Sensory-Motor Primitives as a Basis for Imitation: Linking Perception to Action and Biology to Robotics

Maja J Matarić
Computer Science Department and Neuroscience Program
University of Southern California
941 West 37th Place, SAL 228, MC 0781
Los Angeles, CA 90089-0781
mataric@cs.usc.edu

X.1 Introduction

Imitation is a powerful means of skill acquisition, social interaction, and cultural transfer. Although popular notions of "monkeying" and "parroting" make the phenomenon appear ubiquitous, the more formal definition of "true imitation" limits it to very few species, as it requires the ability to acquire arbitrary new skills from observation (Tomasello, Kruger & Rather 1993, Byrne & Russon 1998). The difficulty in teasing true imitation apart from related phenomena, including stimulus enhancement, social facilitation, emulation, and priming, indicates that the underlying mechanism for imitation likely involves a complex interaction of learning mechanisms. We consider imitation to be a broad capability based on a spectrum of behavior and learning systems of increasing complexity, only the most sophisticated of which correspond to the definition of true imitation.

In our work, we view imitation as the capability to acquire new skills by observation, based on the the imitator's existing behavioral repertoire, which may be simple or sophisticated. Specifically, we focus on motor behaviors, and on their underlying mechanisms. In this chapter, we present a model of imitation that results from an incrementally evolved capability based on an evolutionarily old motor control structure (motor primitives), utilized by a mechanism for movement understanding and simple mimicry (mirror neurons), driven by a specialized attentional mechanism. The entire system is based on the notion of *primitives*, structures that directly link the visual and motor systems in the context of complete generalized movements. The capacity for mimicry is based on this direct mapping mechanism, and itself serves as a basis for more complex, less direct forms of imitation and social learning. Thus primitives are the fundamental substrate for a powerful system that couples visual perception and motor execution, and unifies recognition, classification, and learning of movement skills.

A "primitive" can refer to something that is old in an evolutionary sense, such as structures that organize the underlying mechanisms of movement, including spinal fields (Bizzi, Mussa-Ivaldi & Giszter 1991) and central pattern generators (Stein 1997). In a computational sense, primitives can be viewed as a basis set of motor programs that are sufficient, through combination operators, for generating the entire movement repertoire. Our imitation model combines the two interpretations; it involves a system that is innately primed with a set of primitives sufficient for a large movement repertoire, and is capable of learning arbitrary new skills that can, in a hierarchical fashion, be used as new compositional elements.

Our model is inspired and constrained by the following lines of evidence from psychophysics and neuroscience:

1) The existence of a *selective attention* mechanism for extracting biological movement information from the visual stream. The mechanism in our model is based on data from attentional studies and our own eye-tracking work, described in Section X.4.

2) The existence of *mirror neurons*, which map observed movements into executable ones, thus performing sensory-motor integration and transformation. The mechanism in the model is driven by the features from the attentional mechanism and maps onto the corresponding motor primitive(s), discussed in Section X.3.

3) The existence of *motor primitives* for structuring movement, constituting a motor vocabulary composable into a broad movement repertoire. In the model, the primitives are the general basis for movement, invoked by the mirror neurons as a part of the direct mimicry system underlying imitation, as discussed in Section X.5.

4) The existence of a *classification-based learning* system. In the model, this system encodes new motor skills by using sequencing and superposition of motor primitives, as described in Section X.6.

We bring these four lines of evidence together in a general imitation model involving movement perception, learning, and reproduction. The model is based on a direct mapping between the sensory and motor systems, organized in terms of primitives. We do not model the involved mirror and motor structures at the neural level. Instead, we propose an information processing model based on key structural and organizational constraints arising from neuroscience and physchophysical evidence. This level of modeling allows us to experiment with various specific implementations to validate our theories and test their practical applications. As our goal is to gain insight into imitation in natural systems and enable it in artificial ones, we relate the components of our model to relevant work in neuroscience, cognitive science, and robotics.

The rest of the chapter is organized as follows. We begin in Section X.2 with the motivation for this work, which bridges the fields of cognitive science and neuroscience to robotics. Next, in Section X.3, we survey the evidence from neuroscience, focusing on mirror neurons, and touch on relevant inspiration from cognitive science and developmental psychology. In Section X.4, we review work on movement perception, and describe the results of our eye-tracking study that serve as the basis for the attentional mechanism in our model. In Section X.5 we discuss motor control in biology and robotics, focusing on motor primitives. Section X.6 brings together the discussed components of the model into a unified framework. Section X.7 describes our experimental methods and test-beds, and summarizes some of the experimental results. Section X.8 gives an overview of other imitation work in robotics. Section X.9 summarizes the chapter and discusses continuing work.

X.2 Motivation

Before describing how we model imitation, we first address our dual motivation for pursuing it. Imitation is a challenge for both natural and artificial systems, as it involves the integration of multiple cognitive/computational systems, namely perception, memory, action, and learning. Thus, both analysis and synthesis of imitation are difficult, and any insights gained are likely to be relevant to both biology and robotics.

Historically, the majority of work on imitation in nature emphasized precise classification of the phenomenon into distinct types and stages (Piaget 1962, Davis 1973, Tomasello et al. 1993) and addressing its role in social behavior (Bandura & Walters 1963, Bandura 1977, Meltzoff & Moore 1994). More recently, the evolution of imitation, and its relation to other forms of learning are being explored (Moore 1996, Moore 1992). Our work bears on the evolutionary history of imitation in that it outlines a potentially phylogenetically old mechanism as the core of the phenomenon.

Imitation is not only a powerful learning mechanism, but is also thought to be a keystone in the evolution of communication and language (Arbib 2000, Arbib & Rizzolatti 1996, Rizzolatti, Gadiga, Gallese & Fogassi 1996*b*, Jeannerod, Arbib, Rizzolatti & Sakata 1995, Donald 1993). Thus, it "vertically" integrates cognitive systems from the lowest levels of perception and motor control to the highest levels of cognition. The notion of vertical integration is a foundational principle of *behavior-based robotics*, a branch of robotics and Artificial Intelligence that aims at developing autonomous, adaptive, embodied systems that exist in the physical world and cope with the demands in the "here and now" (Brooks 1986, Brooks 1991, Matarić 1997, Arkin 1998). Unlike traditional robotics, where thinking (computation) takes more time than does action, behavior-based robotics focuses on fully integrated systems that act in real time in their environment.

Imitation presents a potentially powerful mechanism for automated programming and control of robots. In the past, the idea of "teaching by showing" or "learning by demonstration" was limited to reproducing the goal or output state of the activity, but not the complete process that achieved it, or the underlying mechanism. In contrast, current efforts focus on modeling the imitation mechanism itself, in order to develop a general tool for improved robot control and more natural interaction with robots. Section X.8 provides a more detailed overview of relevant work on imitation in robotics. In addition to programming robots more directly and efficiently, a functional model of imitation has a variety of other pragmatic applications for more natural human-machine interaction. An important one is aiding skill (re)learning in children, athletes, the handicapped, and those recovering from stroke or injury. Thus, the motivation for studying imitation comes from the challenge of understanding this complex and powerful natural phenomenon, as well as applying it for programming of artificial system toward a more natural human-machine interaction.

X.3 Inspiration

A strong connection between perception and action is a key component of our imitation model, and is inspired by evidence from neuroscience, cognitive science, and developmental psychology. Decety (1996) and Jeannerod & Decety (1995) demonstrated that a shared neural substrate is used both in imagined and executed movements. Subjects involved in imagination and visualization of motor tasks demonstrated activation of motor pathways similar to that during execution of movement. Similarly, increased activity in the motor cortex was measured during observation of movement, both in humans watching another human (Jeannerod et al. 1995, Fadiga, Fogassi, Pavesi & Rizzolatti 1995) and in monkeys watching a human (di Pellegrino, Fadiga, Fogassi & Rizzolatti 1992).

The direct connection between movement perception and its execution is supplied by the so-called mirror neurons found in area F5 of the monkey pre-motor cortex (Rizzolatti et al. 1996*b*). These neurons fire when the monkey observes a specific movement as well as when it executes the same movement. Since their discovery, mirror neurons have become a subject of intense study, as they are a natural candidate for neural theories of imitation. It is reasonable to postulate that the function these neurons perform, the connection between observed and performed behavior, is the substrate for imitative capability. The mirror system likely enables mimicry in monkeys, who are not thought capable of "true imitation". The presence of mirror neurons in other species has not yet been tested extensively. However, similar facility was found in the left Broca's area (area 45) and the dorsal pre-motor area 6 of the human brain (Iacoboni, Woods, Brass, Bekkering, Mazziotta & Rizzolatti 1999, Rizzolatti, Fadiga, Matelli, Bettinardi, Perani & Fazio 1996*a*). It is currently not known what movements are represented by the mirror system (besides the few that were tested directly), and whether the movements are learned or innate.

The neuroscience evidence discussed so far has addressed the connection between sensory and motor systems. Psychophysical experiments addressing the cognitive aspect of this connection provide complementary evidence. Vogt (1995) tested human subjects for the existence of an intermediate coding level between perception and action during an imitation task consisting of copying drawings with closed eyes. Subjects showed no improvement after immediate rehearsal, and no effect from distractions, possibly due to a lack of intermediate coding and the existence of a generative perception module that automatically prepares for action. Similarly, Pelisson, Goodale & Prablanc (1978) showed that adjustments of hand pointing to visual targets do not require visual information about the moving arm, implying the existence of internal models of movement that transform visual information into motor representations.

Further evidence supporting sensory-motor integration and generative perception can be found in lesion studies (Goodale, Jakobson, Milner, Perrett, Benson & Hietanen 1994), studies of athletic performance (Abernethy 1990, Elkins 1996, Savelsvertgh & Bootsma 1994), and in developmental psychology (Bertenthal 1996). Developmental psychology work in infant imitation also supports sensory-motor integration as the basis of imitation. Meltzoff and Moore argue that imitation is a fundamental human capability, found in newborns and based on an innate link between perceptual and motor systems (Meltzoff & Moore 1988, Meltzoff & Moore 1994), enabling young children to imitate facial expressions and hand movements without visual feedback. They hypothesize an imitation mechanism involving innate, integrated, supra-modal representations of human movements and postures (Meltzoff & Moore 1995, Meltzoff & Moore 1997).

The above biological evidence serves as powerful inspiration for our philosophy and approach to modeling imitation. We postulate that imitation is based on a simpler mimicry system encoded through mirror neurons and involving direct mappings between observed and executable movements. The executable movements are encoded using the phylogenetically older system of motor primitives. Thus, mimicry involves mapping observed movements to an existing motor primitive repertoire. Imitation, in contrast, involves a more complex process capable of creating arbitrary new motor skills by composing complex hierarchical combinations of primitives and other already learned skills.

Our imitation model thus employs a series of cognitive systems that build upon each other in precisely the incremental fashion one would expect from an evolved capacity. The complete system utilizes four key components: 1) movement perception through a specialized selective attention system, 2) direct sensory-motor mapping between the perceived and executable movement, 3) movement generation through a structured system of composable motor primitives, and 4) learning of new movements and skills by building on the existing repertoire of primitives, through classification and combination. Each of these components is now addressed in detail.

X.4 Movement Perception

We postulate that imitative capability takes advantage of a specialized attentional system attuned to biological movement. This system predates imitation, and possibly evolved for social interaction with conspecifics. Its ability to recognize biological movement and, given a complex visual input, selectively attend to the areas containing task-related information, is a key aspect of imitation. We will propose that this specialized movement perception and attention system is not only driven by external visual input, but also by the motor primitives, which serve as predictors of what movement will occur next.

To explain and support this claim, we overview relevant evidence from neuroscience, and then describe an eye-tracking study we conducted with the goal of addressing precisely the question of selective attention in imitation.

Neuroscience Evidence

Johansson (1973) demonstrated the human capability to recognize biological motion from a small number of structured visual cues such as light dots in key locations on the limbs. This well-known experiment points toward a specialized mechanism for body-motion detection. Perrett and co-authors performed a series of studies of movement perception in the macaque monkey. The results support the existence of specialized neural detectors and predictors for specific postures, movements, and goal-oriented behaviors. Perrett, Smith, Mistlin, Chitty, Head, Potter, Broennimann, Milner & Jeeves (1985) show macaque neurons that selectively respond to specific movement types and stimulus forms, such as particular movements of the whole body or body parts. Perrett, Harries, Bevan, Thomas, Benson, Mistlin, Chitty, Hietanen & Ortega (1989*a*) demonstrate similar, view-independent results, implying the potential for viewer-centered and goal-centered descriptions of the movement. Perrett, Harries, Mistlin & Chitty (1989*b*) describe a population of neurons involved in the recognition of specific actions of others, and Perrett, Harries, Mistlin, Hietanen, Benson, Bevan, Thomas, Oram, Ortega & Brierley (1990) show that expectation of movement indirectly correlates with the amount of resulting neural firing. Taken together, this evidence supports the existence of a specialized biological movement detection system also capable of prediction.

Eye-Tracking Evidence

Little is known about visual attention involved specifically in imitation. Based on our interest in visuo-motor integration theories of imitation, we conducted a set of eye-tracking experiments aimed at addressing the following questions:

1. Is there a difference between watching a movement with the intention to imitate and just watching?
2. When watching with the intention to imitate, what features are fixated on?

In a standard cognitive model of perception and motor control, one would expect the answer to the first question to be positive, denoting different processing involved in observation alone compared to observation combined with or followed by motor planning. However, if we adopt a strong sensory-motor integration assumption, which postulates a direct mapping between movement observation and reproduction, we would expect to see no difference in fixation patterns between the imitation and no-imitation conditions. The expectation would be that, in both cases, the sensory-motor system maps the observed movement to the executable motor programs; in the former condition the execution of the movements results in imitation while in the latter movement is inhibited.

The second question was aimed at gaining insight into the attentional mechanism, and its relation to the underlying motor control system. Given that subjects have limited attentional resources and are thus forced to fixate selectively, we were interested in finding out what they would attend to and how the visual information fixated on would serve as a basis for subsequent imitation.

In the experiment, fixation patters of 40 subjects were recorded while being shown 25 3-second long videos of different types of unfamiliar finger, hand, and arm movements on a computer screen. Before each set of videos, the subjects were told whether they would subsequently imitate each video after viewing it. The computer screen was black between the stimuli and during the imitation phase. In the imitation condition, the subject's right arm was extended and imitation took place without visual feedback.

Regardless of the type and size of presented stimulus, and regardless of whether the subjects intended to subsequently imitate, they attended to the hand most frequently and for the longest time periods; fixation frequency and duration were both maximized at the hand. The only difference between the subjects' behavior between the imitation and no-imitation conditions was in pupil dilation. In the imitation condition, dilation was significantly larger, while the fixation behavior remained the same in both conditions. Finally, in spite of focusing overt attention mostly on the hand, subjects were effective at imitating the presented movements. For a detailed description of the experiments, analysis, and results, see Matarić & Pomplun (1998).

These data provide a decisive negative answer to the first question, testing for the existence of any difference between observation with and without intention to imitate. Thus, the results add overt fixation behavior measured through eye-tracking to data from other modalities that all support the direct sensory-motor mapping theory. The results are consistent with the neuroscience and developmental psychology work discussed above, all supporting generative perception of movement as a general property of visuo-motor processing, and thus of imitation.

The data also allow some interesting interpretations for the second question above, asking what the location of fixation tells us about the underlying motor control mechanism. The results demonstrating consistent selective attention on the hand allow for several hypotheses that impact imitation

1. Perception of biological motion focuses overt attention at the end-effector; this insight can be used to greatly simplify models of attention.

2. Internal models, possibly in the form of motor primitives, are used to fill in the motor control details which are not directly observed.

3. Those internal models also influence the attentional mechanism by providing movement prediction or expectation.

4. An efficient transformation between the visual (extrinsic) and motor (intrinsic) representations and coordinate frames converts the observed end-effector information into executable movement.

The results of this study, in conjunction with other evidence we overviewed, serve as the basis for the attentional mechanism of our imitation model, described in Section X.6. The data also support our primitives-based framework, described in the next section.

X.5 Motor Control and Motor Primitives

Our model postulates that the motor system is organized into a collection of primitives, motor programs that encode complete stereotypical movements. These can be executed sequentially or concurrently (through parametrized superposition), to create a large movement repertoire. This general organization of the motor system is then used by the mirror system (which is phylogenetically newer) to instantiate direct sensory-motor mappings into executable motor programs. In this section we describe the neuroscience inspiration for the motor primitives aspect of our model.

Understanding the modular organization of biological movement and implementing it in artificial systems are outstanding challenges in biology and robotics. The theory of motor primitives has served as an inspiration for work both in behavior-based control in general and in robot models of imitation in specific. Mussa-Ivaldi & Giszter (1992) and Giszter, Mussa-Ivaldi & Bizzi (1993) suggested the existence of spinal force-field motor primitives that converge to single equilibrium points and produce high-level behaviors such as reaching and wiping. In experiments with spinalized frogs and rats, when the spine is stimulated with an electrode, a particular field is activated, causing the leg to execute a complete behavior and come to rest at the field's equilibrium point. When two or more fields are stimulated at the same time, either a linear superposition of the fields is obtained, or one dominates (Mussa-Ivaldi, Giszter & Bizzi 1994). In either case, a meaningful movement of the leg results. Only a small number of such distinct fields was found in the frog's spine; thus only a dozen primitives may be sufficient for coding the frog's entire motor repertoire, through sequencing and superposition of supra-spinal inputs through descending pathways (Bizzi et al. 1991).

This suggests elegant and modular organization for motor control, in which entire behaviors are coded with primitives whose composition results in more complex motor outputs. Our previous work (Matarić & Marjanović 1993, Matarić 1995, Matarić 1997) used this organization to structure motor control of single and multiple mobile robots into a collection of *basis behaviors* that fit the notion of primitives directly into the modular behavior-based control paradigm (Arkin 1998). The approach was demonstrated to be effective for various individual and group behaviors, including navigation, foraging, following, flocking, and homing. Arkin (1987) applied schema theory (Arbib 1992) to behavior-based mobile robots and employed a similar notion of composable behaviors, also stemming from neuroscience foundations (Arbib 1981).

The idea of using primitives for motor control has also been studied in articulated, humanoid robotics. Williamson (1996) and Marjanović, Scassellati & Williamson (1996) developed a 6 DOF (degrees of freedom) robot arm controller. In contrast to biological and mobile robotics movement primitives, theirs coded for four static postures, three for reaching and for resting. Interpolation was applied to reach any desired end-point position within the 3D space encompassed by the primitives. Schaal & Sternad (1998) used nonlinear attractor dynamics to create two types of motor primitives, and demonstrate them on a rhythmic drumming task with a 7 DOF robot arm.

Our own previous work on articulated humanoid control (Matarić, Zordan & Williamson 1999), applied to the control of 7 DOF arms of a physics-based 20 DOF humanoid simulation, implemented two versions of the force-field primitives. One closely modeled the frog data described above, using an intrinsic, joint-space representation. The other used another biologically-inspired approach, impedance control (Hogan 1985), which operates in the extrinsic end-point (hand) coordinate frame. Both were shown to provide an effective substrate for control of a complicated sequential motor task, but each had limitations in terms of ease of control for certain types of movements. We proposed a combination approach, employed in our current model and discussed in the next section.

The issue of coordinate frame choice is just one of the many that must be resolved in order to practically apply the primitives-based approach to motor control. While an expressive and composable vocabulary of motor primitives is an appealing idea, *what* the primitives in such a vocabulary should be remains a difficult problem. In the next section we discuss this issue, describe our choice of primitives, and their role in the imitation model.

X.6 Our Model

Attempts to endow robots with biologically-inspired imitation capabilities are recent, and those aimed at studying biological imitation through robotics are few. Our effort brings together insights and methods from these two fields into a testable model of imitation with practical applications. We combine the discussed biological inspirations into the following model components:

1) A selective attentional mechanisms for extracting salient movement information from the visual stream by focusing on the effectors (hands or tools). We assume that this mechanism is based on an evolutionarily old capacity to recognize kin and conspecifics (Parr & de Waal 1999) and interpret their actions and situation-specific intentions (Cheney & Seyfarth 1990).

2) A mirror-like sensory-motor mapping system as a means of representational integration and transformation from visual input into executable motor programs. This mechanism performs the functions variously referred to as direct matching, mirroring, or resonance (Iacoboni et al. 1999, Rizzolatti et al. 1996*b*), found in the monkey and human.

3) A motor control system consisting of primitives that constitute the basic motor vocabulary composable into a broad movement repertoire through sequencing and superposition. These can be thought to perform the function of spinal force-field primitives (Bizzi et al. 1991) as well as cortical central pattern generators (Brooks 1981, Stein 1997).

4) A classification-based learning mechanism that learns from the match between the observed and executable movements and is biased toward parametrized reuse of the existing motor programs through hierarchical composition.

Figure 1 shows the model, with the information processing flow indicated among its main components. We now describe each of the model components in turn.

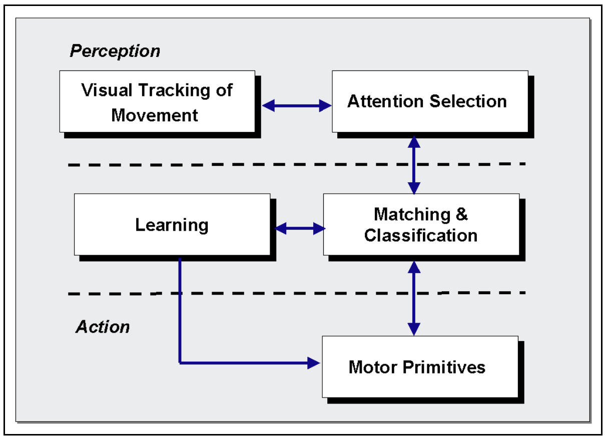

Figure 1: *The information processing flow in our imitation model, structured around perceptual-motor primitives used for movement segmentation, classification, and generation.*

Attention in Movement Tracking

Deciding what to pay attention to for subsequent imitation is a difficult problem. Based the literature on visual attention (Pashler 1999), and on our own eye-tracking results (Matarić & Pomplun 1998), we assume that the focus of attention is not driven directly by the intention to imitate. Instead, in our model we postulate an attentional mechanism that takes into consideration both intrinsic bottom-up image features (such as velocity changes and the position of the end-point) and extrinsic, top-down task goals (imitation or otherwise) in selecting what to attend to.

Our attentional mechanism performs a segmentation of the retinal space, then evaluates the current interest level of each feature, and finally displaces the center of attention accordingly. We transform all of the features into an egocentric coordinate frame (Stein 1992, Flanders, Tillery & Soechting 1992, Flanders & Soechting 1995) suitable for subsequent higher-level processing. The frame is initially anchored on a fixed feature (e.g., the head) or on a visual landmark, if one is observed. The retinal segmentation is fovea-based, using a log-polar grid (Hubel & Wiesel 1977). As the distance between a feature and the center of attention increases, the resolution of the feature decreases.

The evaluation of the current interest level for a feature is based on three factors: retinal movement, captivation, and task-level bias. Retinal movement is the amount a feature has recently moved in the retinal image. Captivation is based on the recent interest level of the feature, and on the overall level of "boredom", inversely proportional to the novelty of the image. A feature loses interest when it is given attention and produces little or no newly occupied cells in the retina. Finally, task-level bias is any preference for specific features (e.g., the feet for learning dance steps). The three factors are normalized and weighted so that retinal movement is the basic driver of attention, sufficient captivation supersedes it, and task-level bias, when present, overrides both.

Retinal movement, captivation, and task-level bias represent dimensions in feature-interest space. Each is represented as a point in that space, with its magnitude corresponding to its interest level. Displacement of the center of attention is based on the computed and thresholded interest levels of the features and their distance, in interest space, from the current center of attention. The feature with the most interest captures the focus of attention and inhibits the others. If no features meet the criteria, attention de-focuses on either the entire body or a subset of recently interesting features. Figure 2 illustrates the attentional system. More detail about the attentional part of the model can be found in Jenkins (2000).

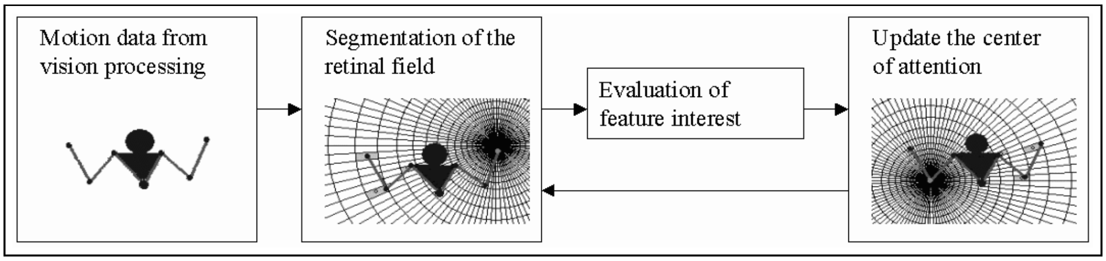

Figure 2: *The attentional system.*

The selected features serve as the basis for movement classification into primitives.

Primitive-Based Sensory-Motor Integration

Combining perceptual and motor properties into a unified representation is a key challenge of the direct mapping hypothesis that drives our work. The motor primitives in this model, much like the basis behaviors in our previous work (Matarić 1995, Matarić & Marjanović 1993), are derived from a combination of top-down and bottom-up constraints, reflecting evolutionary interaction between morphology and behavior. The bottom-up constraints stem from the physical properties of the motor system, its kinematics and dynamics. The top-down constraints are shaped by the motor behaviors the system performs, based both on the evolutionary bias for a given species and by the specific skills learned during lifetime.

At the perceptual end, primitives are shaped by the structure of the visual system and the inputs it receives. In general, this input is underconstrained and defies classification into a parsimonious set of task-independent categories. Fortunately, the mirror system is specialized for mapping observations of biological movements, in particular those from systems with similar morphology, i.e., conspecifics. As a result, the driving forces that shape the primitive matching system from both the perceptual and motor sides are consistent.

Our approach does not involve a detailed model of the mirror system as insufficient neuroscience data is currently available. Instead, we focus on the selection of the motor primitives, and the way in which they fundamentally shape movement perception, discussed in the following sections.

What Type of Primitives?

Research in motor control typically addresses either discrete or oscillatory movements, but rarely both within a unified framework. A primitive-based motor control system must account for both types and provide an efficient means of encoding both.

We have proposed three types of primitives: 1) discrete, consisting of unobstructed straight-line movements based on reaching; 2) oscillatory, consisting of repetitive movements based on central pattern generators; and 3) postural, involving large subsets or whole-body arrangements of joints (Demiris & Mataric 1998). Assuming this division, we implemented each primitive in the most "natural" frame of reference, i.e., one that minimized coordinate transformation computations. We used intrinsic (joint) space representations for postural and oscillatory movements, and extrinsic (Cartesian) space for discrete ones (Matarić et al. 1999). We have also explored various representations, including convergent vector fields in both joint and Cartesian space, impedance control (well suited for discrete movements), interpolated joint-space control, and central pattern generators (CPGs). To address the problem of self-collisions, we introduced an *avoidance* primitive, which was superimposed with any currently executing primitive (Matarić, Zordan & Mason 1998). These primitives were validated in a task of executing a dance (Matarić et al. 1999).

The discrete and oscillatory primitives are both driven by the end-effector movement(s) of the demonstrator, while postures require task-based bias to draw attention to the specific body parts. We postulate that, from an evolutionary perspective, imitation of postures is less useful than imitation of end-effector behaviors, which include the use of tools and gesture-based communication. Since these convey a selective advantage, they could have resulted in the adaptation of the mirror system, the mechanism that performs the direct mapping between the observed and executed end-effector movements.

One of the issues we are exploring with this work concerns the extent to which end-point trajectories are sufficient for movement imitation, given that the attentional mechanism provides that information with highest fidelity. In extensions of the above work, we are exploring the encoding of movement ranging from end-point only, end-point with added postural constraints, and finally adding the wrist and elbow positions. The last presents a higher-dimensional problem and an added challenge for an invariant primitive representation. We are currently exploring encodings involving wavelets and principal components analysis (PCA) methods for reducing the dimensionality of the movement data and developing a more general representation. Human planar discrete movements have been shown to reduce to a small number of principal components (Sanger 2000); we are pursuing the question of how that property scales to free movements (in 3D) and to movement reconstruction.

The motor primitives in our model are a set of parametric motor controllers capable of generating stereotypical discrete, oscillatory, and postural movements. We first address the issue of parametrization, and then the more overarching question of what constitutes a "stereotypical" movement, i.e., what and how many primitives should there be.

Parametrizing Primitives

The parameters in the motor primitive representation represent the invariants in the perceptuo-motor mapping: in order to keep the primitives general, we make them invariant to exact Cartesian position, rate of motion, size, and perspective. However, in order to produce an accurate imitation, these values must be correctly reconstructed. Thus, they are provided directly to the motor primitives from the perceptual system as parameters, but are not stored within the primitive representation. As a result, the primitives determine the high-level, generic description and parametrization of an observed movement; the specific parameter values are used to instantiate the generic primitive description into a specific executable trajectory.

The generic encoding of parametric primitives provides a parsimonious representation for control. From a robotics perspective, the notion of endowing a system with a collection of behavioral primitives and thus avoiding on-line run-time trajectory planning is very appealing, and is gaining popularity (Matarić et al. 1999, Schaal & Sternad 1998). Our past work has addressed the process of designing primitives, or basis behaviors, for a specific system and set of tasks in the domain of mobile robotics (Matarić 1995) and humanoid articulated control (Matarić et al. 1999). Such "innate" primitives are effective but are determined by the designer; we next discuss how they can be automatically generated.

Learning Primitives from Movement Data

To study the question of what the primitives should be in a bottom-up fashion, we are using three types of human movement data: 1) visual data of torso movement, 2) 3D marker data of the human arm, and 3) joint angle data from the full body (each described in more detail in the Experimental Validation Section). These data are given as input into statistical clustering methods, in order to explore any natural groupings contained therein. We are currently exploring two different clustering approaches. In one, we bias our clusters by psychophysical evidence about known movement properties. Thus, discrete and oscillatory movements are separated, by directly selecting reaching-type invariants in the data. In the second approach, we are exploring the means for converting the spatio-temporal movement data into a more abstract, symbolic form, and applying string matching algorithms to find patters therein. A key challenge of this approach is the reduction of the high dimensionality of the continuous movement data into an abstract form.

We postulate that given sufficient data, most of the clustering techniques should extract similar movement classes, based on the inherent kinematic and dynamic properties of the underlying movements. Once the primitive groupings are extracted, we apply generalization methods to create the desired higher-level, parametric representations for them, in order to use them as the basis set for observed movement classification.

Classification

Instead of assuming a one-to-one mapping between observed and executed movements, which greatly limits the imitative ability of a system, we favor a small set of additive basis primitives as a substrate. The classification mechanism finds the best fit between the observed behavior and any combination of the primitives, resulting in a richer mapping. Through learning, new skills are created, and themselves serve as primitives.

In the most basic version of the model, which corresponds to a simple interpretation of the mirror system, a small set of primitives is built in, and used to map a subset of observed behaviors in a one-to-one fashion. Reaching for objects, observed in the monkey mirror system (Rizzolatti et al. 1996*b*) for example, is a likely primitive in such a model, and one we use in our system. In a more general version of our model, one we have experimented with so far, the observed movements are mapped to a basis set of three primitives: discrete, oscillatory, and postural, and are parametrized and sequenced to represent more complex movements. We present results from implementations of this model in the Experimental Validation Section. The most general version of the model involves two extensions. The first is the superposition of primitives; so far we have only used superposition of the avoidance primitive with other behaviors. We are currently exploring representations for primitives that facilitate superposition. The second extension is learning new primitives, and using those as a part of an continually-expanded basis for classification and imitation, as described in the next section.

New Skill Learning

In our model, learning is a part of the comparison and classification processes that map the observed onto the known. A simple and direct match that generates a close approximation of the observed behavior results in the reinforcement of an existing skill, already encoded in the form of a parametrized combination of the primitives. On the other hand, a novel match that generates a significant performance error requires rehearsal and training. This organization leads us to the following specific expectations, consistent with the general performance of human skill acquisition:

1. The more similar skills are learned, the easier learning becomes, as the probability of a match to an existing motor program is higher.
2. Observation and learning of a skill similar to one that is known is faster than that of a completely new skill.
3. Refinement of a skill that is similar to an already known one is slower due to interference.

Thus, the process of skill learning by imitation in our model occurs during classification as well as during and after performance, based on the feedback through proprioception and external observation (by the learner and a teacher, if present). This type of learning lends itself to modeling with reinforcement-based methods (Kaelbling, Littman & Moore 1996, Sutton & Barto 1998). The interaction of the classification and rehearsal learning systems is an important part of our research.

When sufficient training is provided, a novel pattern becomes stereotyped and encoded as a part of the composable building blocks. This results in a hierarchical structure allowing classification to be done at a high level, if a similar skill already exists, or at a lower, more basic level, for novel observations. Such a hierarchy of increasingly more abstract representations allows for simpler and more compact coding of behaviors, and their use not only for motor execution but also for communication. This structure also facilitates symbol grounding, a key problem in Arificial Intelligence and robotics, involving the challenge of connecting high-level symbolic representations with low-level sensory and motor control (Harnad 1990).

X.7 Experimental Validation

We are validating our imitation model using a large body of input movement data, applied to a collection of experimental test-beds, in a series of imitation learning tasks, described next.

Human Movement Data

We briefly overview the methods we use for movement data gathering. We have developed a motion tracking system for extracting features from a video stream (Weber 2000). This system is simple and can easily be replaced by more sophisticated commercial versions, but is capable of selecting a collection of features from the moving image, based on a constrained (un-occluded and unambiguous) initial position and kinematic model of the object/body being imitated. Our work so far assumes that a human is being imitated, and the kinematic model used is that of a generic adult human. However, the system is general and can use any kinematic model. This greatly simplifies establishing the initial match between the features in the visual image and the underlying body. The match enables tracking of the body over time, using a planar rectangle representation of the limbs. This allows for fast computation and updating of limb position, and for simple prediction of future position, which is in turn used to speed up recognition. Input from the motor primitives could provide further predictive capability, but we are yet to properly explore it.

We have also gathered a large body of data of human arm movements using the FastTrak system that provides 3-D Cartesian position of the markers. These data were gathered in a collaborative experiment conducted at the National Institutes of Health Resource for the Study of Neural Models of Behavior, at the University of Rochester. Similar to our eye-tracking experiment (Section X.4), subjects watched and imitated short videos of arm movements, but were, in this case, also fitted with a set of markers on their imitation arm. Four markers were used, located near the shoulder, elbow, wrist, and middle finger joint. In addition to the subject data, marker data are also available for the demonstrator. These data are used for validation of the automatic primitive generation and primitives-based classification systems, as well as for testing implementations of the model and comparing its performance to that of human subjects on the same tasks.

Finally, we are also using full-body human movement data captured by the Sarcos SenSuit, a flexible exoskeleton that simultaneously measures 35 degrees of freedom, including the shoulder, elbow, wrist, hip, knee, waist, and ankle joints.

Experimental Test-Beds

We are using several distinct experimental test-beds to validate various implementations of our imitation model.

Physics-based humanoid simulation: Adonis, a physics-based humanoid simulation[1], is actuated from the waist up, and has 20 degrees of freedom (DOF): 3 in the neck, 3 in the waist, and 7 in each arm (3 in the shoulder and wrist, and one in the elbow). Adonis uses external sensing (such as vision input), and internal joint angle state sensors. This test-bed was used for our past motor control work and continues to be the closest test-bed to human body control. Figure 3 shows a snapshot of Adonis.

Humanoid avatars: Two COSIMIR[2] humanoid avatars are fully actuated and consist of 37 body DOF and 14 hand DOF in actuated fingers. Various models of the dynamics, muscles, and control can be applied. All sensing is external. Figure 4 shows a snapshot of the two avatars.

Robot dog: Laika, an Aibo[3] dog robot, has 18 degrees of freedom: 2 in the neck and 4 in each leg. It also has passively compliant paws. Its sensors include a color CCD camera, stereo microphones, joint encoders in the legs, and touch sensors on the paws and the head. The motivation for using Laika is to study imitation between different body morphologies and kinematics. Laika will be used to test our imitation model on imitating a human. Additionally, Laika's movements will be used as input to the other experimental test-beds. Figure 5 shows a snapshot of Laika.

[1]Developed in the Animation Lab at Georgia Institute of Technology, using S/D Fast.

[2]Developed at the University of Dortmund, described in Romann (1999)

[3]Developed and distributed by Sony.

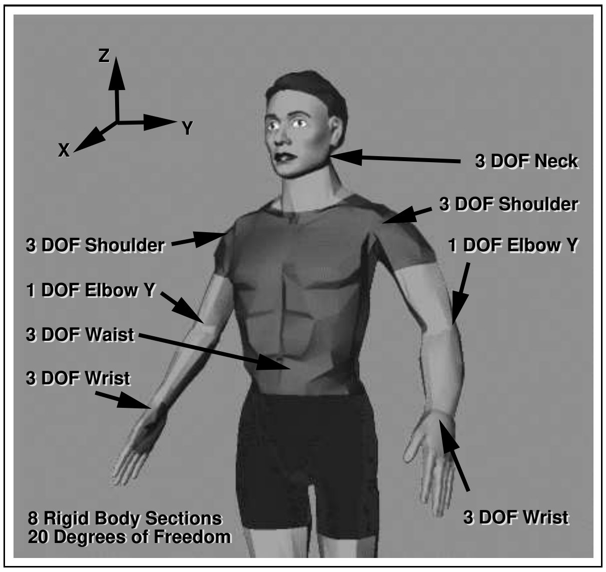

Figure 3: *Adonis, a 20 degree-of-fredom dynamic humanoid simulation test-bed used for validating our model.*

Wheeled mobile robots: Increasing numbers (currently 9) of Pioneer[4] mobile robots, equipped with pan and tilt color cameras, front and rear ultrasound sensors, and differentially steered bases, are used for group robotics research in our laboratory. Although these robots are too different from any biological system to be used for studying the direct mechanisms of true imitation, they are convenient for addressing task-level imitation, which is a challenging problem in its own right given the limitations of robot sensors and effectors. Figure 6 shows a collection of our mobile robots.

Full-body humanoid robot: Our future plans include implementation and evaluation of the model on the Sarcos full-body humanoid robot, through collaboration with the Kawato Dynamic Brain Project at the ATR Human Information Processing Labs in Kyoto, Japan. The humanoid has 30 DOF: 3 in the neck, 7 in each arm, 3 in the trunk, 3 in each leg, and 2 in each eye.

The purpose behind using such a broad variety of test-beds is to validate the generality of our model, as well as put it to practical use on various complex agents and robots with potential real-world programming and control applications.

Experimental Validation of the Model

Our experiments to date have demonstrated several aspects of the model. As described earlier (Section X.6), we have experimentally validated and compared several types of movement primitives along different evaluation criteria (Matarić et al. 1999).

In Weber, Jenkins & Matarić (2000), we report on using a small set of three primitives (discrete straight and curved movements and oscillatory closed-curve movements) as a basis for recognition and reconstruction of various human upper-body movements, including dance (e.g., aerobics) and athletics (e.g., throwing a baseball) on the 20 DOF humanoid simulation. The particular primitives were chosen because they cover the space of possible observed end-point trajectories. To provide spatio-temporal invariance for the end-point data, we used a sliding horizon window of normalized

[4]Distributed by ActivMedia.

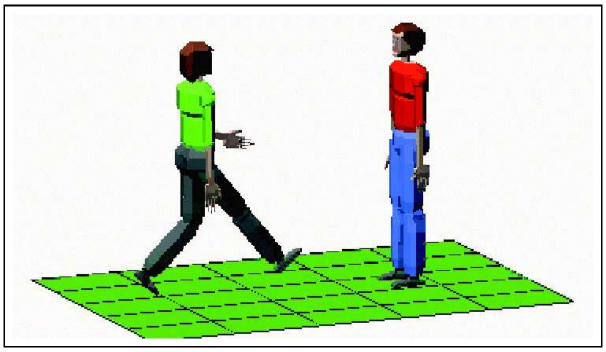

Figure 4: *Two cosimir avatars, 65 degree-of-fredom humanoid simulation test-beds used for validating our model.*

Figure 5: *Laika, a 20 degree-of-fredom robot dog used for validating our model.*

movement gradients rather than absolute end-point positions. Each instance of the window is a high-dimensional vector classified into one of the three primitives based on a Euclidean distance metric. Temporally consecutive windows are grouped based on classification results and produce a sequence of parametrized primitives; the parameters are extracted directly from the movement data. In the imitation phase, the hand of the humanoid executes the sequence of primitives using impedance control (Hogan 1985), in a similar fashion as demonstrated in our previous work (Matarić et al. 1999). The motor controller is provided with the 3D Cartesian end-point position and a default upper-arm orientation, resulting in a natural-looking trajectory reconstruction for all of the various types of movements. Figure 7 overviews the implementation. Details of the primitives, segmentation, and videos of the imitation can be found at http://www-robotics.usc.edu/~agents/imitation.html.

Besides demonstrating sequencing of primitives, we have experimented with superposing primitives as well. We combined an avoidance primitive with reaching and postural primitives in order

Figure 6: *A variety of mobile robots used for validating our model.*

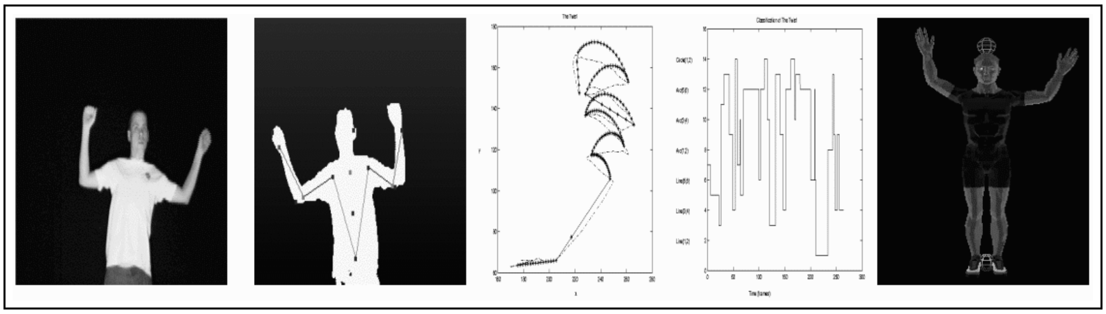

Figure 7: *An overview of the implementation of the imitation system.*

to produce collision-free hand trajectories around the torso and head of the Adonis simulation performing the Macarena dance (Matarić et al. 1998). Details and videos can be found at http://www-robotics.usc.edu/∼ agents/macarena.html.

We have also implemented a connectionist version of the imitation model, using central pattern generators as primitives for arm and leg movements, and a recurrent neural network architecture as the mechanism for learning novel behavior sequences (Billard & Matarić 2000*b*). The implementation was validated on the pair of COSIMIR avatars imitating various arm and finger movements from visual motion tracking (Billard & Matarić 2000*a*), as well as on the dog robot imitating a visually tracked wooden puppet (Billard & Matarić 2000*a*).

X.8 Imitation in Robotics

To relate our approach to others, we now provide a brief survey of related work on modeling imitation in physical robotic systems. We begin with an overview of various methods for visual perception and understanding of human movement, as it is a critical aspect of imitation as well as a major research challenge in machine vision.

Computational theories and implementations of movement perception have largely dealt with classifying movement by employing tools from time-series analysis. Hidden Markov Models (HMMs) (Rabiner & Juang 1986), originally used to automatically segment and recognize speech (Rabiner 1989), are the most popular modeling tool used for analyzing visual data of human and robot movement. Pook & Ballard (1993) used HMMs to analyze a video of a robot flipping a plastic fried egg. Brand, Oliver & Pentland (1997) used a coupled HMM for recognizing and analyzing time-extended behaviors such as object assembly. Lee & Xu (1996) used a similar HMM-based approach to recognize several hand-shaped letters. Yang, Xu & Chen (1997) applied the HMM representation to human skill learning by observation and skill transfer to articulated robot arms applied to recognition and copying of written digits and parts replacement by tele-operation. Finally, our own work has used a simplified form of HMMs to capture the dynamics of robot interactions over time as model of basis behavior execution patterns (Goldberg & Matarić 1999, Goldberg & Matarić 2000). The field of human-computer interfaces and its various relatives have addressed movement perception and greatly broadened the spectrum of approaches being studied. People-tracking and gesture recognition from video are two major applications of interest. A detailed survey of movement perception is beyond the scope of this chapter; for more information see Essa (1999).

The earliest robotics work to address imitation was focused on achieving assembly tasks from observation. Typically, a series of still visual images of a human performing a simple object moving/stacking tasks was recorded, segmented, interpreted, and then repeated by a robotic arm (Ikeuchi, Suehiro, Tanguy & Wheeler 1990, Ikeuchi, Kawade & Suehiro 1993, Kuniyoshi, Inaba & Inoue 1994, Hovland, Sikka & McCarragher 1996, Kaiser 1997). The work was aimed at task-level imitation (Byrne & Russon 1998), achieving the final outcome of the observed behavior, not at imitating the process that brought the result about. Motor control consisted of applying a conventional trajectory planner to reach each extracted sub-goal (e.g., pick up block1, put on top of block2). Recent efforts, including our own (Matarić 1994), have been increasingly oriented toward analyzing the underlying mechanisms of imitation in natural systems and modeling those on artificial ones. Schaal (1997) used imitation for "priming" a model-based reinforcement learning system in the task of balancing a pole by a 7 DOF robot arm observed from a brief human demonstration. Demiris, Rougeaux, Hayes, Berthouze & Kuniyoshi (1997) demonstrated a robot head with a pair of cameras that observed and imitated a small set of head movements of a human demonstrator, by using a built-in system that directly mapped a set of possible observed head movements to the robot's own movements. Using a mobile robot, Hayes & Demiris (1994) employed a similar direct-mapping imitation mechanism to learn maze navigation. The strategy of following a demonstrator robot was also used in Dautenhahn (1994) and Dautenhahn (1995), who investigated the social dimension of learning by imitation. Demiris (1999) describes a model related to ours, which also employs classification-biased recognition and prediction, as a part of a dual-route active-passive system. Haruno, Wolpert & Kawato (1999) and their related work (Haruno, Wolpert & Kawato 1998, Kawato 1995) describe a more general bi-directional approach to human motor control and learning, based on multiple pairs of forward (predictive) and inverse (controller) models, organized hierarchically. Experienced movement is segmented into predictive forward models, and the learning process is based on the error in the prediction. This is similar to our model, where segmentation is done at the level of primitives and known skills, which drive the classification and prediction. Unlike the above, we assume a smaller set of internal models, i.e., primitives.

X.9 Summary and Continuing Work

We have presented a model of imitation based on phylogenetically old mechanisms of mirror neurons and motor primitives. Combining evidence from neuroscience and cognitive science, we have postulated that imitation uses selective visual attention driven by end-point movement, a mirror-like matching and classification system for mapping the observed movement onto a set of motor primitives, a basis movement vocabulary consisting of those primitives, and a learning mechanism for creating novel sequences and superpositions of the primitives into a continually expanding movement repertoire. We then briefly described several experimental validations of our model and related it to other robotic imitation implementations.

We continue to develop and validate this imitation model on a variety of robotic test-beds, and hope to provide both novel insights into biological imitation and novel techniques for controlling and interacting with robots.

Acknowledgements

The work described here is supported by the National Science Foundation Career Award IRI-9624237. The author is grateful to the anonymous reviewers, to her imitation research group (Aude Billard, Chad Jenkins, and Ajo Fod), and especially to Richard Roberts, for invaluable comments on earlier drafts of this chapter.

X.9 Bibliography

References

Abernethy, B. (1990), 'Expertise, visual search, and information pick-up in squash', *Perception* **19**, 63–78.

Arbib, M. (1981), Perceptual Structures and Distributed Motor Control, *in* V. B. Brooks, ed., 'Handbook of Physiology: Motor Control', The MIT Press, pp. 809–813.

Arbib, M. (1992), Schema Theory, *in* S. Shapiro, ed., 'The Encyclophedia of Aritificial Intelligence, 2nd Edition', Wiley-Interscience, pp. 1427–1443.

Arbib, M. (2000), The Mirror System, Imitation, and the Evolution of Language, *in* K. Dautenhahn & C. Nehaniv, eds, 'Imitation in Natural and Artificial Systems', MIT Press.

Arbib, M. & Rizzolatti, G. (1996), 'Neural Expectations: A Possible Evolutionary Path From Manual Skills to Language', *Communication and Cognition* **29**(2–4), 393–424.

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.

Arkin, R. C. (1998), *Behavior-Based Robotics*, MIT Press, California.

Bandura, A. (1977), *Social Learning Theory*, Prentice-Hall, Inc., Englewood Cliffs, N.J.

Bandura, A. & Walters, R. H. (1963), *Social Learning and Personality Development*, Holt, Rinehart and Winston, Inc, New York.

Bertenthal, B. I. (1996), 'Origins and early development of perception, action, and representation', *Annual Review of Psychology* **47**, 431–459.

Billard, A. & Matarić, M. J. (2000*a*), Betty: Robot, play with me! Robot: O.K. How do we play? Betty: You watch me and do like I do. Look!, *in* 'Proceedings, Workshop on Interactive Robotics and Entertainment (WIRE-2000)', Pittsburgh.

Billard, A. & Matarić, M. J. (2000*b*), A biologically inspired robotic model for learning by imitation, *in* 'Proceedings, The Fourth International Conference on Autonomous Agents (Agents 2000)', Barcelona, Spain, pp. 372–380.

Bizzi, E., Mussa-Ivaldi, F. A. & Giszter, S. (1991), 'Computations Underlying the Execution of Movement: A Biological Perspective', *Science* **253**, 287–291.

Brand, M., Oliver, N. & Pentland, A. (1997), Coupled hidden Markov models for complex action recognition, *in* 'Proceedings, CVPR', IEEE Press, pp. 994–999.

Brooks, R. A. (1986), 'A Robust Layered Control System for a Mobile Robot', *IEEE Journal of Robotics and Automation* **RA-2**, 14–23.

Brooks, R. A. (1991), Intelligence Without Reason, *in* 'Proceedings, IJCAI-91', Morgan Kaufmann Publishers, Inc., Sydney, Australia, pp. 569–595.

Brooks, V. B. (1981), *Handbook of Physiology: Motor Control*, The MIT Press, Cambridge, Massachusetts.

Byrne, R. W. & Russon, A. E. (1998), 'Learning by Imitation: a Hierarchical Approach', *The Journal of Behavioral and Brain Sciences* **16**, 3.

Cheney, D. L. & Seyfarth, R. M. (1990), *How Monkeys See the World*, The University of Chicago Press, Chicago.

Dautenhahn, K. (1994), Trying to Imitate - a Step Towards Releasing Robots from Social Isolation, *in* 'Proceedings of Perception to Action', IEEE Computer Society Press, Lausanne, Switzerland, pp. 290–301.

Dautenhahn, K. (1995), 'Getting to Know Each Other - Artificial Social Intelligence for Autonomous Robots', *Robotics and Autonomous Systems* **16**, 333–356.

Davis, J. M. (1973), Imitation: A Review and Critique, *in* Bateson & Klopfer, eds, 'Perspectives in Ethology', Vol. 1, Plenum Press, pp. 43–72.

Decety, J. (1996), 'Do imagined and executed actions share the same neural substrate?', *Cognitive Brain Research* **3**, 87–93.

Demiris, J. (1999), Active and Passive Routes to Imitation, *in* K. Dautenhahn & C. Nehaniv, eds, 'Proceedings of the AISB'99 Sumposium on Imitation in Animals and Artifacts', Edinburgh.

Demiris, J. & Mataric, M. (1998), Perceptuo-motor primitives in imitation, *in* K. Dautenhahn & G. M. Hayes, eds, 'Proceedings of the AA98 Workshop on "Agents in Interaction: Acquiring Competence through Imitation"', Minneapolis/St. Paul, USA.

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.

di Pellegrino, G., Fadiga, L., Fogassi, L. & Rizzolatti, G. (1992), 'Understanding motor events: a neurophysiological study', *Experimental Brain Research* **91**, 176–180.

Donald, M. (1993), 'Precis of Origins of the modern mind: Three stages in the evolution of culture and cognition', *The Journal of Behavioral and Brain Sciences* **16**, 737–791.

Elkins, J. (1996), *The object stares back: On the nature of seeing*, Simon & Schuster, NY.

Essa, I. A. (1999), 'Computers Seeing People', *AI Magazine* **20**(2), 69–82.

Fadiga, L., Fogassi, L., Pavesi, G. & Rizzolatti, G. (1995), 'Motor facilitation during action observation: a magnetic stimulation study', *J. Neurophysiol.* **73**(6), 2608–2611.

Flanders, M. & Soechting, J. F. (1995), 'Frames of Reference for Hand Orientation', *Journal of Cognitive Neuroscience* **7**(2), 182–195.

Flanders, M., Tillery, S. I. H. & Soechting, J. F. (1992), 'Early stages in a sensorimotor transformation', *Behavioral and Brain Sciences* **15**, 309–362.

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.

Goldberg, D. & Matarić, M. J. (1999), Coordinating Mobile Robot Group Behavior Using a Model of Interaction Dynamics, *in* O. Etzioni, J. P. Muller & J. M. Bradshaw, eds, 'Proceedings, The Third International Conference on Autonomous Agents (Agents '99)', ACM Press, Seattle, Washington, pp. 100–107.

Goldberg, D. & Matarić, M. J. (2000), Reward Maximization in a Non-Stationary Mobile Robot Environment, *in* 'Proceedings, The Fourth International Conference on Autonomous Agents (Agents 2000)', Barcelona, Spain.

Goodale, M. A., Jakobson, L. S., Milner, A. D., Perrett, D. I., Benson, P. J. & Hietanen, J. K. (1994), 'The Nature and Limits of Orientation and Pattern Processing Supporting Visuomotor Control in a Visual Form Agnosic', *Journal of Cognitive Neuroscience* **6**(1), 46–56.

Harnad, S. (1990), 'The Symbol Grounding Problem', *Physica D* **42**, 335–346.

Haruno, M., Wolpert, D. M. & Kawato, M. (1998), Multiple Paired Forward-Inverse Models in the Cerebellum, *in* 'Proceedings of the 5th International Conference on Neural Information Processing (ICONIP '98)', Vol. 13, pp. 1177–1180.

Haruno, M., Wolpert, D. M. & Kawato, M. (1999), Multiple Forward-inverse models for human motor learning and control, *in* 'Advances in Neural Information Processing Systems 11', MIT Press.

Hayes, G. & Demiris, J. (1994), A Robot Controller Using Learning by Imitation, *in* A. Borkowski & J. L. Crowley, eds, 'Proceedings of the International Symposium on Intelligent Robotic Systems', LIFIA–IMAG, Grenoble, France, pp. 198–204.

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

Hovland, G. E., Sikka, P. & McCarragher, B. J. (1996), Skill Acquisition from Human Demonstration Using a Hidden Markov Model, *in* 'Proceedings, IEEE International Conference on Robotics and Automation', Minneapolis, MN, pp. 2706–2711.

Hubel, D. & Wiesel, T. (1977), 'Functional Architecture of Macaque Monkey Cortex', *Proceedings of the Royal Society of London* **198**, 1–59.

Iacoboni, M., Woods, R. P., Brass, M., Bekkering, H., Mazziotta, J. C. & Rizzolatti, G. (1999), 'Cortical Mechanisms of Human Imitation', *Science* **286**, 2526–2528.

Ikeuchi, K., Kawade, M. & Suehiro, T. (1993), Assembly Task Recognition with Planar, Curved, and Mechanical Contacts, *in* 'Proceedings of IEEE International Conference on Robotics and Automation', Atlanta, GA.

Ikeuchi, K., Suehiro, T., Tanguy, P. & Wheeler, M. (1990), Assembly Plan from Observation, Technical report, Carnegie Mellon Univeristy Robotics Institute Annual Research Review.

Jeannerod, M. & Decety, J. (1995), 'Mental motor imagery: a window into the representational stages of action', *Current Opinion in Neurobiology* **5**, 727–732.

Jeannerod, M., Arbib, M. A., Rizzolatti, G. & Sakata, H. (1995), 'Grasping objects: the cortical mechanisms of visuomotor transformation', *Trends in Neuroscience* **18**, 314–320.

Jenkins, O. C. (2000), An Attention Selection Mechanism for Moving Human Figures, Technical Report IRIS-00-381, University of Southern California, Institute for Robotics and Intelligent Systems.

Johansson, G. (1973), 'Visual Perception of Biological Motion and a Model for its Analysis', *Perception and Psychophysics* **14**, 201–211.

Kaelbling, L. P., Littman, M. L. & Moore, A. W. (1996), 'Reinforcement Learning: A Survey', *Journal of Artificial Intelligence Research* **4**, 237–285.

Kaiser, M. (1997), 'Transfer of Elementary Skills via Human-Robot Interaction', *Adaptive Behavior*.

Kawato, M. (1995), Unidirectional versus bi-directional theory for trajectory planning and control, *in* T. Hida, ed., 'Mathematical Approach To Fluctuations: Complexity and Nonlinearity', World Scientific Ltd., Singapore, pp. 144–180.

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.

Lee, C. & Xu, Y. (1996), Online, Interactive Learning of Gestures for Human/Robot Interfaces, *in* 'Proceedings, IEEE International Conference on Robotics and Automation', Vol. 4, Minneapolis, MN, pp. 2982–2987.

Marjanović, 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. (1994), Learning Motor Skills by Imitation, *in* 'Proceedings, AAAI Spring Symposium Toward Physical Interaction and Manipulation', Stanford University.

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. & Marjanović, M. J. (1993), Synthesizing Complex Behaviors by Composing Simple Primitives, *in* 'Self Organization and Life: From Simple Rules to Global Complexity, European Conference on Artificial Life (ECAL-93)', Brussels, Belgium, pp. 698–707.

Matarić, M. J. & Pomplun, M. (1998), 'Fixation Behavior in Observation and Imitation of Human Movement', *Cognitive Brain Research* **7**(2), 191–202.

Matarić, M. J., Zordan, V. B. & Mason, Z. (1998), Movement Control Methods for Complex, Dynamically Simulated Agents: Adonis Dances the Macarena, *in* 'Autonomous Agents', ACM Press, Minneapolis, St. Paul, MI, pp. 317–324.

Matarić, M. J., Zordan, V. B. & Williamson, M. (1999), 'Making Complex Articulated Agents Dance: an analysis of control methods drawn from robotics, animation, and biology', *Autonomous Agents and Multi-Agent Systems* **2**(1), 23–44.

Meltzoff, A. N. & Moore, M. K. (1988), 'Imitation of Facial and Manual Gestures by Human Neonates', *Science* **198**, 75–78.

Meltzoff, A. N. & Moore, M. K. (1994), 'Imitation, Memory, and the Representation of Persons', *Infant Behavior and Development* **17**, 83–99.

Meltzoff, A. N. & Moore, M. K. (1995), Infants' Understading of People and Things: From Body Imitation to Folk Psychology, *in* J. L. Bermudez, A. Marcel & N. Eilan, eds, 'The Body and the Self', MIT Press/Bradford Books, pp. 44–69.

Meltzoff, A. N. & Moore, M. K. (1997), 'Explaining Facial Imitation: A Theoretical Model', *Early Development and Parenting* **6**(2), 157.1–14.

Moore, B. R. (1992), 'Avian Movement Imitation and a new Form of Mimicry: Tracing the Evolution of a Complex Form of Learning', *Behavior* **122**, 614–623.

Moore, B. R. (1996), The Evolution of Imitative Learning, *in* C. M. Heyes & B. G. Galef, eds, 'Social Learning in Animals: The Roots of Culture', Academic Press, New York, pp. 245–265.

Mussa-Ivaldi, F. A. & Giszter, S. (1992), 'Vector field approximation: a computational paradigm for motor control and learning', *Biological Cybernetics* **67**, 491–500.

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.

Parr, L. A. & de Waal, F. B. M. (1999), 'Visual kin recognition in chimpanzees', *Nature* **399**, 647–647.

Pashler, H. E. (1999), *The Psychology of Attention*, The MIT Press, Cambridge, MA.

Pelisson, D., Goodale, M. A. & Prablanc, C. (1978), Adjustments of Hand Pointings to Visual Targets do not Need Visual Reafference From the Moving Limb, *in* J. K. O'Regan & A. Levy-Schoen, eds, 'Eye Movements, From Physiology to Cognition: Selected/Edited Proceedings of the Third European Conference on Eye Movements', pp. 115—121.

Perrett, D. I., Harries, M. H., Bevan, R., Thomas, S., Benson, P. J., Mistlin, A. J., Chitty, A. J., Hietanen, J. K. & Ortega, J. E. (1989*a*), 'Frameworks of Analysis for the Neural Representation of Animate Objects and Actions', *Journal of Experimental Biology* **146**, 87–113.

Perrett, D. I., Harries, M. H., Mistlin, A. J. & Chitty, A. J. (1989*b*), Recognition of objects and actions: frameworks for neuronal computation and perceptual experience, *in* O. Guthrie, ed., 'Higher Order Sensory Processing', Manchester University Press.

Perrett, D. I., Harries, M. H., Mistlin, A. J., Hietanen, J. K., Benson, P. J., Bevan, R., Thomas, S., Oram, M. W., Ortega, J. & Brierley, K. (1990), 'Social Signals Analyzed at the Single Cell Level: Someone is Looking at Me, Something Touched Me, Something Moved!', *International Journal of Comparative Psychology* **4**(1), 25–55.

Perrett, D. I., Smith, P. A. J., Mistlin, A. J., Chitty, A. J., Head, A. S., Potter, D. D., Broennimann, R., Milner, A. D. & Jeeves, M. A. (1985), 'Visual Analysis of Body Movements by Neurones in the Temporal Cortex of the Macaque Monkey: A preliminary Report', *Behavioral Brain Research* **16**, 153–170.

Piaget, J. (1962), *Play, Dreams and Imitation in Children*, W. W. Norton & Co., New York.

Pook, P. K. & Ballard, D. H. (1993), Recognizing teleoperated manipulations, *in* 'Proceedings of IEEE International Conference on Robotics and Automation', Vol. 2, Atlanta, Georgia, pp. 578–585.

Rabiner, L. R. (1989), 'A tutorial on hidden Markov models and selected applications in speech recognition', *Proceedings of the IEEE* **77**(2), 257–286.

Rabiner, L. R. & Juang, B. H. (1986), 'Introduction to Hidden Markov Models', *IEEE ASSP Magazine* pp. 4–16.

Rizzolatti, G., Fadiga, L., Matelli, M., Bettinardi, V., Perani, D. & Fazio, F. (1996*a*), 'Localization of Grasph Representations in Humans by Positron Emission Tomography: 1. Observation Versus Execution', *Experimental Brain Research* **111**, 246–252.

Rizzolatti, G., Gadiga, L., Gallese, V. & Fogassi, L. (1996*b*), 'Premotor cortex and the recognition of motor actions', *Cognitive Brain Research* **3**, 131–141.

Romann, J., F. E. (1999), 'Projective Virtual Reality: Bridging the Gap between Virtual Reality and Robotics', *IEEE Transaction on Robotics and Automation; Special Section on Virtual Reality in Robotics and Automation* **15**(3), 411–422.

Sanger, T. D. (2000), 'Human Arm Movements Described by a Low-Dimensional Superposition of Principal Components', *Journal of Neuroscience* **20**(3), 1066–1072.

Savelsvertgh, G. J. P. & Bootsma, R. J. (1994), 'Perception-action coupling in hitting and catching', *International Journal of Sport Psychology* **25**, 331–343.

Schaal, S. (1997), Learning from demonstration, *in* M. Mozer, M. Jordan & T. Petsche, eds, 'Advances in Neural Information Processing Systems 9', The MIT Press, pp. 1040–1046.

Schaal, S. & Sternad, D. (1998), Programmable Pattern Generators, *in* 'Proceedings, 3rd International Conferece on Computational Intelligence in Neuroscience', Research Triangle Park, NC, pp. 48–51.

Stein, J. F. (1992), 'The representation of egocentric space in the posterior parietal cortex', *Behavior and Brain Science* **15**, 691–700.

Stein, P. S. G. (1997), *Neurons, Networks, and Motor Behavior*, The MIT Press, Cambridge, Massachusetts.

Sutton, R. S. & Barto, A. G. (1998), *Reinforcement Learning*, MIT Press, Cambridge, Massachusetts.

Tomasello, M., Kruger, A. C. & Rather, H. H. (1993), 'Cultural Learning', *The Journal of Behavioral and Brain Sciences* **16**(3), 495–552.

Vogt, S. (1995), 'Imagery and perception-action mediation in imitative actions', *Cognitive Brain Research* **3**, 79–86.

Weber, S. (2000), Simple Human Torso Tracking from Video, Technical Report IRIS-00-380, University of Southern California, Institute for Robotics and Intelligent Systems.

Weber, S., Jenkins, O. C. & Matarić, M. J. (2000), Imitation Using Perceptual and Motor Primitives, *in* 'Proceedings, The Fourth International Conference on Autonomous Agents (Agents 2000)', Barcelona, Spain.

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.

Yang, J., Xu, Y. & Chen, C. S. (1997), 'Human Action Learning via Hidden Markov Model', *IEEE Transactions on Systems, Man, and Cybernetics–Part A: Systems and Humans* **27**(1), 34–44.