

Behavior-Based Methods for Modeling and Structuring Control of Social Robots

Dylan A. Shell and Maja J Matarić

September 19, 2004

1 Introduction

People and robots are embodied within and act on the physical world. This chapter discusses an action-centered methodology for designing and understanding control, perception, representation, adaptation and learning in physical robots, inspired by evidence from social biological systems. A working robotic implementation based on a biologically plausible model provides strong support for that model. Primarily it provides a demonstration of classes of behavior for which the underlying theory is sufficient. In addition to ensuring that the underlying theory has been specified with algorithmic rigor, it also assures that the model is effective even in the presence of noise and various other effects of a dynamic environment. Shortcomings of the implementation may highlight concrete issues on which a later and further refined theory may focus, thereby playing an important role in the hypothesize-test-rehypothesize cycle. Robotic systems are one of the few ways to provide a complete end-to-end validation of social theories that deal with self-referential notions and require validation.

A central postulate of the action-centered methodology is that intelligent behavior in an embodied system is fundamentally structured by the actions the system is capable of carrying out. In societal systems the individuals' social behavior, including communicative and interaction actions, similarly structure the large-scale behavior. This is supported by neuroscience evidence and has key impact the way human activity and robot control are understood and modeled.

The belief that reasoning agents/systems should be built upon action-centered characteristics such as the physical dynamics and task constraints toward effective cognitive capabilities embodies what is known in the AI community as a bottom-up philosophy, and contrasts with other so-called top-down views. Behavior-based controllers are a means of principally encoding bottom-up robot control strategies through the use of underlying building blocks termed behaviors. Primitives are another key modularity construct of the action-centered methodology, and are essentially behaviors intended for articulated body motor control. The bottom-up formulation, and in particular the action-centered framework, unifies a swathe of robotics research, as demonstrated first in examples from humanoid robotics (Section 2), then in more traditional behavior-based robots (Section 3), and finally in robot groups and teams (Section 4). Those three sections are arranged to present the foundational biological and neuroscientific evidence incrementally when most pertinent and suitably illustrated by robotic examples. Humanoid motor control through motion primitive structures is described first, since it most closely resembles the organization of natural behavior. This is followed by more general behavior-based control, which in turn is followed by a description of robot teams employing the behavior-based controllers. The exact relationship between primitives and behaviors is described in Section 3.3 and is related to properties of *sets* of these structures.

This chapter's focus is on social interactions with robots. The term "social" is interpreted in two complementary ways. The first refers to rich interactions between people and robots, and is illustrated with a method whereby human movement is used to enable humanoid robots to act more naturally, and an imitation inspired approach to skill acquisition, that allows teaching rather than programming robots. The second refers to the bottom-up use of artificial societies or communities of robots as models of social behaviors. Together, these components of the chapter are intended as a starting point for those studying models or

implementations of neuro-cognitive theories in robots, and for whom social implications and issues of those theories are of primary interest.

2 Control and Imitation in Humanoid Robots

Humanoid robots are designed to resemble the human form. This morphology is expected to allow such robots to effectively utilize human tools and exploit man-made environments in a manner originally intended for humans, in order to perform a variety of tasks. Furthermore, humanoid robots are expected to interact with humans in a natural manner. These two major goal – that humanoids are intended as generalists, and that they are expected to be socially capable – present major research challenges.

Moving beyond task-specific robots requires complex Artificial Intelligence (AI) questions be addressed, notably those dealing with representational issues and programming/control methodologies (Russell and Norvig 2002). How to go about structuring information and knowledge, and computing on those structures, in a way that is sufficiently open-ended yet time-efficient is a difficult and unsolved problem. The fast (technically real-time) responses required of real robots makes these traditionally difficult problems harder still.

Building sociable robots present host of new issues, which include (and are not limited to) social learning, imitation, gesture and natural languages communication, emotion and interaction recognition (Fong et al. 2003). Perception is notoriously challenging and makes recognition of interaction cues difficult, while limited actuation often results in the robot’s actions appearing overly simplistic. Nevertheless, even current robots that were designed with social competency in mind far surpass the more standard research robots in terms of appeal and ability to engage people.

Such robots, when socially embedded, can help to inform theories about human behavior in two ways. The first is in the design of the robots themselves. Various theories are applied during the design process, and the validity of those theories is considered in terms of the final result. Breazeal (2002) describes the application of ideas from developmental psychology in the design of a social robot. Here the robot implementation provides verification of the psychological theory, and interaction with the robot may provide pertinent insights or suggest further issues for study both in robotics and psychology. The second way of informing theories of human behavior is through the development of a robot as an artifact that people can interact with, in order to study the people themselves. In this case the robot allows experimental conditions to be controlled more precisely than would otherwise be possible. Work by Scassellati (2001) uses robotics to control experimental conditions in the process of diagnosing autistic children, with the potential of providing both more consistent diagnosis and a novel means of socializing the patients. Breazeal (2002) shows specific instances of the robot performing simple tracking of salient features in the environment with its eyes, and appearing to have a sophisticated understanding of the content being communicated by a human subject. This may provide insight into the types of non-verbal cues that account for a significant portion of channel bandwidth in human face-to-face interpersonal communication (Argyle 1988).

The remainder of this section discusses a body of research that uses neuroscience evidence about motor control and movement recognition in order to structure movement for humanoid robots. The work is focused on an action-centered philosophy that brings together movement, activity recognition, and generation within a unified representation. The main component of that representation is the notion of complete generalized movement behaviors, called *primitives*.

Such primitives are the main focus of this section because they lay the foundation for a central idea of this chapter: that a sparse set of exemplar structures can be used to capture the inherent regularity of a larger repertoire of behavior. This sparse set can be learned and can be used to recover that broader repertoire. The approach is useful for reducing complexity in the study of social interactions at a number of different levels, as discussed in later sections.

In this section the focus is on the definition of *motor primitives* and their use for control of humanoid robots at the kinematic level. (Kinematic refers to the relationship of motion between the robot constituent parts, ignoring the effects of masses or forces and torques). The empirical evidence for both primitives and *mirror neurons* is discussed; mirror neurons are another concept from neuroscience that has had a key

impact on action-centered methodologies, especially in the area of *imitation learning*, an important form of social interaction. Section 2.4 elaborates the import of this unified representation through a discussion of the implications of an imitation-based methodology for human-robot interaction. Thus, the primitives-based approach to structuring robot (and human) behavior can be used in the development social robots, and hence studies of human behavior through social robotics.

2.1 From Biological Evidence and Neuroscience Inspirations

Robots can serve as helpful scientific tools, used to validate or reinforce research hypotheses about the behavior of natural systems. Matarić (1998, pg. 5) cites examples of robotic (and faithful simulation) models of a variety of natural behavior ranging from general models of reflexive behavior selection (Connell 1990) and evolutionary approaches for single (Nolfi et al. 1994) and multi-robot (Agah and Bekey 1997) systems. Subjects include crickets (Webb 1994), lobsters (Ayers and Crisman 1992; Grasso et al. 1996), frogs (Arbib 1989), flies (Franceschini 1975; Cliff 1990) and various other insects (Deneubourg et al. 1990; Kube and Zhang 1993; Beckers et al. 1994). Section 4 will deal specifically with the design and analysis of systems like these, which feature many, often cognitively simple, interacting constituents.

In manner very similar to the examples of social robots already described, humanoid robots are embodied models that can provide strong validation and insights for neuroscientists interested in primate behavior. Thus these robots in particular can be seen as scientific tools in order to study human behavior, social or otherwise (Atkeson et al. 2000). This section and the next will deal with single robots designed to be relatively competent in social circumstances. Social constraints are, for the most part, outside the roboticists realm of controllable design parameters; consequently, robots that operate within existing social systems provide insights about interfaces to those systems.

The primitives-based approach holds that the foundation for robot control should be a set of parsimonious and composable primitives, or motor control programs. This scheme for modularizing control is an excellent example of cross-pollination of ideas and approaches between biology and engineering, and more specifically between neuroscience and robotics. Engineers are interested in building robots with better performance, but the complexity of physically embodiment means that natural analogues find their role in implementations based on solely pragmatic goals.

The widely cited biological evidence from Mussa-Ivaldi and Giszter (1992) and Giszter et al. (1993), which provides the foundation for the theory of motor primitives, points to the existence of spinal field motor primitives that encode a complete movement (or behavior). Experiments with spinalized frogs and rats have shown that when a individual field is activated, the limb executes a complete behavior, such as a reaching or wiping. Similar movements are exhibited by numerous individuals from the same species. Thus, these motor primitives are referenced, or indexed, through a spatial organization that is species-universal and species-specific. Mussa-Ivaldi et al. (1994) further showed that simultaneous simulation of multiple fields results in either linear superposition or the domination of one field over the others. Finally, only a small number of such distinct fields has been identified in the frog's spine; indicating that only about a dozen primitives may be sufficient for encoding the frog's entire motor repertoire, through sequencing and superposition of supra-spinal inputs through descending pathways (Bizzi et al. 1991).

Extensive research into limb biomechanics has resulted in an organization of the motor control system that is both modular and elegant. A model in which entire behaviors are used as motor primitives, and composition permits the production of a host of additional outputs, captures the essence of this paradigm. Additional principles derived from the frog and rat data include the following:

1. The set of primitives is parsimonious.
2. The primitives span and partition the range of achievable actions.
3. Each of the primitives consists of exemplar movements. When triggered individually, each produces an archetypal motion.
4. Linear superposition and temporal sequencing are used to produce more complex, composite movements.

While motor primitives have served as the inspiration and foundation for modularization of movement, both in animals and in robots, another movement-related line of evidence, for mirror neurons, has also had a major impact on theories of natural and artificial movement, and in particular imitation. Gallese et al. (1996) describe neurons in area F5 of the macaque monkey that appear to functionally connect vision and motor control system. The neurons are dubbed “mirror neurons” because they fire when the monkey observes a specific movement in another monkey (or human) as well as when it executes the same movement itself. Rizzolatti et al. (1996) describe the similar faculty in the human brain, and Iacoboni et al. (1999) establishes a link with the imitation mechanism. The mirror system thus becomes the foundation of the capability to mimic and imitate, both fundamentally social skills. The notion of *true imitation*, which involves acquiring entirely novel skills by observation, is (currently) recognized only in chimpanzees, dolphins, and humans (Bryne and Russon 1998). Monkeys and other species are believed to only be capable of mimicry, involving the mirror system as the means of mapping the observed to the known, but not the ability to expand the known repertoire through this process. Arbib (2002) argues for the mirror system as a keystone in the evolution of the social mechanisms of communication and language. The mirror system is said to provide a mechanism for understanding the grounding of observed actions in a manner consistent with other neural theories.

The combination of a structured representation grounded in action (as enabled by motor primitives) and the bridging of perception and action (as enabled by the mirror system) allows for a general and powerful action-centered model of representation, control, and learning. Humanoid robots employing such a scheme could recognize, imitate, and learn motor skills through natural interaction with other robots and humans.

2.2 Extracting Natural Structure in Human Motion

The preceding sections have described the original biological inspiration for the notion of primitives, and their role in modularizing robot control and social interaction. Next, considerations regarding the origin or selection of an appropriate set of primitives are addressed. An obvious approach commonly employed by roboticists is to manually design a set of primitives which capture an effective structure and modularity of the system. However, the ability to automatically derive such a set is highly desirable. Fod et al. (2002b) and Jenkins and Matarić (2004) showed how such automated derivation of primitives is possible, in the context of humanoid upper-body control, by using human movement data as input.

Jenkins and Matarić (2004) showed empirically that natural human movement forms a smooth manifold, at least for certain activities, because even a small number of representative sample trajectories can give sufficient support for a range of expressive actions. This is the key requirement that allows for automated derivation. The next step was to develop a methodology that was sufficiently powerful to statistically extract the underlying structure in human movement, without the use of programmer bias or pre-programming.

Fod et al. (2002b) described the application of principle components analysis (PCA) to reduce the dimensionality of human upper-body joint angle data. Segmented data are converted into vectors, and eigenvectors extracted. The resulting “eigenmovements” can be combined as low-level primitives to synthesize general motion. Jenkins and Matarić (2004) developed Spatio-temporal Isomap (ST-Isomap) to address the shortcomings of using PCA, specifically those issues arising in parameterization of motion toward controller synthesis. Consequently, prototypical motions that exploit the structure and redundancy inherent in human movement were successfully automatically derived using this method.

Experiments performed on human subjects also inform the process of determining a set of primitives. Matarić (2002a, pg. 7) describes the importance of an appropriate choice of parameterization frame for primitives for a given system and discusses the use of extrinsic end-effector coordinates in the control of humanoid robots. The selection of reference frame depends on the requirements of the task and activity being performed. Matarić and Pomplun (1998) describes human psychophysical data indicating that people pay attention to and track (as indicated by the fixations of their eyes) the end-effector of a limb being observed (such as a finger of a hand, or the tip of a pointer).

2.3 Robot Control and Perception

Although control of manipulators (i.e., arms and arm-like effectors) is a well-studied problem in robotics, the industrial automation roots of traditional approaches render them ill-equipped for the challenges facing robots working in human, social environments. The large number of degrees of freedom (DOF) in any realistic humanoid makes for a complex control problem. The human arm has seven DOF, the hand 23, and the spine an order of magnitude more. Reasonable safety measures dictate that robots should be submissive to external impulses, which fundamentally alters the acceptable engineering approaches that can be employed. Furthermore, issues of usability advocate that effective humanoid movement should also be human-like, a quality that is difficult to formalize and codify.

A number of distinct humanoid motor control methods exist. Alternatives to primitive-based control involve run-time trajectory planning, requiring search or optimization algorithms that are currently too computationally complex for real-time on-line computation required of a physical robot. The primitives-based approach has emerged in part as a response to this need for time-efficient movement computation and in part in response to the principles from biology and neuroscience. It attempts to capture the crucial properties of the functionality of biological systems, rather than their exact mechanisms. The primitives-based model, which is consistent with behavior-based control, postulates that the motor system is structured as a collection of primitives, or motor programs, that encode complete stereotypical movements (Mataric 2002a). These can be executed sequentially or concurrently (through parameterized superposition), to create a large movement repertoire. This general organization of the motor systems is then used by the mirror system to instantiate direct sensory-motor mappings into executable motor programs.

An early humanoid torso using the primitives-based philosophy is described by Williamson (1996) and Marjanović et al. (1996). Four static postures were used as the foundational set of primitives to generate motor controls for a six-DOF robot arm. Schaal and Sternad (1998) used nonlinear dynamics to create attractors that define types of primitives. This approach has more recently been demonstrated in complex simulated humanoids (with between 20 and 132 DOF with dynamics (Mataric et al. 1998, 1999; Drumwright and Mataric 2003), as well as articulated a Sony Aibo dog. Additionally, automatically derived primitives are being deployed on platforms like Vanderbilt University's ISAC humanoid (Kawamura et al. 2004).

Consistent with the action-centered philosophy, movement primitives have also been proposed as effective mechanisms for movement prediction (Mataric 2002a). Jenkins and Mataric (2004) describes the use of future state prediction with primitives encoded as dynamical systems. Drumwright and Mataric (2003) used primitives encoded as a set of parameterized joint space exemplars to enable highly accurate humanoid upper-body behavior recognition and classification using a probabilistic estimation framework.

2.4 Skill Acquisition for Social Robots

As noted above, the mirror neuron system provides another powerful line of biological evidence arguing for action-centered perception and imitation. It is postulated to have played a foundational role in the evolution of social interaction and even language (Arbib 2002). Consequently, endowing robots with similar capabilities effectively sets the stage for social robotics.

Mataric (2002a) describes two additional lines of evidence, also from neuroscience, that enable imitation: *selective attention* and *classification-based learning*. The latter is a machine learning term for an algorithm with the ability to map various observed instances into one of a finite set of classes, frequently improving the mapping through experience. Together, these components permit a model of imitation that captures movement, perception, and learning. Learning in this sense refers to the expansion of the skill vocabulary, through the attainment of new behaviors, not only the improvement of fixed abilities.

If true imitation is attainable, it presents far more than just a mechanism for natural programming of robots. Rather, it presents an open-ended, generative means of activity representation and learning, crucial for social interaction. Imitation integrates cognitive systems from the lowest levels of perception and motor control to the highest levels of cognition. The motivation for studying imitation thus comes from the challenge of understanding this complex and powerful natural phenomenon, and from its potential in enabling social robotics (Mataric 2002a). Much can be done with simpler (non-humanoid) robots in terms

of engaging humans and pursuing nontrivial interactions. The section that follows takes up these sorts of issues while focusing on problems with higher cognitive requirements than the motor control examined here.

3 General Behavior-Based Control

Questions regarding representation and behavioral organization are of pivotal concern in AI and robotics. Control architectures provide a means of principally constraining the space of possible solutions, often focusing on particular representational or planning methodologies, in order to render practical problems achievable. A variety of architectures with different underlying principles have been proposed and demonstrated for robot control.

This section discusses the behavior-based methodology (Mataric 1997a; Arkin 1998) and its connections with the primitive-based philosophy described above. Behavior-based control modularizes complex control problems into a set of concurrently executing modules, called *behaviors*, each processing input and producing commands for the robot's effectors and/or for other behaviors in the system. The dynamics of interaction among the behaviors and the physical world result in the robot's aggregate performance. The behavior-based approach favors a parallel, decentralized approach to control, while still allowing for substantial freedom of interpretation (Mataric 1998).

The behavior-based and movement primitives philosophies stem from the same biological evidence. Early behavior-based work, specifically motor schemas (Arkin 1989), was based on the same neuroscience evidence found in frogs (Bizzi et al. 1991) that has guided the work in movement primitives. Subsequently, the behavior-based methodology has further generalized and adapted the conceptual organization toward a variety of control domains, and away from direct motor control.

The behavior-based methodology has been widely misunderstood, largely due to its lingering confusion with reactive control methods. Criticism frequently focuses on the role of representation, and specifically the capabilities commonly typified by high-level symbolic reasoning. Arguments typically stem from the fact that a variety of early behavior-based work was minimalist in nature, typically reactive and thus incapable of lasting representation and learning. Additionally, the behavior-based approaches foundations lie in observations of biological motor control, a rudimentary low-level mechanism requiring seemingly minimal mental competency. In the last decade, clear distinctions have been drawn between simple reactive and more complex behavior-based systems (Mataric 2002b), shown to be as expressive as planner-based methods.

This section presents a summary of work that has extended previously accepted limits on the behavior-based paradigm, particularly in terms of representational capabilities. Methods used to understand behavioral composition and coordination of multiple behaviors and arbitration mechanisms are also described. The section also illustrates how the behavior structure provides an effective substrate for higher-level capabilities such as path planning for navigation, and learning behavior coordination.

3.1 Behavioral Structure and Artificial Intelligence

Historically, the AI community had worked on disembodied agents, with the robots being unconventional exceptions. Unfortunately, the assumption that subsystems could be "ported" to robots when they were ready proved to be unrealistic. The challenges faced by an embodied agent, including uncertainty in perception and action and a dynamic and unpredictable world, were the very same challenges that had been abstracted away, and thus remained unaddressed. Deliberation alone was not a mechanism that would enable a robot to deal with the contingencies of the real world; most effective criticisms came from practitioners who experimented directly with physical robots (Brooks 1991).

Brooks' widely cited paper (Brooks 1986) describes the difference between traditional control architectures – consisting of functional modules performing perception, modeling, planning, and motor control – and a new decomposition into task-achieving modules layered to produce increasing competence and complexity. Crucially, the task-achieving modules connect perception to action, guaranteeing a reactive, timely response to the physical world. Brooks (1991, pg. 3) further outlines a justification for focusing on "being and reacting" from an evolutionary timescale based argument. The proposed "Subsumption Architecture"

was a means of structuring a reactive system and is the forerunner of contemporary behavior-based robotics, which has evolved since. Work by Arkin (1989) constrained behaviors to perform in a manner much closer to the biologically-inspired vector-fields described in Section 2.1. The notion of behavior has been significantly broadened. Fundamentally, constraints on the behaviors reduce their expressiveness in favor of special-purpose efficacy. The behavior-based methodology attempts to conserve *organizational* principles from biology and neuroscience at an abstract “informational” level, so that the constituent behaviors are minimally constrained, for maximal system flexibility. This flexibility, illustrated by a wide variety of implemented behavior-based systems, has unfortunately also fostered an ongoing confusion about the nature and limitations of the methodology.

Current behavior-based controllers are still characterized by a bottom-up construction, with each module corresponding to an observable pattern of interaction between the robot and its environment. The modules, typically called behaviors, operate concurrently and at similar time-scales, and interact with one another. This forms the substrate for embedding representation into behavior-based systems (Mataric 1992); the representation is inherently distributed. Behavior-based systems are best suited for environments with significant dynamic changes, where fast response and adaptivity are necessary, but the ability to do some looking ahead and avoid past mistakes is also useful (Mataric 2002b).

System decomposition by activity (Brooks 1991) ensures a concrete connection between perception and action, a principle already described above, in the context of primitives. Behaviors represent “activities” not because of the semantics of their labels, as in classical AI, but because they capture inherent regularity of the interaction dynamics of the robot and the world. Behaviors are thus encapsulations of dynamics, and are made general through parametrization.

The repercussions of the new behavioral organization continue to have an impact on robotics and AI. What had essentially been suggested was that the direction that had been taken since the field’s inception was incorrect and founded on the idea that would not carry over to physical robots. In many cases the traditional behavioral organization was assumed at the time to be the *only* way to structure systems. Many had assumed that scaling issues (e.g., to large or continuous state-spaces), issues of partial-observability, non-stationarity, and uncertainty, could (and further should) be addressed from within their traditional representation.

3.2 Representational Issues

The pioneering role of the Subsumption Architecture, the title of Brooks (1991), and an unfortunate lexical collision with “behaviorism” in psychology have all resulted in the broadly accepted misconception that behavior-based systems do not permit representation. As early as Mataric (1992), representation was introduced into behavior-based systems, in that case in the context of topological spatial mapping and path-finding.

Mataric (1992, 1997a) describe the work with *Toto*, a mobile robot that was first to use dynamic behaviors, created and activated whenever needed to represent landmarks in the environment. Planning, previously absent in behavior-based systems, is achieved in *Toto* through spreading activation within the network of map behaviors. Mataric (1992) describes high-level competencies for landmark detection using unique time-extended sensory signatures. The landmark behaviors are used to fill behavioral-slots, resulting in a graph of active map locations. This coupling of action, perception, and representation is similar to the mirror-neuron and motor primitive model already described, but *Toto*’s most significant behaviors are several cognitive degrees of separation from basic motor control. Mataric (1992) mentions that this mechanism falls under the broad umbrella of ideas termed *cognitive maps*, and that it is representative of a particular interpretation of the organization and function of the rat hippocampus. The faithfulness of this type representation to the actions that the robot could perform and the constraints and dynamics that structure its action space are summed up in the maxim: “behavior-based systems think the way they act.” (Mataric 2002b).

Decety (1996) and Jeannerod and Decety (1995) provide evidence indicating that biological systems may operate in the same manner, showing that both imagined and executed movements share the same neural substrate. When simply imagining or visualizing a movement, subjects’ motor pathways exhibited activation similar to that which occurs during actually performing the movement. This evidence points to principles

employed by behavior-based architectures from a organization level, as well as to their embodied approach to representation.

The traditional view holds that higher-level cognitive capabilities are best modeled using symbolically. As an alternative, Nicolescu and Mataric (2002) describe a hierarchical behavior-based architecture that enables behaviors to represent more abstract concepts. The inclusion of both external and sequential preconditions allows their *abstract behaviors*, to cope with temporal sequencing while maintaining the conventional concurrent execution. Their *network abstract behavior* hides the details of an entire network of behaviors, and presents an external interface as if it were a single behavior; recursive application enables general hierarchical representations. The work thus allows for representing temporal sequences and hierarchical structures, without using plan operators or symbolic mechanisms. As a result, there is no need to produce a “middle layer” to bridge the difference-in-kind between reactive and symbolic layers in hybrid systems, the common alternative to behavior-based systems. Nicolescu (2002, pg. 68) describes expressive power of the framework in terms of a particular human-robot interaction task.

Here, and in the Toto work, and generally in behavior-based systems, representations are stored in a distributed fashion, so as to best match the underlying modularity that produces the robot’s observable behavior. If a robot needs to make high-level decisions (e.g., plan ahead), it does so in a network of communicating behaviors, rather than a single components (e.g., a centralized planner). Thus, the very same characteristics that structure the robot’s action space also impact the way the robot represents and reasons about the world, as in biological evidence indicated above.

3.3 Behavior Composition

The previous sections has, for the most part, discussed only single behaviors; important questions arise when collections of behaviors are considered. Typically, behaviors are hand-designed to perform a particular activity, attain a goal, or maintain some state. It is, of course, impossible to define an “optimal” behavior set. Nevertheless, practical experience has established a number of consistent guiding principles.

Mataric (1995) describes *basis behaviors* as a useful tool for structuring and thus simplifying behavior synthesis, i.e., design. Basis behaviors are a set of behaviors that are *necessary* in the sense that each either achieves, or helps achieve, a relevant goal that cannot be achieved by other members of the set. Furthermore, a basis behavior set is *sufficient* for achieving the goals of the system/robot. Other desirable properties include simplicity, stability, robustness, and scalability.

Section 2 noted that primitives are to be 1) parsimonious and 2) required for the generation of appropriate behavior. These two properties distinguish basis behaviors from a larger space of feasible behaviors. In a sense, basis behaviors are homomorphic with primitives. The word “basis” was selected to be indicative of the similar notion within a linear algebra¹. The property of parsimony (or necessity) is analogous to the idea of linear independence; the property of sufficiency is similar to the linear algebraic concept of a span.

An additional third condition, namely that behaviors each individually generate prototypical interactions, suggests that not all sets of basis behaviors will necessarily be equal. Thus, a set of *natural basis behaviors* can be defined such that, continuing the algebraic metaphor, the additional constraint of most natural individual interactions is respected. This third provision attempts to capture the fact that it may be beneficial (certainly in a biological system) to have frequently performed actions become part of the basis set. Under realistic conditions, a particular “projection” may be better than another; this idea of a natural basis undertakes to express that notion. Motion primitives described above, and particularly those that are the desired output of data-driven derivation methods, are thus simply described as a natural basis set intended for motor control of humanoid robots and other articulated manipulators.

Another recognized organizational principle is *orthogonality*. Two behaviors are orthogonal if they do not interfere with one another, each inducing no side-effects in the other (see further discussion on page 12). Toto’s obstacle-avoiding navigation system is an example; its four low-level, reactive navigation behaviors were triggered by mutually exclusive sensory conditions (Mataric 1992). Another method is to have different

¹In various other references, including Mataric (1992, 1995, 1997b), the words “basic” and “basis” have been used interchangeably. Afterthought suggests that the latter is in fact preferable.

behaviors control separate degrees of freedom (DOF) of the system. This form of factorization is only feasible when the robot’s dynamics do not inhibit the separability. An extreme example of highly coupled DOF are helicopters; Saripalli et al. (2002) have demonstrated that behavior-based control is still feasible for particular behavioral regimes.

In addition to the choice of behaviors themselves, an important issue in behavior-based systems is the action selection mechanism. This is a well studied problem that deals with coordinating individual behaviors in a manner that produces the most coherent global results in spite of possible disparities or conflicts between behavioral units. Two frequently employed methods are *prioritization* and *fusion*. Prioritization, wherein one behavioral unit is given the ability to overrule the output of others, is particularly useful for modeling those situations in which safety or survival dictates fast actions. Fusion involves the combination of output from various behaviors so that effects of each are incorporated, often using a weighted average, or related ideas. Fusion is useful for performing actions that make progress toward the achievement of multiple concurrent goals. In spite of the simplicity of the prioritization and fusion operations, they are the often employed. Many more sophisticated alternatives exist; Pirjanian (1999) provides an extensive review of significant work addressing the challenges of behavior selection.

3.4 Adaptation and Learning

Adaptation and learning are some of the most fundamental yet difficult properties for intelligent systems. Robotic domains present serious challenges for learning: the physical world involves very large, if not infinite, state-spaces, the robot’s world is partially observable, the robot’s actions are nondeterministic, and feedback from the environment (including people and other robots) may be grossly delayed or entirely unavailable.

A reinforcement learning robot is given an external reward signal (in robotics typically from a benevolent system designer) from which it must separate good actions from bad ones. The correlations between rewards and actions can be nontrivial because of possible delays between action and reward, because of environmental dynamics that constrain or aid a robot’s actions and alter the reward critically, and because of various further complications in social situations where the actions of multiple individuals may be responsible for the a single global reward, or factor into each robots’ own reward.

Matarić (1997a) describes a reformulation of the reinforcement learning methodology, using basis behaviors as the atomic representation. The use of goal achieving behaviors that hide low-level control details allows the full state space to be replaced with a smaller set of behavior conditions. Since behaviors, by their definition, capture the underlying task dynamics and its intrinsic structure, they capture only the details necessary for the learning system. Conditions are much fewer than states, so their use diminishes the robot’s learning space and speeds up any reinforcement learning algorithm.

The non-stationary character of real environments, coupled with uncertainty in sensing and action, require additional mechanisms in order to make learning feasible within a suitable time span. Mahadevan and Connell (1992) demonstrated the first major success in applying reinforcement learning to a mobile robot, which learned to push a box. The problem was made tractable through the use of a behavior-based structure, in which the robot learned not in the prohibitive global state space, but within the context of each behavior. The behavior structure thus provides a means of accelerating learning. Matarić (1997a) used a related idea; *shaped reinforcement*, which provides informational cues for the robot during the execution of a given behavior, and upon its termination, is another means of taking advantage of the modularity afforded through the behavior structure.

Beyond learning to coordinate behaviors more appropriately, another focus has been on learning behaviors themselves, in particular from demonstration, so through some form of imitation. The idea of “putting through” exploits the robot’s embodiment; new skills are acquired through the robot’s own experience, and the demonstrator – typically a person who will play the role of teacher – guides the robot. Nicolescu (2002) describes an example of this approach: the robot is allowed to learn the effect of its actions on the perceived world, while guided by a human coach. To be accurate, the robot learns about an entire task and task representations, not just a single behavior, or even how to coordinate existing behavior.

Behaviors present an effective substrate for robot learning. Different behavioral levels can each be tuned independently, each with their own shaped rewards or some similar mechanism. The learning mechanism

can operate at either a meta-level, above an operational system, but responsible for adjusting various parameters. Alternatively it may effect the very heart of the behaviors themselves, and the process of behavior construction. Mataric (2001) provides an overview of work in behavior-based robot learning.

Within the behavior-based paradigm, task dynamics effect the behavior choice. The structure which defines the space of goal-driven actions biases the means for perception, the form of representation, and what is subsequently later by the system. Thus, the way behavior-based robots act not only effects the way they think, but also the way they learn.

4 Collective Behavior-Based Robotics

Practitioners who build behavior-based systems often do so in order to demonstrate some behavioral theory, and may occasionally focus on biologically plausible techniques. This bottom-up methodology relies on the ability to generate selected behavior; synthetically producing appropriate behavior demonstrates that a model includes all essential details. For example, the above described Toto experiments in Mataric (1992) demonstrated that a particular high-level model of the rat hippocampus was indeed sufficient to achieve certain spatial awareness.

The same focus on generating behavior has been relatively successful with groups of robots, where the focus is on demonstrating collective phenomena. The research typically concentrates on employing minimalist robots in order to demonstrate particular social phenomena. This area is broadly known as *swarm robotics*, due to the focus on synergistic effects from simple constituents. An elegant example is presented by Beckers et al. (1994), effectively demonstrating how physical dynamics of simple non-communicating robots enable repeatable and robust clustering global behavior in a manner even simpler than the uncomplicated mechanism postulated by Deneubourg et al. (1990).

Coordinating a set of robots is a demanding problem, and adding robots to a system makes the problem more complex. A multi-robot system, like a humanoid system, consist of many degrees of freedom, but, unlike a humanoid, is loosely coupled without the same physical constraints. Individual components in a multi-robot system may (and usually do) only have local perception, communication, and computation capabilities. Such a system may also deal with possibly unpredictable (Darley 1994) emergent collective dynamics. Multi-robot coordination involves issues of distributed task allocation, communication, and action selection, all under uncertainty. The same problems, but without physical uncertainty properties, are also studied by the multi-agent research community, resulting in some shared and some quite divergent insights (Gerkey and Mataric 2004).

This section discusses the application of basis behaviors to multi-robot coordination analysis and synthesis. Composition of basis behaviors enables a small set of local rules, known to produce predictable spatial patterns, to be used in the generation of a repertoire of global behaviors. The classic work that demonstrated the applicability of basis behaviors to groups is described, which demonstrates the production of social skills from other far simpler behaviors. Arguments for basis behaviors as an appropriate level of description for learning can be extended to learning in groups; validating experiments are also summarized below.

4.1 Behavior Composition

Section 3.3 described basis behaviors as those which are both sufficient and necessary for achieving some system task or goal. The notion has been demonstrated to be particularly useful at the collective level; it was originally implemented in that context.

Mataric (1995) describes work done with the *Nerd Herd*, a homogeneous troupe of 20 autonomous mobile robots with very limited sensing and computational abilities. The robots were equipped with a simplistic gripper used for holding and releasing metal pucks. Each robot was programmed with an identical small set of behaviors: **following** (following a single robot in front), **wandering** (moving about without collisions and without a particular goal), **homing** (moving toward a particular Cartesian goal location), **aggregation** (moving toward the nearest neighbors), and **dispersion** (moving away from nearest neighbors).

Although the behaviors consisted of reactive rules based on local sensory information, their labels prescribe actions based on only locally sensed information, they are named for larger scale group phenomena they collectively produce. Thus, when together, individuals executing these local behaviors produced global patterns, such as aggregation into a single group, spreading out over an area, herding, surrounding, docking, flocking, and foraging (Mataric 1995).

For **flocking**, the vectorial components of four behaviors (**following** being the exception) are simple weighted sums, with empirically obtained values for the weights. In the case of **surrounding**, sums of the outputs from **aggregation** and **following** are sufficient. The case of **herding** is interesting because it involves **flocking** and **surrounding**, both of which are compound behaviors.

To produce **foraging**, temporal sequencing and behaviors triggered by non-mutual sensor states are used. Four basis behaviors (**aggregation** is excluded), triggered by different sensing states (an example would be: switch to **homing** when a puck is held by the gripper), produce collective foraging. Mataric (1995) provides details and a pseudo-algorithmic description for the basis behaviors, as well as **foraging** and **flocking** compounds.

The basis behaviors were designed to conserve energy by minimizing interference between robots. Interference is one of the chief determinants of performance in multi-robot systems and is often the defining impediment, resulting reduction in performance with the addition of robots past a particular critical number (Arkin et al. 1993; Balch and Arkin 1994). Goldberg (2001) describe a methodology for evaluation and minimization of interference and interference causing circumstances.

Mataric (1997b) describes the problem of learning *social rules* in order to minimize the negative effects of groups and maximize synergy. Two functionally distinct types of competition are identified: the first, goal competition, is the result of individuals who possess conflicting goals and hence must compete; the second, resource competition, is a consequence of multiple embodied robots;

Goal competition arises when dealing with multiple agents, even those that are not situated within a world with dynamics of its own. It is studied generally within the distributed AI and multi-agents community. The chapter by N. Schurr, S. Okamoto, RT. Maheswaran, P. Scerri and M. Tambe in this volume provides one such example. Resource competition is the result of physical coexistence; robots take up space, consume energy, etc. Resource issues tend to become further exacerbated by the addition of robots, and thus they can be the factors which limit the scalability of a particular system.

In the foraging studies the space around the “Home” area (to which foraged pucks must be returned) rapidly becomes congested. Goldberg (2001) studied methods for reducing this resource contention, including spatial (territorial) and temporal (time-sharing) division. He also developed a statistical mechanism that allows robots to model interaction dynamics online (i.e., through experiences within an experiment) so that behavioral decisions could explicitly consider (and estimate) interactions. This is an attractive way to deal with interaction dynamics, because the robot can sense those effects (emergent or otherwise) that are most important and reason about or act on them. Section 4.2 discusses more of the specifics for behavior-based learning in multi-agent systems.

Mataric (1995) enumerates a number of examples of biological groups that exhibit those five basis behaviors, showing that they are plausible activities for individuals within a group to perform. An interesting unanswered question is if some other less obvious basis set exists that may produce the same collective behaviors. Of course, foraging and flocking are also social activities that occur in nature and clearly involve more competence than merely the outlined basis behaviors. Certainly the sequencing of distinct activities, as in **foraging**, introduces a degree of complexity. In the case of **flocking**, intuitively it seems acceptable behavior is far more constrained than any of the four constituent basis behaviors. For example, the **dispersion** and **aggregation** competencies provide lower and upper bounds on robot-to-robot distances.

This demonstrates a central behavior-based principle first described in Brooks (1986): the addition of competencies frequently *decrease* size of the feasible set of actions because behaviors provide an increasing number of constraints. A “blank-slate” robot has no limitations on which actions it can perform in a given situation. A robot that has an **obstacle-avoid** behavior has the constraint that it may not perform a forward action while in front of a wall. Thus, adding layers, and hence increasing the cognitive capabilities of the robot, refines the suitable actions for particular situations.

Another interesting example of collective behavior modeled using simple rules and through computational methods is the work by Helbing et al. (2000). In this case the social phenomena of interest is crowd behavior, and the work demonstrates that simple force laws between the modeled people is sufficient to produce a number of naturally observable macroscopic effects, like self-organized lane formation, oscillatory behavior at bottlenecks and herding during panic. Again, interference patterns are a fundamental sign of non-linearities. Shell and Mataric (2004) suggests that robots interacting with large numbers of people, and attempting to direct the ensemble behavior, may result in useful systems in the future.

The two cases of the **foraging** and **flocking** robot behaviors are combined in a manner that thoughtfully employs a high-degree of orthogonality. Basis behaviors are orthogonal in the sense that they have a degree of independence and minimal interference with one another. The term is adopted from the mathematical notion of orthogonal functions, where two functions have an inner product that is zero. In the **flocking** behavior orthogonality is achieved through the distance thresholds, which were selected so as to have least possible overlap. In the **foraging** case, temporal sequencing allows certain concurrent execution, but minimizes the basis behaviors that directly conflict (i.e. maximizing orthogonality). In both cases, orthogonality makes it is easy to construct an intuitive argument for the prescribed collective behavior. While this makes the modeling easier, it is very hard to achieve in general; for compelling examples see the unexpected complications in Holland and Melhuish (1999) versus serendipitous success in Beckers et al. (1994). Small effects may be multiplied through positive feedback or autocatalytic effects toward either fortuitous or undesirable consequences.

Mataric (1995) describes ethnologically inspired foraging with homogeneous robots. Working with physical robots and real hardware means that no two agents are truly identical. Unique sensory and actuator properties and other variability between robots may become multiplied over time. Flaws, such as wheels on one side of a robot having slightly different frictional properties than the others, can result in a general sub-optimal behavior like turning left more often than right, can be averaged out with multiple robots. Systematic flaws are likely to be amplified, because, as in this example, all robots would be similarly biased and may produce globally consistent consequences.

Balch (1998) describes a methodology for quantitatively measuring diversity in teams of robots, so that a single metric for degree of heterogeneity can be applied to behaviorally distinct individuals. When a task demands heterogeneous robots rather than groups of homogeneous ones is not well understood yet. Variability, either intentional as in a team with differently equipped robots, or unintentional through physical inconsistencies, create a demand for robust and adaptive behavior and provides stringent tests for basis sets.

4.2 Learning

Section 3.4 mentioned the basis behavior approach to reinforcement learning, enabling tractable learning for the noisy, nondeterministic, and uncertain environments. Mataric (1997c) proposed this methodology as a means of managing the state-space size in the multi-robot case, as well as the interactions between robots, including interference, that make the environment far more stochastic than facing a typical single robot (Mataric 1997a).

Adaptations of reinforcement learning have been made effective in the single- and multi-robot domains. The use of behaviors as a substrate has in particular been effective in making learning feasible in complex multi-robot domains. The method has been validated within the foraging domain, using the same herd of twenty robots.

Three structural mechanisms were used in the design of the learning system:

Conditions The size of the state-space can be drastically reduced by focusing on only those states that are important for the systems behaviors. Thus, *conditions* are that small subset of the states which are necessary and sufficient for triggering behaviors. The basis-behavior set provides a perspective from which relevant states can be identified and irrelevant one rejected.

Progress Estimators Estimators of behavioral progress are maintained for feasible behaviors allowing useful feedback for the learning mechanism during behavioral execution. For example while moving

toward a region labeled “Home”, a useful progress estimator is the distance still remaining. Estimators are internal to the robots, but are tied to the physical task achievement. Thus, distance already traveled is not sufficient because unforeseen circumstances may invalidate that progress.

Shaped Reinforcement The provision of reinforcement to behaviors rather than the entire robot helps bypass the credit assignment problem (which is the challenge of deciding which actions are responsible for a particular reward). The combination of internal progress estimators (temporal information) and behavioral reinforcement are together called Shaped Reinforcement.

Without the additional information provided by the progress estimators and shaped reinforcement, or the space reduction of conditions, the real-world challenges could not be faced head-on by the learning algorithm. The demonstration that these three additions simplify the reinforcement learning problem enough to make it feasible for multi-robot foraging with physical robots, provides strong support for the behavior-based approach to collective robot control.

Matarić (1997b) reports on the challenges of using these mechanisms to produce altruistic behavior aimed at a collective good out of robots that often act greedily, and perhaps even irrationally. Robots acting within in a social setting have additional sources of information. For example, observation of a peer performing a successful action constitutes a reinforcement signal. Particularly when coupled with internal progress estimators this signal can provide useful feedback. If robots may assume that their peers are running similar learning algorithms, then the observation behavioral changes in other robots also allows inference of previously received signals. This type of vicarious reinforcement is only feasible for many robots – permitting a partially parallel exploration of the state and reward spaces.

A learning mechanism involving high-level social interactions between robots is described in Nicolescu (2002, pg. 96). An implementation of learning via demonstration for skill exchange between robots is demonstrated; one robot teaches another through an imitation-like mechanism. The same machinery used to learn from a human teacher (Nicolescu and Matarić 2001) can be seamlessly employed for robot-to-robot interactions.

This framework has a number of advantages. First, the robots learn the dynamics of the world *in situ*, through their own experiences. Second, the rich interactions make a small number of trials sufficient for learning and generalization. Task learning from even a single example has been demonstrated. Third, the mechanism is natural for human users. Finally, the work is a demonstration of learning and generalization at a cognitive level through a physically grounded experience. Such a methodology can permit two robots, each with different histories, internal representations, and differing morphologies, to impart knowledge on one another. This and related reproduction of this sort of effect within synthetic systems suggest that these tools are becoming suitable for modeling and studying an extensive range of social phenomena, including those involving sophisticated interactions.

4.3 Activity Modeling

In much the same flavor as the data-driven methods described in Section 2.2 for deriving humanoid movement primitives, recent work has aimed at tracking the spatial patterns of people form during natural daily activities. Fod et al. (2002a) describe an algorithm for fusing the data from several laser ranging devices, permitting people’s motions to be tracked while they move. Yan and Matarić (2002) demonstrate the use of proxemics-inspired spatial features for modeling activity tracked from the laser devices. The work also provides an empirical methodology for validating the proxemics theory; the use of a natural dimensionality reduction techniques may produce the same prescriptive rules through more mechanistic means.

Panangadan et al. (2004) describe a statistical approach to estimating when a type of interaction occurs between people, and demonstrates it on tracking data from a table tennis game. Intuitively, the method infers a degree of interaction by estimating the degree of “influence” a person’s actions have on those of others; thus the model does not rely on physical proximity as a sign of interaction and is able to detect interactions even in spatially distal scenarios. Models of people’s natural social interactions are important for robots in human environments. There may be applications in which robots may need to interact with

large groups of people (Shell and Matarić 2004), or where such models can be useful for evacuation and other related uses (Shell and Matarić 2003).

5 Discussion

The common thread connecting the body of work described above is the unifying action-centered bottom-up philosophy and approach taken to address problems of control, perception, and representation in physically embodied agents. The robot’s capabilities, i.e., behaviors, form the foundational structure from which all other capabilities are derived. The large space of feasible actions is factorized into a smaller set of behaviors whose purpose is to span the action space while reflecting its structure. This is the essence of the basis behavior methodology.

Composition, through operations like superposition and sequencing, is used to recover the complete action space. In practice, the basis set and compositional operators may not fully reproduce the space of actions, and hence the approach imposes limitations on the robot. The reduction of dimensionality also involves some loss of information, but, importantly, not a loss of structure.

The basis behavior methodology has been applied to the so-called *inverse problems*. For example, motion primitives can be used for producing joint-angle trajectories for humanoids (Drumwright and Matarić 2004), thereby sidestepping the traditional inverse kinematics calculation by trading off computation for lookup. Other inverse problems, like achieving a desired swarm behavior through local rules, have also been studied in this context (Kube and Zhang 1993). In general inverse problems are often under-constrained, making the basis behavior methodology an elegant solution, but only after a required amount of designer creativity.

Alternative composable mechanisms, including vector fields and motor programs, have been proposed for use in robotics (Arkin 1989). The notion of basis behaviors, in contrast, is more general, and thus applicable at various levels of abstraction. The basis behavior approach has been demonstrated in a range of circumstances from low-level motor control to social interactions. The dynamics and task-level constraints, which behaviors are designed to exploit, are present in and responsible for structuring many levels of behavior. This reoccurring structure makes the behavior-based approach particularly suitable for analysis of a range of complex systems, where choosing a correct level of system description traditionally is a key difficulty.

The applicability of these same ideas is suitably exemplified by comparing the imitation inspired mechanisms in work by Jenkins and Matarić (2004) and Nicolescu (2002). The former considers “skill learning”, the derivation of motor primitives for control that can later be used for performing more complex behaviors. The latter considers “task learning,”, focusing on higher-level behavior notions and generation of abstract behaviors. Both use a basis behavior set, but focus on behaviors which fall in contrasting parts of the cognitive spectrum. Jenkins et al. (2004) addresses the distinctions between these levels of learning.

Neither the vast majority of behaviors nor systems described herein are linear in nature. The nomenclature adopted throughout this paper, however, is linear algebraic in origin. The described properties, such as spanning, irreducibility, and independence, are not formally proved in the discussed work, but the algebraic analogy may prove useful. Similarly, in the generation of movements for humanoids, a trajectory formed through composition of two natural exemplars need not look natural at all, in theory. In practice, however, for many operational regimes it does. Nonlinearities have been observed and are in some cases the origin of the complex behavior that make these systems interesting. One example mentioned above is the failure of PCA to produce a natural set of basis behaviors from human motion data, while a non-linear method (ST-Isomap) proved highly effective (Jenkins and Matarić 2004).

6 Summary

This chapter has described a united action-centric methodology for generating a wide range of robot behavior. The produced behavioral complexity ranges from natural motor control for humanoids to effective collective behavior for robot teams. The reviewed basis behavior methodology advocates a sparse set of composable behaviors whose achieved goals are not reducible to the other behaviors in the basis set. Thus, for a given goal,

the cardinality of the basis set is bounded below through the requirement of sufficiency, and bounded above by necessity. The methodology is inspired by a spectrum of biological evidence ranging from neuroscience to cognitive science to ethology. The most compelling support stems from the neural foundation for motor control; evolutionary and sociological justifications were also discussed. The basis behavior approach is an outgrowth of the behavior-based methodology, which is rooted in the tight perception-action loops justified by biological evidence.

Central to the discussion throughout has been the applicability of the basis behavior methodology for designing “social” robots. These have been demonstrated through two different, but complementary, approaches. The first is through the synthesis and study of complete artificial micro-societies of robots that work together, producing global effects through local operations. In this case, basis behaviors have been an illuminating mechanism sufficient for linking the local action to their global results. The second approach considered robot interaction with humans; imitation was presented as a principled and expressive mechanism for human-robot interaction and learning. For humanoid robots this can take the form of learning “natural” skills. The effectiveness of learning from demonstration in mobile robots, through rich task descriptions gleaned from situated interactions, was also described.

The work described herein demonstrates the belief that the concept of action-centric behaviors or primitives, stable prototypical interactions that exploit the dynamics of a given system to achieve or maintain goals, can be generalized through the levels of adaptive control, from low-level motor actions to social interactions, thereby providing a methodology for design and analysis of complex behavior.

Acknowledgments

The work reported here was supported in part by DARPA MARS DABT63-99-1-0015, TASK F30602-00-2-0573 and MARS2020 5-39509-A grants, and via the NASA subcontract Grant NAG9-1444, and in part by the ONR MURI N00014-01-1-0890 and DURIP N00014-00-1-0638 grants, and in part by the NSF CAREER Grant to M. Matarić No. 9896322.

References

- Arvin Agah and George A. Bekey. Phylogenetic and ontogenetic learning in a colony of interacting robots. *Autonomous Robots*, 4:85–100, 1997.
- Michael Arbib. Visuomotor coordination: Neural models and perceptual robotics. In *Visuomotor Coordination: Amphibians, Comparisons, Models and Robots*, pages 121–171. Plenum Press, 1989.
- Michael Arbib. The Mirror System, Imitation, and the Evolution of Language. In Kerstin Dautenhahn and Chrystopher L. Nehaniv, editors, *Imitation in Animals and Artifacts*, pages 229–280. MIT Press, Cambridge, MA., USA, 2002.
- Michael Argyle. *Bodily Communication*. Methuen & Co., London, UK, second edition, 1988.
- Ronald C. Arkin. Motor Schema-Based Mobile Robot Navigation. *The International Journal of Robotics Research*, 8(4):92–112, 1989.
- Ronald C. Arkin. *Behavior-Based Robotics*. MIT Press, Cambridge, MA., USA, 1998.
- Ronald C. Arkin, Tucker Balch, and Elizabeth Nitz. Communication of Behavioral State in Multi-agent Retrieval Tasks. In *Proceedings IEEE/RSJ International Conference on Robotics and Automation (ICRA’93)*, volume 3, pages 588–594, Atlanta, GA., USA, May 1993.
- Christopher G. Atkeson, Joshua G. Hale, Frank Pollick, Marcia Riley, Shinya Kotosaka, Stefan Schaal, Tomohiro Shibata, Gaurav Tevatia, Ales Ude, Sethu Vijayakumar, and Mitsuo Kawato. Using Humanoid Robots to Study Human Behavior. *IEEE Intelligent Systems*, 15(4):46–56, 2000.

- Joseph Ayers and Jill Crisman. The Lobster as a Model for an Omnidirectional Robotic Ambulation Control Architecture. In Randall D. Beer, Roy E. Ritzmann, and Thomas McKenna, editors, *Proceedings, Biological Neural Networks in Invertebrate Neuroethology and Robots*, pages 287–316. Academic Press, New York, NY., USA, 1992.
- Tucker Balch. *Behavioral Diversity in Learning Robot Teams*. PhD thesis, Georgia Institute of Technology, College of Computing, December 1998.
- Tucker Balch and Ronald C. Arkin. Communication in Reactive Multiagent Robotic Systems. *Autonomous Robots*, 1(1):27–52, 1994.
- Ralph Beckers, Owen E. Holland, and Jean-Louis Deneubourg. From local actions to global tasks: Stigmergy and collective robotics. In *Artificial Life IV. Proc. Fourth International Workshop on the Synthesis and Simulation of Living Systems*, pages 181–189, Cambridge, MA., USA, July 1994.
- Emilio Bizzi, Ferdinando A. Mussa-Ivaldi, and Simon F. Giszter. Computations Underlying the Execution of Movement: A Biological Perspective. *Science*, 253(5017):287–291, July 1991.
- Cynthia L. Breazeal. *Designing Sociable Robots*. MIT Press, Cambridge, MA., USA, 2002.
- Rodney A. Brooks. A Robust Layered Control System for a Mobile Robot. *IEEE Journal of Robotics and Automation*, RA-2(1):14–23, April 1986.
- Rodney A. Brooks. Intelligence without Representation. *Artificial Intelligence*, 47(1-3):139–159, 1991.
- Richard W. Bryne and Anne E. Russon. Learning by Imitation: A Hierarchical Approach. *Journal of Behavioral and Brain Sciences*, 21(5):667–721, October 1998.
- Dave Cliff. The computational hoverfly: a study in computational neuroethology. In Jean-Arcady Meyer and Stewart W. Wilson, editors, *Proceedings, From Animals to Animats, First International Conference on Simulation of Adaptive Behavior (SAB-90)*, pages 87–96. MIT Press, September 1990.
- Jonathan H. Connell. *Minimalist Mobile Robotics: A Colony Architecture for an Artificial Creature*. Academic Press, Boston, MA., USA, 1990.
- Vince Darley. Emergent Phenomena and Complexity. In Rodney A. Brooks and Pattie Maes, editors, *ALife IV*. MIT Press, July 1994.
- Jean Decety. Do imagined and executed actions share the same neural substrate? *Cognitive Brain Research*, 3:87–93, 1996.
- Jean-Louis Deneubourg, Simon Goss, Nigel R. Franks, Ana B. Sendova-Franks, Claire Detrain, and Laetitia Chrétien. The Dynamics of Collective Sorting: Robot-Like Ants and Ant-Like Robots. In Jean-Arcady Meyer and Stewart W. Wilson, editors, *Proceedings, From Animals to Animats, First International Conference on Simulation of Adaptive Behavior (SAB-90)*, pages 356–363. MIT Press, September 1990.
- Evan Drumwright and Maja J. Matarić. Generating and Recognizing Free-Space Movements in Humanoid Robots. In *Proceedings IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'03)*, pages 1672–1678, Las Vegas, NV., USA, October 2003.
- Evan Drumwright and Maja J. Matarić. Exemplar-Based Primitives for Humanoid Movement Classification and Control. In *Proceedings IEEE/RSJ International Conference on Robotics and Automation (ICRA '04)*, pages 140–125, New Orleans, LA., USA, April/May 2004.
- Ajo Fod, Andrew Howard, and Maja J. Matarić. A laser-based people tracker. In *Proceedings IEEE/RSJ International Conference on Robotics and Automation (ICRA '02)*, pages 3024–3029, Washington DC, USA, May 2002a.
- Ajo Fod, Maja J. Matarić, and Odest C. Jenkins. Automated Derivation of Primitives for Movement Classification. *Autonomous Robots*, 12(1):39–54, January 2002b.
- Terrence Fong, Illah Nourbakhsh, and Kerstin Dutenhahn. A survey of socially interactive robots. *Robotics and Autonomous Systems*, 42:143–166, 2003.

- Nicolas Franceschini. Sampling of the visual environment by the compound eye of the fly. In Allan Snyder and Randolph Menzel, editors, *Photoreceptor Optics*, pages 98–125. Springer, Berlin, 1975.
- Vittorio Gallese, Luciano Fadiga, Leonardo Fogassi, and Giacomo Rizzolatti. Action recognition in the premotor cortex. *Brain*, 119:593–609, 1996.
- Brian P. Gerkey and Maja J. Matarić. Are (explicit) multi-robot coordination and multi-agent coordination really so different? In *Proceedings of the AAAI Spring Symposium on Bridging the Multi-Agent and Multi-Robotic Research Gap*, pages 1–3, Palo Alto, CA., USA, March 2004.
- Simon F. Giszter, Ferdinando A. Mussa-Ivaldi, and Emilio Bizzi. Convergent force fields organized into the frog’s spinal cord. *Journal of Neuroscience*, 13(2):467–491, 1993.
- Dani Goldberg. *Evaluating the Dynamics of Agent-Environment Interaction*. PhD thesis, University of Southern California, Computer Science Department, 2001.
- Frank Grasso, Thomas Consi, David Moutain, and Jelle Atema. Locating odor sources in turbulence with a lobster-inspired robot. In Pattie Maes, Maja J. Matarić, Jean-Arcady Meyer, Jordan Pollack, and Stewart W. Wilson, editors, *Proceedings, From Animals to Animats 4, Fourth International Conference on Simulation of Adaptive Behavior (SAB-96)*, pages 101–112. MIT Press, 1996.
- Dirk Helbing, Illés Farkas, and Tamás Vicsek. Simulating dynamical features of escape panic. *Nature*, 407:487–490, September 2000.
- Owen E. Holland and Chris Melhuish. Stigmergy, Self-Organization, and Sorting in Collective Robotics. *Artificial Life*, 5(2):173–202, 1999.
- Marco Iacoboni, Roger P. Woods, Marcel Brass, Harold Bekkering, John C. Mazziotta, and Giacomo Rizzolatti. Cortical Mechanisms of Human Imitation. *Science*, 286(5449):2526–2528, December 1999.
- Marc Jeannerod and Jean Decety. Mental motor imagery: a window into the representational stages of action. *Current Opinion in Neurobiology*, 5:727–732, 1995.
- Odest C. Jenkins and Maja J. Matarić. Performance-Derived Behavior Vocabularies: Data-driven Acquisition of Skills from Motion, 2004. To appear in the International Journal of Humanoid Robotics.
- Odest C. Jenkins, Monica N. Nicolescu, and Maja J. Matarić. Autonomy and supervision for robot skills and tasks learned from demonstration, 2004. To appear in the AAAI-04 Workshop on Supervisory Control of Learning and Adaptive Systems.
- Kazuhiko Kawamura, R. Alan Peters II, Robert E. Bodenheimer, Nilanjan Sarkar, Juyi Park, Charles A. Clifton, Albert W. Spratley, and Kimberly A. Hambuchen. A parallel distributed cognitive control system for a humanoid robot. *The International Journal of Humanoid Robotics*, 1(1):65–93, March 2004.
- C. Ronald Kube and Hong Zhang. Collective robotics: From social insects to robots. *Adaptive Behavior*, 2(2):189–219, 1993.
- Sridhar Mahadevan and Jon Connell. Automatic Programming of Behavior-Based Robots Using Reinforcement Learning. *Artificial Intelligence*, 55((2-3)):311–365, 1992.
- Matthew J. Marjanović, Brian Scassellati, and Matthew M. Williamson. Self-Taught Visually-Guided Pointing for a Humanoid Robot. In Pattie Maes, Maja J. Matarić, Jean-Arcady Meyer, Jordan Pollack, and Stewart W. Wilson, editors, *Proceedings, From Animals to Animats 4, Fourth International Conference on Simulation of Adaptive Behavior (SAB-96)*, pages 35–44. MIT Press, 1996.
- Maja J. Matarić. Integration of Representation into Goal-Driven Behavior-Based Robots. *IEEE Transactions on Robotics and Automation*, 8(3):304–312, June 1992.
- Maja J. Matarić. Designing and Understanding Adaptive Group Behavior. *Adaptive Behavior*, 4(1):51–80, December 1995.

- Maja J. Matarić. Behavior-Based Control: Examples from Navigation, Learning, and Group Behavior. *Journal of Experimental and Theoretical Artificial Intelligence*, 9(2-3):323-336, 1997a. special issue on Software Architectures for Physical Agents.
- Maja J. Matarić. Learning Social Behaviors. *Robotics and Autonomous Systems*, 20(2-4):191-204, June 1997b.
- Maja J. Matarić. Reinforcement Learning in the Multi-Robot Domain. *Autonomous Robots*, 4(1):73-83, March 1997c.
- Maja J. Matarić. Behavior-Based Robotics as a Tool for Synthesis of Artificial Behavior and Analysis of Natural Behavior. *Trends in Cognitive Science*, 2(3):82-87, March 1998.
- Maja J. Matarić. Learning in Behavior-Based Multi-Robot Systems: Policies, Models, and Other Agents. *Cognitive Systems Research*, 2(1):81-93, April 2001. special issue on Multi-disciplinary studies of multi-agent learning.
- Maja J. Matarić. Sensory-Motor Primitives as a Basis for Imitation: Linking Perception to Action and Biology to Robotics. In Kerstin Dautenhahn and Chrystopher L. Nehaniv, editors, *Imitation in Animals and Artifacts*, pages 391-422. MIT Press, Cambridge, MA., USA, 2002a.
- Maja J. Matarić. Situated Robotics. In *Encyclopedia of Cognitive Science*. Nature Publishing Group, Macmillan Reference Limited, November 2002b.
- Maja J. Matarić and Marc Pomplun. Fixation Behavior in Observation and Imitation of Human Movement. *Cognitive Brain Research*, 7(2):191-202, October 1998.
- Maja J. Matarić, Victor B. Zordan, and Zachary Mason. Movement Control Methods for Complex, Dynamically Simulated Agents: Adonis dances the Macarena. In *Autonomous Agents*, pages 317-324, Minneapolis/St. Paul, MN., USA, May 1998. ACM Press.
- Maja J. Matarić, Victor B. Zordan, and Matthew M. Williamson. Making Complex Articulated Agents Dance: an analysis of control methods drawn from robotics, animation and biology. *Autonomous and Multi-Agent Systems*, 2(1):23-44, June 1999.
- Ferdinando A. Mussa-Ivaldi and Simon F. Giszter. Vector field approximation: a computational paradigm for motor control and learning. *Biological Cybernetics*, 67(6):491-500, 1992.
- Ferdinando A. Mussa-Ivaldi, Simon F. Giszter, and Emilio Bizzi. Linear combination of primitives in vertebrate motor control. *Proceedings of the National Academy of Sciences of the United States of America*, 91:7534-7538, August 1994.
- Monica Nicolescu. *Natural Methods for Learning and Generalization in Human-Robot Domains*. PhD thesis, University of Southern California, Computer Science Department, 2002. Also available as *Center for Robotics and Embedded Systems* Technical Report CRES-02-006.
- Monica Nicolescu and Maja J. Matarić. Learning and Interacting in Human-Robot Domains. *IEEE Transactions on Systems, Man, Cybernetics, special issue on Socially Intelligent Agents - The Human in the Loop*, 31(5):419-430, September 2001.
- Monica N. Nicolescu and Maja J. Matarić. A Hierarchical Architecture for Behavior-Based Robots. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AGENTS)*, pages 227-233, Bologna, Italy, July 2002.
- Stefano Nolfi, Dario Floreano, Orazio Miglino, and Francesco Mondada. How to evolve autonomous robots: Different approaches in evolutionary robotics. In *Artificial Life IV. Proc. Fourth International Workshop on the Synthesis and Simulation of Living Systems*, pages 190-197, Cambridge, MA., USA, July 1994.
- Anand Panangadan, Maja J. Matarić, and Gaurav S. Sukhatme. Identifying Human Interactions in Indoor Environments. In *Autonomous Agents and Multi-Agent Systems*, New York City, NY., USA, July 2004.
- Paolo Pirjanian. Behavior Coordination Mechanisms: State-of-the-Art. Technical Report IRIS-99-375, Institute for Robotics and Intelligent Systems, University of Southern California, October 1999.

- Giacomo Rizzolatti, Luciano Fadiga, Vittorio Gallese, and Leonardo Fogassi. Premotor cortex and the recognition of motor actions. *Cognitive Brain Research*, 3(2):131–141, March 1996.
- Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice-Hall, Inc., Upper Saddle River, NJ., USA, second edition, 2002.
- Srikanth Saripalli, David J. Naffin, and Gaurav S. Sukhatme. Autonomous Flying Vehicle Research at the University of Southern California. In Alan Schultz and Lynn E. Parker, editors, *Multi-Robot Systems: From Swarms to Intelligent Automata, Proceedings of the First International Workshop on Multi-Robot Systems*, pages 73–82. Kluwer Academic Publishers, 2002.
- Brian Scassellati. Investigating Models of Social Development Using a Humanoid Robot. In Barbara Webb and Thomas R. Consi, editors, *Biorobotics*. MIT Press, August 2001.
- Stefan Schaal and Dagmar Sternad. Programmable Pattern Generators. In *Third International Conference on Computational Intelligence in Neuroscience*, pages 48–51, Research Triangle Park, NC., USA, October 1998.
- Dylan A. Shell and Maja J. Matarić. Human motion-based environment complexity measures for robotics. In *Proceedings IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'03)*, pages 2559–2564, Las Vegas, NV., USA, October 2003.
- Dylan A. Shell and M.J. Matarić. Directional audio beacon deployment: an assistive multi-robot application. In *Proceedings IEEE/RSJ International Conference on Robotics and Automation (ICRA'04)*, pages 2588–2594, New Orleans, LA., USA, April/May 2004.
- Barbara Webb. Robotic experiments in cricket phonotaxis. In Dave Cliff, Philip Husbands, Jean-Arcady Meyer, and Stewart W. Wilson, editors, *Proceedings, From Animals to Animats 3, Second International Conference on Simulation of Adaptive Behavior (SAB-94)*, pages 45–54, Brighton, UK., August 1994. MIT Press.
- Matthew M. Williamson. Postural Primitives: Interactive Behavior for a Humanoid Robot Arm. In Pattie Maes, Maja J. Matarić, Jean-Arcady Meyer, Jordan Pollack, and Stewart W. Wilson, editors, *Proceedings, From Animals to Animats 4, Fourth International Conference on Simulation of Adaptive Behavior (SAB-96)*, pages 124–131. MIT Press, 1996.
- Helen Yan and Maja J. Matarić. General Spatial Features for Analysis of Multi-robot and Human Activities from Raw Position Data. In *Proceedings IEEE/RSJ International Conference on Robotics and Intelligent Systems (IROS'02)*, pages 2770–2775, Lausanne, Switzerland, October 2002.