authors thank Victor Zordan for help with the details of Adonis.

## References

Arbib, M. A., Schweighofer, N. & Thatch, W. T. (1995), 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', North-Holland

Elsevier Science, pp. 11–36.

Arkin, R. C. (1987), Motor Schema Based Navigation for a Mobile Robot: An Approach to Programming by Behavior, *in* 'IEEE International Conference on Robotics and Automation', Raleigh, NC, pp. 264–271.

Atkeson, C. G. (1989), 'Learning Arm Kinematics and Dynamics', *Annual Review of Neuroscience* **12**, 157–183.

Baker, D. R. & Wampler II, C. W. (1988), 'On the Inverse Kinematics of Redundant Manipulators', *International Journal or Robotics Research* **7**(2), 3–21.

Brady, M., Hollerbach, J. M., Johnson, T. L., Lozano-Perez, T. & Mason, M. T., eds (1982), *Robot Motion: Planning and Control*, The MIT Press, Cambridge, Massachusetts.

Colgate, E. & Hogan, N. (1989), An analysis of Contact Instability in Terms of Passive Equivalents, *in* 'Proceedings of the 1989 IEEE International Conference on Robotics and Automation', pp. 404–409.

Connell, J. H. (1988), A Behavior-Based Arm Controller, MIT AI Memo 1025, MIT Artificial Intelligence Lab, Cambridge, MA.

Craig, J. J. (1989), *Introduction to Robotics: Mechanics and Control*, second edn, Addison-Wesley, Reading, Massachusetts.

Demiris, J., Rougeaux, S., Hayes, G. M., Berthouze, L. & Kuniyoshi, Y. (1997), Deferred Imitation of Human Head Movements by an Active Stereo Vision Head, *in* 'Proceedings of the 6th IEEE International Workshop on Robot Human Communication (RoMan97)', IEEE Press, Sendai, Japan.

Flash, T. & Hogan, N. (1985), 'The coordination of the arm movements: an experimentally confirmed mathematical model', *Journal of Neuroscience* **7**, 1688–1703.

Giszter, S. F., Mussa-Ivaldi, F. A. & Bizzi, E. (1993), 'Convergent force fields organized in the frog's spinal cord', *Journal of Neuroscience* **13**(2), 467–491.

Hogan, N. (1985), 'Impedance Control: An Approach to Manipulation', *Journal of Dynamic Systems, Measurement, and Control* **107**, 1–24.

Kuniyoshi, Y., Inaba, M. & Inoue, H. (1994), 'Learning by watching: extracting reusable task knowledge from visual observation of human performance', *IEEE Transactions on Robotics and Automation* **10**(6), 799–822.

Loeb, E. P., Giszter, S. F., Borghesani, P. & Bizzi, E. (1993), 'Effects of Dorsal Root Cut on the Forces Evoked by Spinal Microstimulation in the Spinalized Frog', *Somatosensory and Motor Research* **13**(2), 467–491.

Lozano-Pérez, T. (1982), Task Planning, *in* M. Brady, J. M. Hollerbach, T. L. Johnson, T. Lozano-Pérez & M. T. Mason, eds, 'Robot Motion: planning and control', MIT Press, Cambridge, MA, pp. 473–488.

Malcolm, C. & Smithers, T. (1990), Symbol grounding via a hybrid architecture in an autonomous assembly system, *in* P. Maes, ed., 'Designing Autonomous Agents', MIT Press, pp. 145–168.

Manz, A., Liscano, R. & Green, D. (1991), A comparison of real-time obstacle avoidance methods for mobile robots, *in* 'Proceedings, Experimental Robotics II, the 2nd International Symposium', Touluse, France.

Marjanovic, M., Scassellati, B. & Williamson, M. (1996), Self-Taught Visually-Guided Pointing for a Humanoid Robot, *in* P. Maes, M. Matarić, J.-A. Meyer, J. Pollack & S. Wilson, eds, 'Fourth International Conference on Simulation of Adaptive Behavior', The MIT Press, Cape Cod, MA, pp. 35–44.

Matarić, M. J. (1992), Designing Emergent Behaviors: From Local Interactions to Collective Intelligence, *in* J.-A. Meyer, H. Roitblat & S. Wilson, eds, 'From Animals to Animats: International Conference on Simulation of Adaptive Behavior'.

Matarić, M. J. (1995), 'Designing and Understanding Adaptive Group Behavior', *Adaptive Behavior* **4**(1), 50–81.

Matarić, M. J. (1997), 'Behavior-Based Control: Examples from Navigation, Learning, and Group Behavior', *Journal of Experimental and Theoretical Artificial Intelligence* **9**(2–3), 323–336.

Matarić, M. J., Zordan, V. B. & Mason, Z. (1998), Movement Control Methods for Complex, Dynamically Simulated Agents: Adonis Dances the Macarena, *in* 'Autonomous Agents', Minneapolis, St. Paul, MI.

Miyamoto, H., Schaal, S., Gandolfo, F., Gomi, H., Koike, Y., Osu, R., Nakano, E., Wada, Y. & Kawato, M. (1996), 'A Kendama Learning Robot Based on Bi-Directional Theory', *Neural Networks*.

Morrow, J. D., Nelson, B. J. & Khosla, P. K. (1995), Vision and Force Driven Sensorimotor Primitives for Robotic Assembly Skills, *in* 'Proceedings of the IEEE International Workshop on Intelligent Robots and Systems (IROS-95)', Vol. 3, Pittsburg, PA, pp. 234–240.

Mussa-Ivaldi, F. A., Giszter, S. F. & Bizzi, E. (1994), 'Linear combinations of primitives in vertebrate motor control', *Proceedings of the National Academy of Sciences* **91**, 7534–7538.

Nelson, R. C., Jägersand, M. & Fuentes, O. (1995), Virtual Tools: A Framework for Simplifying Sensory-Motor Control in Complex Robotic Systems, *in* 'Proceedings of the IEEE International Workshop on Vision for Robots (IROS-95)', Pittsburg, PA, pp. 70–79.

Paul, R. P. (1981), *Robot Manipulators: Mathematics, Programming, and Control*, The MIT Press, Cambridge, Massachusetts.

Schaal, S. & Atkeson, C. C. (1994), 'Robot Juggling: An Implementation of Memory-Based Learning', *Control Systems Magazine* **14**, 57–71.

Steels, L. (1994), 'The Artificial Life Roots of Artificial Intelligence', *Artificial Life* **1**(1), 75–110.

Whitney, D. E. (1982), The mathematics of coordinated control of prosthetic arms and manipulators, *in* M. Brady, J. M. Hollerbach, T. L. Johnson, T. Lozano-Pérez & M. T. Mason, eds, 'Robot Motion: planning and control', MIT Press, pp. 287–304.

Williamson, M. (1996), Postural Primitives: Interactive Behavior for a Humanoid Robot Arm, *in* P. Maes, M. Matarić, J.-A. Meyer, J. Pollack & S. Wilson, eds, 'Fourth International Conference on Simulation of Adaptive Behavior', The MIT Press, Cape Cod, MA, pp. 124–131.