

The New Robotics—towards human-centered machines

Stefan Schaal¹

¹Computer Science and Neuroscience, University of Southern California, 3710 S. McClintock Avenue—RTH 401, Los Angeles, California 90089-2905 and ATR Computational Neuroscience Laboratories, 2-2 Hikoridai, Seika-cho, Soraku-gun, Kyoto 619-02

(Received 26 December 2006; accepted 21 May 2007; published online 16 July 2007)

Research in robotics has moved away from its primary focus on industrial applications. The New Robotics is a vision that has been developed in past years by our own university and many other national and international research institutions and addresses how increasingly more human-like robots can live among us and take over tasks where our current society has shortcomings. Elder care, physical therapy, child education, search and rescue, and general assistance in daily life situations are some of the examples that will benefit from the New Robotics in the near future. With these goals in mind, research for the New Robotics has to embrace a broad interdisciplinary approach, ranging from traditional mathematical issues of robotics to novel issues in psychology, neuroscience, and ethics. This paper outlines some of the important research problems that will need to be resolved to make the New Robotics a reality. [DOI: 10.2976/1.2748612]

CORRESPONDENCE

Stefan Schaal: sschaal@usc.edu

In 1921, Karel Capek's play *Rossum's Universal Robots*—largely credited to be the beginning of the era of robotics [e.g., [Sciavicco and Siciliano \(1996\)](#)—provided two interesting visions. First, robots were human-like, and second, they interacted with humans in a natural, i.e., human-like, way. For more than half a century after Capek's conception of a robot, realizations lagged far behind his ideas: robots were minimalist mechanical devices, far from achieving any general human-like performance, and there was essentially no interaction between humans and robots, except for a programmer uploading code. Of course, all this made sense, as robots were merely used for highly repetitive and high precision positioning tasks in industry. And it is in this domain that, so far, robots have made a difference and contributed in useful ways to the advance of technology and industrial productivity.

The beginning of the 21st century, however, witnessed a remarkable

change in the research focus of robotics. As initiated by artificial intelligence in the 1980's and 1990's, there has always been great interest in generating more general purpose and intelligent robotic machines. As there was little application for such robots in industry—i.e., industrial special purpose robots are hard to beat in their productivity on the task they were designed for—work on intelligent robots was initially largely confined to university research and some unique research programs like space exploration, underwater exploration, etc. With the advance of information technology and mechanical design, however, Capek's vision of human-like, autonomous, and interactive robots has gradually become within the reach of research prototypes. In the late 1990's, a new field of robotics became salient, humanoid robotics, primarily through pioneering work at Waseda University, Sarcos Inc., and the Honda Corporation [e.g., for an overview, see [Menzel and D'Alusio](#)

(2000)]. In these places, the first full-body humanoid robots were built and introduced to the research community and popular press. Waseda University's Wabian series of human robots primarily focused on biped walking skills, as did Honda's Asmio robot, which, due to its association with a very large company and its very refined appearance, has received a lot of attention all over the world. The Sarcos humanoid robot "DB" focused on autonomous manipulation skills, often paired with learning abilities that were derived from ideas from computational neuroscience (Atkeson *et al.*, 2000). By now, numerous humanoid robots exist, primarily in Asia and Japan in particular, but also in Europe and the US. Since the year 2000, there is even a specialized conference devoted to humanoid robotics, the IEEE/RAS International Conference on Humanoid Robotics.

Besides the long-standing scientific challenge of understanding how the human brain works and how to create an artificial system with similar capabilities as found in humans, is there any real use to having humanoid robots? From the viewpoint of traditional manufacturing in industry, the answer would probably be negative, as the cost/benefit ratio of humanoid robots is—so far and in the medium future—quite unattractive. But traditional industry is most likely not what the New Robotics is going to impact on. Rather, humanoid and similar robots will be part of our normal life, i.e., they will take over functions in our society where there are simply too few humans to cope with societal needs. This issue has been recognized and put forward by many countries in the world, often sparked by the trends of increasing proportions of older people in the society and the associated danger of a future shortage of adequate care giving. However, there are many other areas where robots can be helpful. Work in hazardous environments, and search and rescue missions after disasters are among the most commonly mentioned issues. Robots can also act as physical therapists to help patients exercise properly and regain lost motor functions—for instance, the lack of physical therapists provides stroke patients in the US with only about 40 min of training a day, while about 6 h a day would be needed in the immediate poststroke period to regain motor functions (Taub and Wolf, 1997; van der Lee *et al.*, 1999; Wolf and Blanton, 1999)—wouldn't it be nice if one could "check out" a robot physical therapist from the local pharmacy? Of course, robots could help in reaching, carrying, and manipulating objects for people in convalescence and with disabilities—this topic is often discussed under the name of personal robotics or assistive robotics. These robots can also take over mentoring and cognitive functions, like reminding people to take their medications in time, or to not forget other scheduled activities. There are also interesting opportunities for robots in the education and the daily life of children. For instance, autistic children seem to react very positively to educational games with robots (e.g., Robins *et al.*, 2005). And robots could also become playmates and educational helpers in




Figure 1. An artist's sketch of future involvement of robots in the society: (a) assistive robots, (b) playmate robots in child education, (c) robots for mentoring and assistance in manipulation tasks, (d) robots that teach movement exercises, (e) personal robots for the elderly, (f) robots for surveillance and protection of children and adults.

school scenarios such that children become and remain interested in higher education in science and technology. Many more scenarios could be enumerated here—Fig. 1 provides an artist's sketch of some of these ideas.

It is important to emphasize that the New Robotics is going to be human-centered¹, i.e., robots will work with humans in normal human environments. And it is in this role of assistive robots that many new and interdisciplinary components become part of robotics research. For instance, if a robot is supposed to be accepted by children and adults, it needs to adhere to certain social behaviors and standards that we as humans find acceptable. Thus, the psychology and ethics of human-robot interaction needs to be explored. If a robot is going to act as a physical therapist, the robot needs to incorporate knowledge about the neural and musculo-skeletal deficiencies of a patient, and tailor exercises and suggestions for improvements accordingly. Essentially, such a robot must have some knowledge of how the human brain controls movement, and how neural deficiencies, e.g., due to stroke, affect the movement abilities of a patient. Or, if a robot is supposed to work in a human environment with the same objects and tools that humans use in daily life, it may be beneficial to understand how humans accomplish motor skills and how the human skill repertoire is structured. Thus,

¹There are also some dedicated workshops for human-centered robotics that have been started recently, e.g., the International Workshop on Human-Centered Robotics.

roboticists, traditionally deeply grounded in applied mathematics, mechanics, control theory, and optimization theory, are suddenly confronted with a wealth of new topics ranging from psychology, biology, and neuroscience to ethics. The New Robotics will require a new kind of scientist that can traverse a very broad range of different disciplines.

Figure 2 is an attempt to structure research topics in this new, human-centered robotics into three major categories: (i) action, (ii) interaction, and (iii) engagement. Action research is primarily concerned with issues of motor control in both humans and robots. Interaction research emphasizes human-robot or robot-robot interaction, i.e., the basic research of how two systems can act together in order to accomplish a goal—at one extreme, this means that one system, e.g., the human, instructs the other system, the robot, what to do. At the other extreme, the two systems could work together to accomplish a goal, e.g., as in jointly carrying a heavy object. Finally, engagement research examines the cognitive and psychological components of human-centered robotics. Naturally, the structure suggested in Fig. 2 is not crisp, i.e., there is a large amount of overlap between action, interaction, and engagement research. Nor is the list of research topics in Fig. 2 comprehensive. But for the purpose of giving a perspective on the future of robotics research and in order to review some relevant recent literature, Fig. 2 provides a useful guideline for the coming sections of this paper. It should be noted that several other perspectives of the future of robotics can be found in the literature (e.g., [Khatib et al., 2004](#); [Brock et al., 2005](#)).

ACTION RESEARCH

Action research is grounded in theory of motor control, which was largely laid out in classical engineering fields like cybernetics ([Wiener, 1948](#)), optimal control ([Bellman, 1957](#)), and control theory ([Slotine and Li, 1991](#); [Narendra and Annaswamy, 2005](#)). These fields addressed many crucial issues of movement systems, including negative feedback

control, feedforward control, nonlinear control, movement planning with optimization criteria, stability, control with time delayed system, adaptive control, stochastic control, state estimation, control of overactuated and underactuated systems, etc. Figure 3 illustrates a classical control diagram for motor control that highlights the different stages of a control circuit, each of which is associated with a large number of research topics. Many of these topics have received ample attention in robotics research such that they have become textbook knowledge (e.g., [Russell and Norvig, 1995](#); [Sciavicco and Siciliano, 2000](#)). Several issues, however, deserve special attention. It should be noted that a large amount of research in biological motor control is focused on topics similar to those in Fig. 3 ([Sabes, 2000](#); [Schaal and Schweighofer, 2005](#); [Shadmehr and Wise, 2005](#)).

MOTOR PRIMITIVES AND MOTOR LIBRARIES

While issues like coordinate transformations and motor command generation in robotics (Fig. 3) are fairly well understood by now (e.g., [Sciavicco and Siciliano, 2000](#)), they rely on someone providing the robot with an appropriate plan of what to do. In many robotics projects, this part is taken over by a human designer who manually creates the plan for a particular behavior. While this approach is highly successful in static environments, the New Robotics is inherently about robots acting in normal human, i.e., dynamic environments. Thus, the question arises of how to create flexible and reactive motor control and planning. Most commonly, an answer is sought by creating a toolbox of elementary movements, which in sequence or superposition create a sufficiently large spectrum of complex motor skills. We call such elementary movements motor primitives, but in the literature, they have also been denoted as schemas, basis behaviors, options, macros, etc. (for a review, see [Mataric, 1998](#); [Schaal, 1999](#); [Schaal et al., 2003](#)).

One can distinguish at least two major trends on research on movement primitives. First, movement primitives can be

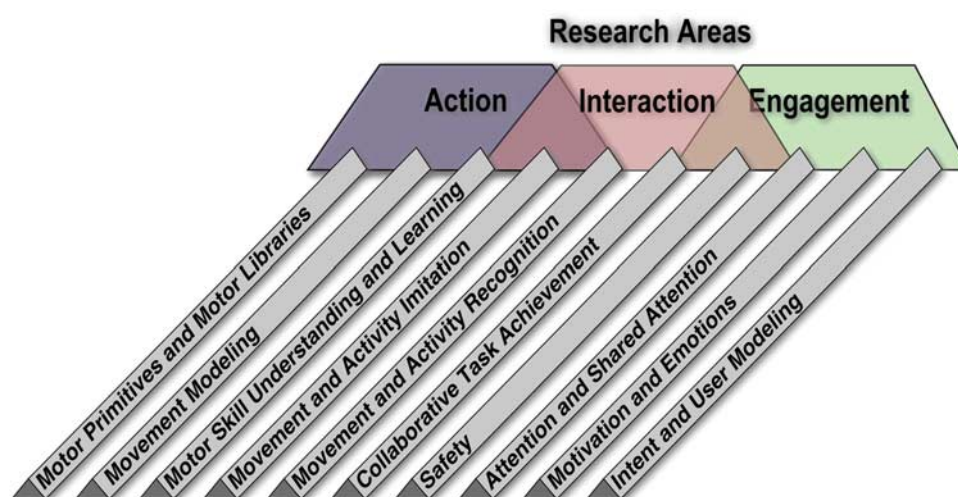


Figure 2. Research topics in human-centered robotics.

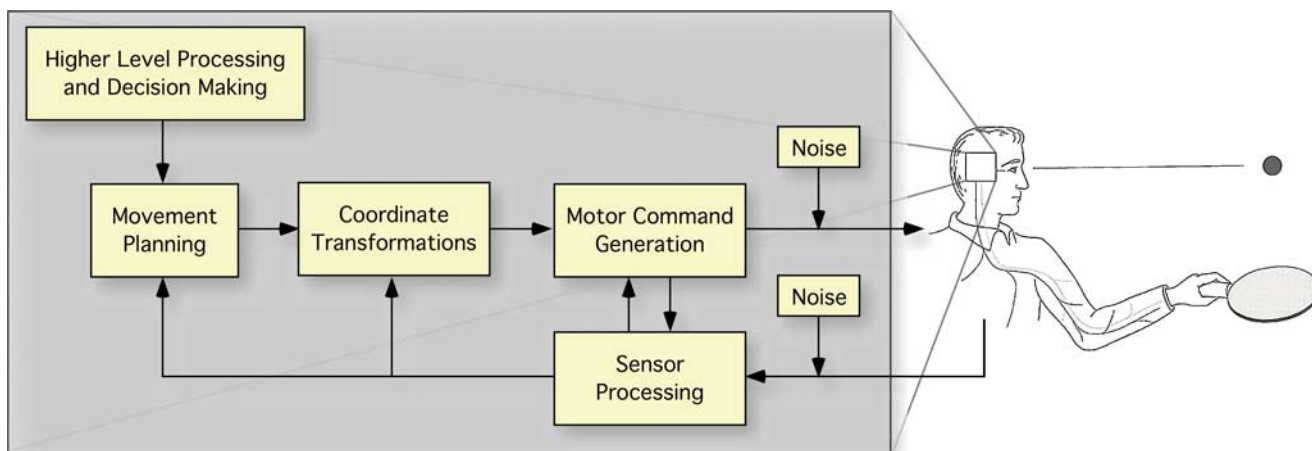


Figure 3. Sketch of a generic control diagram that is typically used in robotics research.

modeled as dynamic systems, like point attractors and limit cycles (e.g., Bullock and Grossberg, 1988; Burridge *et al.*, 1999; Schaal *et al.*, 2004; Righetti and Ijspeert, 2006). The dynamic systems approach allows addressing various important issues in human-like movement in a natural way, e.g., stability issues, perception action coupling, obstacle avoidance with potential fields, etc. Frequently, dynamic systems motor primitives are easily interpreted as behavioral units, e.g., reaching for a target, biped locomotion, etc. As a drawback, modeling with nonlinear differential equations is often quite complex and not always analytically tractable.

The second major approach to movement primitives is by means of statistical techniques (e.g., Atkeson *et al.*, 2000; Haruno *et al.*, 2001; Doya *et al.*, 2002; Jenkins, 2003; Inamura *et al.*, 2004). As an advantage, statistical methods offer well-understood tools and mechanisms to extract movement modules from empirical data, such that movement primitives can be grounded in behavioral recordings from humans. As a disadvantage, it seems that statistically extracted movement primitives are frequently not behaviorally meaningful, such that it is not clear how well such movement primitives generalize to new behaviors. Moreover, phenomenal like perception action coupling, obstacle avoidance, etc., are not naturally included in statistical models of movement primitives.

So far, comprehensive implementations of any movement primitive approach are lacking, and it is hard to objectively favor one approach over the other.

MOVEMENT MODELING

In order to operate in environments that were built for humans, it is often advantageous for robots to employ similar movement strategies as humans. Additionally, if robots move like humans, humans accept them more easily as partners due to our inherent tuning to natural looking movement (Johansson, 1973; Viviani and Schneider, 1991; Viviani *et al.*, 1997). Thus, one needs to ask the question of what principles (or models) humans use in movement generation, and

what are the underlying strategies of human motor skills. These questions are hard, as they have been addressed by the field of behavioral psychology for motor control and neuro-motor control for several decades (e.g., see Schmidt, 1988; Arbib, 1995a; Arbib, 1995b; Shadmehr and Wise, 2005) and definite conclusions are still missing.

However, it is possible to highlight several characteristics of human movement, in particular, if these characteristics are largely missing in robotic systems. One of these is that human movement generation is highly compliant, while most (even modern) robots use rather stiff (high gain) control. The reason why roboticists favor high gain control is that it reduces the need for accurate models of the robot and its environment. High gain control simply means that position and velocity errors between a desired and an achieved trajectory are multiplied with a large number, and these results are then added to the motor commands. Thus, small deviations from the desired movement lead to strong corrective motor commands, i.e., the robot feels very stiff. The problem with stiff control is that any unforeseen perturbation is fought with very high effort, potentially so high that either the robot itself or the environment may get damaged. Humans and animals, on the other hand, softly give in to perturbation, and thus avoid hurting the environment (which may be another human) or themselves. Given that in a dynamic environment collisions with the environment are unavoidable, compliant control seems to be crucial for the New Robotics.

An alternative to high gain control is low gain control with feedforward models (An *et al.*, 1988; Miall and Wolpert, 1996; Kawato, 1999), i.e., models of the dynamics of the robot that support predicting the optimal motor command for any position, velocity, and acceleration of the robot's degrees of freedom. In theory, the design specifications of a robot allow generating these models fairly accurately. But in practice, theoretical models and actual dynamics of the robot often do not match, due to complex actuator dynamics, wires, hoses, protective shells, etc., on the robot. Thus, low gain

model-based control on complex human-like robots is actually quite hard and still a topic of ongoing research (Vijayakumar and Schaal, 2000; Vijayakumar *et al.*, 2005). The topic becomes even more complex if the interaction dynamics between a robot and its environment are taken into account—every walking system, for instance, has these interaction dynamics due to ground-reaction forces. Interactions with the environment create constraints, and the mathematics of model-based control becomes significantly more complex (Peters *et al.*, 2005a; Nakanishi *et al.* submitted). As a consequence, compliant control in human-like robots has hardly been achieved, neither theoretically nor experimentally (for some exceptions, see Albu-Schäffer *et al.*, 2007), and, potentially, a lot can be learned from human movement studies (Gomi and Kawato, 1996; Franklin *et al.*, 2003; Osu *et al.*, 2004) and the properties of the human musculo-skeletal system (Brown *et al.*, 1999; Brown and Loeb, 1999; Loeb, 2001).

Another fascinating feature of human motor control is that most of the time, we accomplish several tasks in parallel, e.g., balancing on two feet and grasping for an object. Robotics theory has developed methods of hierarchical task prioritization (Nakamura, 1991; Khatib *et al.*, 2004; Sentis and Khatib, 2004; Arimoto *et al.*, 2005; Peters *et al.*, 2005b) that, in theory, are quite appealing and capable. In recent work, Sentis and Khatib (Khatib *et al.*, 2002; Khatib *et al.*, 2004; Sentis and Khatib, 2004) demonstrated that rather natural looking behaviors could be generated with appropriate task-space controllers and hierarchies. However, most of the methods rely on accurate models of the robot, and become quite complex when the robot is in contact with the environment, particularly when the contact conditions switch (like in locomotion or object manipulation). Rather little experimental evaluations exist of hierarchical task control on complex robotic systems, and behavioral studies with humans have hardly addressed such complicated problems.

MOTOR SKILL UNDERSTANDING AND LEARNING

Much more needs to be understood of how to represent motor skill. There is usually a choice as to whether special task coordinates, general external coordinates (e.g., Cartesian coordinates), or internal (joint-angle) coordinates are the most suitable (Saltzman, 1979; Schaal *et al.*, 1992). Special control strategies can make motor skills easier or harder (e.g., Beek, 1989; Schaal *et al.*, 1992). Some ideas exist how to generate generic representations of motor skill (Wada and Kawato, 1995; Miyamoto *et al.*, 1996; Miyamoto and Kawato, 1998; Ijspeert *et al.*, 2003), but, so far, a generic approach to skill acquisition is missing.

Another important component in understanding motor skills lies in representing the task goal, i.e., some kind of a metric that measures task achievement. For some tasks, like reaching for an object, this is straightforward as the metric of success is simply the distance of the hand from the object.

For other tasks, for instance, bipedal locomotion, the task goal is much more abstract and not easily quantifiable. For biped locomotion, the primary task goal is not to fall over, but a secondary task goal is to conserve energy, a tertiary task goal is to reduce wear and tear on the motor system, and another task goal could be to maintain balance while simultaneously reaching for a coffee cup. The goal to balance itself is not easily quantifiable: we could choose the desired position of the center of gravity, but also a desired position of the zero moment point (Kajita and Tani, 1996). Other researchers may claim that such measures are unsuitable for human-like biped locomotion and resort to describing the biped locomotion in terms of a stability measure of the associated limit cycle (McGeer, 1990; Collins *et al.*, 2005). Essentially, specifying task specific optimization criteria is not easy.

If we assume we have a task representation and a metric to measure task success, it is still not clear how to choose appropriate motor commands to accomplish the task. Insights from a human designer are one possibility, but as mentioned before, the New Robotics will require robots that can react to dynamic environments, and it is unlikely that human designers can foresee all possible events that a robot will be exposed to. Thus, learning robots would be very desirable, a field that is called learning control. It is helpful to follow one of the most general frameworks of learning control as originally developed in the middle of the 20th century in the fields of optimization theory, optimal control, and in particular dynamic programming (Bellman, 1957; Dyer and McReynolds, 1970). Here, the goal of learning control was formalized as the need to acquire a task dependent control policy \bullet that maps the continuous valued state vector \mathbf{x} of a control system and its environment, possibly in a time dependent way, to a continuous valued control vector \mathbf{u} :

$$\mathbf{u} = \pi(\mathbf{x}, \alpha, t). \quad (1)$$

The parameter vector α contains the problem specific parameters in the policy \bullet that need to be adjusted by the learning system. Since the controlled system can generally be expressed as a nonlinear function

$$\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u}) \quad (2)$$

in accordance with standard dynamical systems theory (Strogatz, 1994), the combined system and controller dynamics result in

$$\dot{\mathbf{x}} = f(\mathbf{x}, \pi(\mathbf{x}, t, \alpha)). \quad (3)$$

Thus, learning control means finding a (usually nonlinear) function \bullet that is adequate for a given desired behavior and movement system.

A general approach of learning control policies is reinforcement learning (Sutton and Barto, 1998; Schaal, 2002; Barto and Mahadevan, 2003). Unfortunately, reinforcement learning is a rather slow technique, with tremendous computational and theoretical problems when it comes to applying

it to actual robots (Peters *et al.*, 2003a; Peters *et al.*, 2003b; Peters *et al.*, 2005c). Alternatively, one can choose to apply learning only to some parts of the control problem, e.g., as in learning the internal models that are needed in feedforward control. Such tasks can be achieved with supervised learning techniques, which is computationally and practically much more tractable (Atkeson *et al.*, 1986; Schaal and Atkeson, 1994; Atkeson *et al.*, 1997a; Atkeson *et al.*, 1997b; Schaal *et al.*, 2002; Peters and Schaal, 2006). Unfortunately, learning applied to human-like robots has been pursued only to a rather small extent, and a large amount of research both on the theoretical and practical side of machine learning and robotics remains to be done (for a review, see Schaal, 2002; Schaal, *in press*).

INTERACTION RESEARCH

Interaction research adds an interesting component to robotics, i.e., how to coordinate two movement systems to cooperate on the same task goal. Of course, this cooperation is not constrained to a mechanical cooperation but allows for any kind of interaction, e.g., language, gestures, etc. Interaction research has recently started its own research field and associated conferences: human-robot interaction (HRI). To illustrate the complexity of HRI, we could tentatively define it similar to the term “human-computer interaction:”

HRI is where both humans and robots are engaged in a mutual affecting response experience. In technical terms, we have (at least) two-coupled control policies, one working out of a human, one out of a robot

The problem here is that HRI couples at least two nonlinear dynamic systems. One of them, the robot, may be rather well understood. The other, the human, is less well understood, and even at the risk of being hurt if HRI goes wrong. In general, it is well understood that coupled nonlinear dynamic systems can behave in a predictable or unpredictable, highly complex manner (Strogatz, 1994). Thus, research has to be rather careful in designing methods for HRI, always keeping in mind that simple algorithms that are well understood in isolation can cause damage when coupled over multiple systems. That is, ethics and safety concerns are highly important in HRI, as well as a solid understanding of the dynamics of human action and perception, and which stimuli or behaviors could trigger dangerous situations for the human, e.g., seizures from perceptual stimulation, emotional reactions, or inappropriate physical reactions.

MOVEMENT AND ACTIVITY IMITATION

Cultural and imitation learning are among the most fascinating elements of humans, and often believed to be a cornerstone for how humans could develop such a high level of intelligence (cf. reviews in Piaget, 1951; Tomasello *et al.*, 1993; Meltzoff and Moore, 1994; Byrne and Russon, 1998; Rizzolatti and Arbib, 1998; Dautenhahn and Nehaniv, 2002). For more than 20 years, roboticists have tried to add imita-

tion learning, or learning from demonstration as it was called initially, to the repertoire of robot learning. Movement imitation is familiar to everybody from daily experience: a teacher demonstrates a movement, and immediately the student is capable of approximately repeating it. From the viewpoint of learning, a teacher’s demonstration as the starting point of one’s own learning can significantly speed up the learning process, as imitation usually drastically reduces the amount of trial and error that is needed to accomplish the movement goal by providing a good example of a successful movement (Schaal, 1999). Thus, from a robotics point of view, it is important to understand the detailed principles, algorithms, and metrics that subserves imitation, starting from the visual perception of the teacher up to issuing motor commands that move the limbs of the student.

Figure 4 sketches the major ingredients of a conceptual imitation learning system [for a review of many related papers, see Schaal (1999)]. Visual sensory information needs to be parsed into information about objects and their spatial location in an internal or external coordinate system; the depicted organization is largely inspired by the dorsal (what) and ventral (where) stream as discovered in neuroscientific research (van Essen and Maunsell, 1983). As a result, some form of postural information of the movement of the teacher and/or three-dimensional object information about the manipulated object (if an object is involved) should become available. Subsequently, one of the major questions revolves around how such information can be converted into action. For this purpose, Fig. 4 alludes to the concept of movement primitives, which was already discussed above. In general, we think of movement primitives to code complete temporal behaviors, like “grasping a cup,” “walking,” “a tennis serve,” etc. Figure 4 assumes that the perceived action of the teacher is mapped onto a set of existing primitives in an assimilation phase, also suggested in Demiris and Hayes (2002) and (Wolpert *et al.*, *submitted*). This mapping process also needs to resolve the correspondence problem concerning a mismatch between the teacher’s body and the student’s body (Dautenhahn and Nehaniv, 2002). Subsequently, the most appropriate primitives are adjusted by learning to improve the performance in an accommodation phase. Figure 4 indicates such a process by highlighting the better-matching primitives with increasing linewidths. If no existing primitive is a good match for the observed behavior, a new primitive must be generated. After an initial imitation phase, self-improvement, e.g., with the help of a reinforcement-based performance evaluation criterion (Sutton and Barto, 1998), can refine both movement primitives and an assumed stage of motor command generation (see below) until a desired level of motor performance is achieved (e.g., Peters *et al.*, 2003a; Bentevegna *et al.*, 2004a).

Imitation learning requires a high level of motor and perceptual competency, none of which has been realized in robotic systems so far. Nevertheless, the field of imitation in

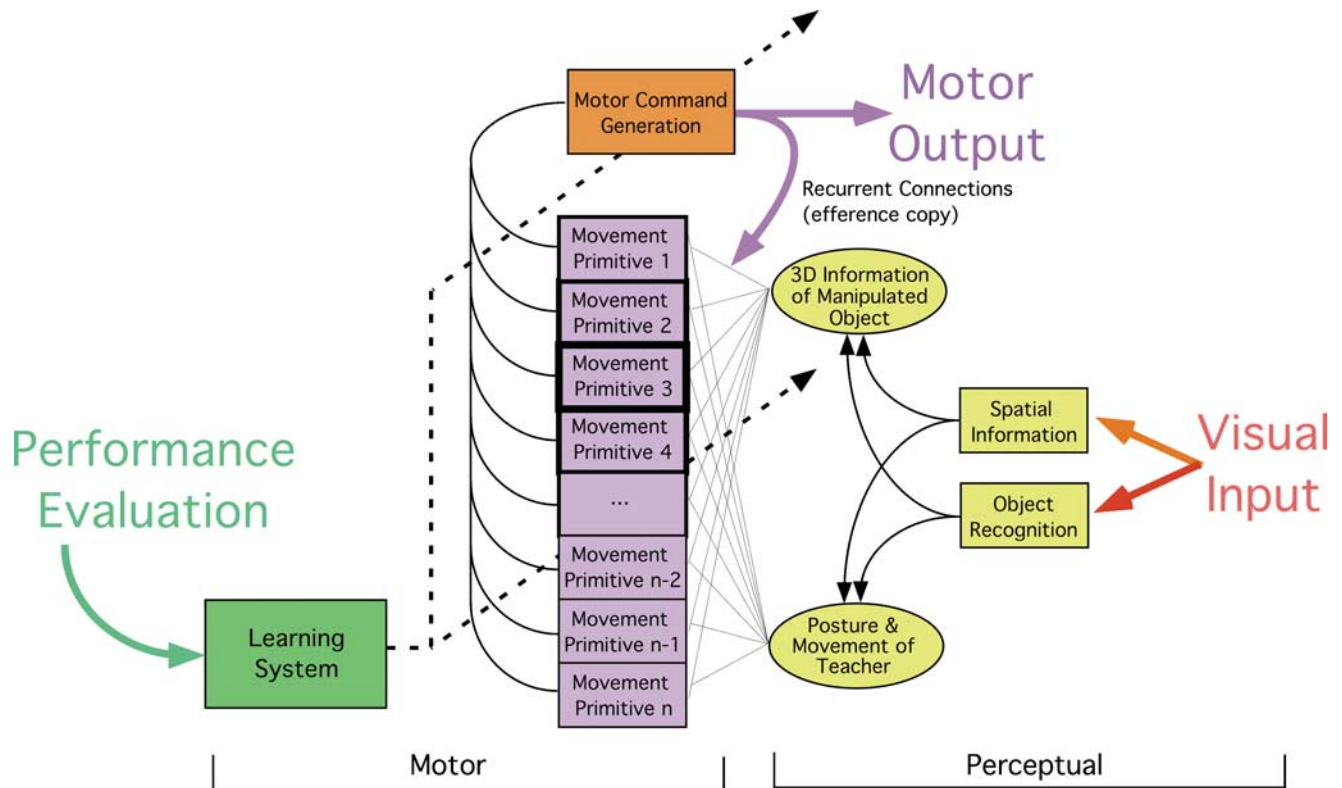


Figure 4. Conceptual sketch of an imitation learning system. The right side of the figure contains primarily perceptual elements and indicates how visual information is transformed into spatial and object information. The left side focuses on motor elements, illustrating how a set of movement primitives competes for a demonstrated behavior. Motor commands are generated from input of the most appropriate primitive. Learning can adjust both movement primitives and the motor command generator.

robotics has seen a large amount of research in the past ten years, and excellent reviews and paper collections exist (e.g., Schaal, 1999; Mataric, 2000; Dautenhahn and Nehaniv, 2002; Schaal *et al.*; Arbib, 2006).

MOVEMENT AND ACTIVITY RECOGNITION

Understanding human movement and activity will be an important component in the New Robotics, as robots will need to infer and react autonomously to the context of what is happening around them. Thus, detecting humans in the environment and recognizing who they are and what they are doing will be crucial. While most related research has originated in computer vision (e.g., for reviews, see Gavrila, 1997; Aggarwal and Cai, 1999) and is beyond the scope of this paper, from the viewpoint of movement generation, a potential shortcoming of those approaches is that representations used for recognition are mostly inappropriate for the control of movement, as they do not relate to the notion of joints, muscles, joint torques, or task goals. Given that robots in the New Robotics will not just be used for surveillance, but rather need to act upon what they perceived, e.g., as in movement imitation or physical therapy, we need action-oriented perception rather than perception in isolation.

Sparked by the discovery of “mirror neurons” in neuro-

biological research, i.e., some neurons that are situated at the interface of perception and action in humans and monkeys (Rizzolatti and Arbib, 1998; Schaal, 1999; Frith and Wolpert, 2004; Arbib, 2006), and that indeed seem to be involved in action-oriented perception, several research projects in robotics and related areas have embarked to develop corresponding theories and algorithms. That is, the modular motor control approach by Wolpert and Kawato (1998), the dynamic systems approach by Ijspeert *et al.* (Ijspeert *et al.*, 2003), and the statistical approach by Inamura *et al.* (Inamura *et al.*, 2004) and Jenkins (Jenkins, 2003) are all methods that perform action-oriented perception and can be applied to robotics. However, it needs to be emphasized that perception in these approaches is usually simplified, i.e., the authors do not use raw video inputs for movement recognition, but rather equip their subjects with markers or other means to reduce the complexity of visual processing. Thus, much more needs to be done to achieve action-oriented perception in truly natural environments.

COLLABORATIVE TASK ACHIEVEMENT

In certain situations, it will be useful if multiple robots, or robots and human, collaborate in order to achieve a particular goal. There are two kinds of collaboration: collaboration

by means of physical contact (e.g., as jointly carrying an object) and collaboration purely by means of information sharing (e.g., as in multiple robots mapping an environment). The latter domain has received a fair amount of attention in recent years due to the RoboCup competition, i.e., mobile robots playing soccer in teams (e.g., see [Bredendfeld et al., 2006](#)), and there is also a technical conference on multi-robot systems (e.g., see [Parker et al., 2005](#)). Many of the techniques developed in this research area will have bearing on the New Robotics, but expanding this paper to this domain would reduce our focus on interdisciplinary human-centered research too much.

Physical collaboration of robots (or robots and humans) is largely addressed under the name of force control in robotics, as the sum of forces acting on a jointly manipulated object are the most central for the object's behavior. A theoretically very appealing framework of such manipulation was developed by Khatib ([Khatib, 1987](#)), initially with the goal of explaining how force could be the most crucial variable in coding motor tasks (in contrast to position, velocity, or acceleration). In [Featherstone and Khatib \(1997\)](#) it was realized that Khatib's framework on force-based task-level control (a.k.a. operational space control) was very suitable for collaborative task achievement, as the special formulation of Khatib avoids complex mathematical interaction between the collaborating motor systems, i.e., each system can individually contribute its force contribution without taking the combined state of the all contributing motor systems into account. Khatib's formulation of task control remains thus one of the most appealing theoretical approaches to the generation of robotic motor skills, as it combines the ability of compliant control, hierarchical control, and even multi-robot coordination. As mentioned before, the only drawback is that these properties require quite accurate models of the robot system. It will be interesting to see in future work whether human motor control can be shown to operate according to similar principles as suggested by Khatib's work, and some work in this direction is under way ([Scholz and Schoner, 1999](#); [Todorov and Jordan, 2002](#); [Todorov, 2005](#)).

SAFETY

While this paper will not address technical issues of safety in the New Robotics, it is nevertheless important to point out that safety will be a major concern and roadblock in the development of the New Robotics. It is inevitable that there will be accidents between humans and robots in a human-centered robotics environment. The robot may not trigger such issues, but who will be able to prevent humans from accidentally pushing the robots, children to climb on them, etc.? The most powerful ways of reducing the risk of injuries is by means of lightweight and compliant robot design (e.g., [Zinn et al., 2004](#)), and by adopting recovery behaviors similar to what we humans do (e.g., [Fujiwara et al., 2002](#)).

ENGAGEMENT RESEARCH

When watching an industrial robot assembling a car, one can be fascinated by the technology for a moment, but after a few repetitions, this interest fades. Engagement research tries to understand what it takes to keep humans interested when interacting with a robot, which will be crucial in the New Robotics when robots are to become educators, physical therapists, personal assistants, or even pets. Engagement research is the least of an original robotics topic and much more rooted in psychology and cognitive sciences. In this paper, we will thus focus only on a few issues of engagement research that relate the most strongly to the technical research of robotics.

ATTENTION AND SHARED ATTENTION

Eye movements are crucial in human behavior in at least two ways. First, we use them to shift the focus of our attention to an object of interest, such that it can be inspected with the high-resolution foveal region of the retina, and second, we use our eyes to inform others about interesting objects, i.e., we share attention and elicit engagement. Additionally, the human oculomotor system has several lower level behaviors that subserve the stabilization of visual input (the vestibulo-ocular reflex, the opto-kinetic reflex, smooth pursuit, and vergence, e.g., see [Shibata et al., 2001](#)). Finally, eye movements seem to be an important component in creating complex motor sequences ([Flanagan and Johansson, 2003](#); [Flanagan et al., 2006](#)), such that sequential action generation may require understanding the principles of visual attention and planning in visual space.

From a robotics point of view, there has been a fair amount of work on creating low-level motor behaviors of vision heads (e.g., [Demiris et al., 1997](#); [Breazeal et al., 2001](#); [Shibata et al., 2001](#); [Panerai et al., 2003](#); [Aryananda and Weber, 2004](#)). So far, however, a complete implementation of all the oculomotor behaviors of the human oculomotor system is missing. As long as robot heads were mounted on static platforms, there was often not much need for a complete oculomotor system, but if the head is supposed to be on a moving humanoid robot, all the stabilization reflexes and higher-level motor behaviors will be required. It is not that any of these behaviors by itself is a technical problem, but rather the challenge lies in the combination and arbitration of all behaviors and some required learning abilities for tuning the behavioral parameters ([Shibata et al., 2001](#)).

Visual attention with robot vision heads has also been explored extensively (e.g., [Braun, 1994](#); [Driscoll et al., 1998](#); [Breazeal et al., 2001](#); [Shibata et al., 2001](#); [Vijayakumar et al., 2001](#)), but it is just recently that a rather comprehensive visual attention system for technical applications was derived from insight from neurobiology ([Koch and Ullman, 1985](#); [Itti and Koch, 2000a](#); [Itti and Koch, 2000b](#); [Itti and Koch, 2001](#); [Itti et al., 2003](#)). This work is currently on its

way to explore how intention (i.e., top-down task relevant biases) can influence attention and decision making (Navalpakam and Itti, 2005; Carmi and Itti, 2006).

Finally, shared attention, which is critical in engagement research, has started to be explored in some robotics projects (Kozima and Ito, 1998; Scassellati, 1998; Scassellati, 2002). Capabilities required for monitoring attention include gaze tracking, posture interpretation, facial expression recognition, etc. (e.g., Breazeal and Scassellati, 2002). It will take a significant amount of research and technological developments until such systems become well understood, comparable in their abilities to human performance, and practical in real-world environments.

INTENT AND USER MODELING

A rather unexplored and hard question of human-robot interaction is how to extract the intent and preferences of somebody's movement. If we understood these issues, appropriate protocols could be developed in interaction and engagement research that appeal to individual users and take their individual inclinations into account. Rather few research projects can be found that relate to these questions. A first step might be to understand human movement and activity in terms of cost functions, i.e., to associate a particular cost function with a particular behavior. Reinforcement learning generates control policies based on such cost functions and deducing the cost function from a given policy (i.e., behavior) could be called inverse reinforcement learning (Ng and Russell, 2000). Alternatively, behavioral statistics can be collected from observed behavior, which subsequently characterize certain user and movement preferences (Bentivegna *et al.*, 2004b; Billard *et al.*, 2004). Much research work needs to be done to find feasible approaches to the topic and intent and user modeling, which essentially tries to approach the topic of a "theory of mind" in algorithmic ways (Gallese and Goldman, 1998; Scassellati, 2002; Wolpert *et al.*, 2003).

MOTIVATION AND EMOTIONS

Closely related to intent and user modeling is the question of how to understand motivation and emotions of humans in an automatic way, and, of course, how to create robot behaviors in the New Robotics that relate to those characteristics of human personalities. So far, there is empirical and explorative work in the robotics community on this topic (e.g., Breazeal and Scassellati, 2002; Kanda *et al.*, 2004; Coradeschi *et al.*, 2006), usually by confronting human subjects with some form of human-like looking robot and by studying their reaction according to some psychological metrics. An interesting new approach is to look into the neuroscientific basis of motivation and emotions, investigated in the context of training and engagement. Research on training and engagement starts from the hypothesis that motor learning itself can generate internal rewards. Specifically, it is proposed that the nervous system computes the rate of motor skill learning,

which serves as an internal reward and adds to external reward terms. According to a preliminary model (Daw and Touretzky, 2002; Schweighofer and Doya, 2003; Schweighofer *et al.*, 2004), the motivation to learn will depend on the balance between successes, failures, and external rewards. In particular, the difficulty of the observed to-be-learned movement is crucial. If too easy or too hard, the rate of learning will soon become zero. Thus, only moderate and decreasing learning errors will yield a high gain in the positive feedback loop that will maximize the motivation to learn. These intuitions can be formalized in training schedules and can lead to methods that optimally keep students engaged for fast learning. A similar research methodology might be applicable to other situations in which human motivation and emotions play important roles.

CONCLUSIONS

This perspective highlighted an interdisciplinary research program that could pave the path towards a New Robotics. The vision is that, at some point in the near future, robots will be ubiquitously among us and help us to fulfill tasks in an increasingly complex society, a society that is more and more confronted by an aging population, educational problems, disasters, diseases, etc. In order to function among us, the robots in the New Robotics will need to have much more human-like abilities in their motor behaviors, and they will also be required to have appropriate social behaviors. Thus, robotics research has to embrace a novel interdisciplinary program ranging from traditional technical topics of control up to neuroscience, psychology, cognitive sciences, and ethics. This paper focused on bringing—in a noncomprehensive way—some structure to the potential research topics in the New Robotics. Some of these topics, e.g., those mentioned under *Action Research*, have already a more solid scientific foundation and are on their way to pilot studies on real robots. Other topics, particularly those enumerated under *Engagement Research*, are more in their infancy and it is the task of research to find out which are the right questions to ask, and which approaches might be promising. In some sense, the New Robotics is nothing but the old dream of artificial intelligence to create an artificial system with similar abilities than those of humans. It is just that the means and approaches to this topic have changed significantly away from the classical ideas of symbolic reasoning, and, in the wake of improved computer technology and insights into the human brain, rather pursue methods of statistical learning, computational neuroscience, and computational psychology. It is also important that the ambitions of the New Robotics are reduced: they can really focus on societal needs and make progress in well-defined problem areas, e.g., robotic physical therapy or robotic education of children with special needs (e.g., autistic children), rather than trying to create the omnipotent robotic system for all problems. In many countries, research agencies have started to fund the seedlings of the

New Robotics, and it will be exciting to witness the development of this new field over the next decade.

ACKNOWLEDGMENTS

This research was supported in part by National Science Foundation Grant Nos. ECS-0325383, IIS-0312802, IIS-0082995, ECS-0326095, ANI-0224419, NASA Grant No. AC#98-516, an AFOSR grant on Intelligent Control, the ERATO Kawato Dynamic Brain Project funded by the Japanese Science and Technology Agency, and the ATR Computational Neuroscience Laboratories.

REFERENCES

- Aggarwal, JK, and Cai, Q (1999). "Human motion analysis: A review." *Comput. Vis. Image Underst.* **73**, 428–440.
- Albu-Schäffer, A, Ott, C, and Hirzinger, G (2007). "A unified passivity based control framework for position, torque and impedance control of flexible joint robots." *Int. J. Robot. Res.* **26**, 23–39.
- An, CH, Atkeson, CG, and Hollerbach, JM (1988). *Model-Based Control of a Robot Manipulator*, MIT Press, Cambridge, MA.
- Arbib, M (1995a). *The Handbook of Brain Theory and Neural Networks*, MIT Press, Cambridge, MA.
- Arbib, MA (1995b). "Schema theory." In *The Handbook of Brain Theory and Neural Networks*, Arbib, MA (ed), pp. 830–834, MIT Press, Cambridge, MA.
- Arbib, MA (2006). *Action to Language via the Mirror Neuron System*, Cambridge University Press, Cambridge, MA.
- Arimoto, S, Sekimoto, M, Hashiguchi, H, and Ozawa, R (2005). "Natural resolution of ill-posedness of inverse kinematics for redundant robots: A challenge to Bernstein's degrees-of-freedom problem." *Adv. Rob.* **19**, 401–434.
- Aryananda, L, and Weber, J (2004). "MERTZ: a quest for a robust and scalable active vision humanoid head robot." In, Vol 2, pp. 513–532.
- Atkeson, CA, Aboaf, EW, McIntyre, J, and Reinkensmeyer, DJ (1986). "Model-based robot learning." Fourth International Symposium on Robotics Research, 2318–2330.
- Atkeson, CG, Hale, J, Kawato, M, Kotosaka, S, Pollick, F, Riley, M, Schaal, S, Shibata, S, Tevatia, G, and Ude, A (2000). "Using humanoid robots to study human behaviour." *IEEE Intell. Syst.* **15**, 46–56.
- Atkeson, CG, Moore, AW, and Schaal, S (1997a). "Locally weighted learning." *Artif. Intell. Rev.* **11**, 11–73.
- Atkeson, CG, Moore, AW, and Schaal, S (1997b). "Locally weighted learning for control." *Artif. Intell. Rev.* **11**, 75–113.
- Barto, AG, and Mahadevan, S (2003). "Recent advances in hierarchical reinforcement learning." *Discrete Event Dyn. Syst.* **13**, 341–379.
- Beek, PJ (1989). *Juggling Dynamics*, Free University Press, Amsterdam.
- Bellman, R (1957). *Dynamic Programming*, Princeton University Press, Princeton, NJ.
- Bentivegna, DC, Atkeson, CA, and Cheng, G (2004a). "Learning tasks from observation and practice." *Rob. Auton. Syst.* **47**, 163–169.
- Bentivegna, DC, Cheng, G, and Atkeson, CG (2004b). "Learning from observation and from practice using behavioral primitives." 11th International Symposium on Robotics Research (ISRR2003), Springer, Ciena, Italy.
- Billard, A, Epars, Y, Calinon, S, Cheng, G, and Schaal, S (2004). "Discovering optimal imitation strategies." *Rob. Auton. Syst.* **47**, 68–77.
- Braun, J (1994). "Visual search among items of different salience: Removal of visual attention mimics a lesion in extrastriate area V4." *J. Neurosci.* **14**, 554–567.
- Breazeal, C, Edsinger, A, Fitzpatrick, P, and Scassellati, B (2001). "Active vision for sociable robots." *IEEE Trans. Syst. Man Cybern.* **31**, 443–453.
- Breazeal, C, and Scassellati, B (2002). "Robots that imitate humans." *Trends Cogn. Sci.* **6**, 481–487.
- Bredendfeld, A, Jacoff, A, Noda, I, and Takahashi, Y (2006). *Robocup 2005: Robot Soccer World Cup IX*, Springer, New York.
- Brock, O, Fagg, AH, Grupen, R, Platt, R, Rosenstein, M, and Sweeney, J (2005). "A framework for learning and control in intelligent humanoid robots." *Int. J. Humanoid Robotics* **2**, 301–336.
- Brown, IE, Cheng, EJ, and Loeb, GE (1999). "Measured and modeled properties of mammalian skeletal muscle. II. The effects of stimulus frequency on force-length and force-velocity relationships." *J. Muscle Res. Cell Motil.* **20**, 627–643.
- Brown, IE, and Loeb, GE (1999). "Measured and modeled properties of mammalian skeletal muscle. I. The effects of post-activation potentiation on the time course and velocity dependencies of force production." *J. Muscle Res. Cell Motil.* **20**, 443–456.
- Bullock, D, and Grossberg, S (1988). "Neural dynamics of planned arm movements: Emergent invariants and speed-accuracy properties during trajectory formation." *Psychol. Rev.* **95**, 49–90.
- Burrige, RR, Rizzi, AA, and Koditschek, DE (1999). "Sequential composition of dynamically dexterous robot behaviors." *Int. J. Robot. Res.* **18**, 534–555.
- Byrne, RW, and Russon, AE (1998). "Learning by Imitation: A hierarchical approach." *Behav. Brain Sci.* **21**, 667–684; discussion 684–721.
- Carmi, R, and Itti, L (2006). "Visual causes versus correlates of attentional selection in dynamic scenes." *Vision Res.* **46**, 4333–4345.
- Collins, S, Ruina, A, Tedrake, R, and Wisse, M (2005). "Efficient bipedal robots based on passive-dynamic walkers." *Science* **307**, 1082–1085.
- Coradeschi, S, Ishiguro, H, Asada, M, Shapiro, SC, Thielscher, M, Breazeal, C, Mataric, MJ, and Ishida, H (2006). "Human-inspired robots." *IEEE Intell. Syst.* **21**, 74–85.
- Dautenhahn, K, and Nehaniv, CL, (eds) (2002). *Imitation in Animals and Artifacts*, MIT Press, Cambridge, MA.
- Daw, ND, and Touretzky, DS (2002). "Long-term reward prediction in T. D. models of the dopamine system." *Neural Comput.* **14**, 2567–2583.
- Demiris, J, and Hayes, G (2002). "Imitation as a dual-route process featuring predictive and learning components, A biologically plausible computational model." In *Imitation in Animals and Artifacts*, Dautenhahn, K, and Nehaniv, CL, (eds), pp. 327–361, MIT Press, Cambridge, MA.
- Demiris, J, Rougeaux, S, Hayes, GM, Berthouze, L, and Kuniyoshi, Y (1997). "Deferred imitation of human head movements by an active stereo vision head." IEEE International Workshop on Robot Human Communication, Sendai, Japan.
- Doya, K, Samejima, K, Katagiri, K, and Kawato, M (2002). "Multiple model-based reinforcement learning." *Neural Comput.* **14**, 1347–1369.
- Driscoll, J, Peters, R, and Cave, K (1998). "A visual attention network for a humanoid robot." Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-98).
- Dyer, P, and McReynolds, SR (1970). *The Computation and Theory of Optimal Control*, Academic.
- Featherstone, R, and Khatib, O (1997). "Load-independence of the dynamically-consistent inverse of the Jacobian matrix." *Int. J. Robot. Res.* **16**, 168–170.
- Flanagan, JR, Bowman, MC, and Johansson, RS (2006). "Control strategies in object manipulation tasks." *Curr. Opin. Neurobiol.* **16**, 650–659.
- Flanagan, JR, and Johansson, RS (2003). "Action plans used in action observation." *Nature (London)* **424**, 769–771.
- Franklin, DW, Osu, R, Burdet, E, Kawato, M, and Milner, TE (2003). "Adaptation to stable and unstable dynamics achieved by combined impedance control and inverse dynamics model." *J. Neurophysiol.* **90**, 3270–3282.
- Frith, CD, and Wolpert, DM (2004). *The Neuroscience of Social Interaction: Decoding, Imitating, and Influencing the Actions of Others*, The Royal Society, Oxford University Press, Oxford.
- Fujiwara, K, Kanehiro, F, Kajita, S, Kaneko, K, Yokoi, K, and Hirukawa, H (2002). UKEMI: falling motion control to minimize damage to biped humanoid robot. In, vol 3, pp. 2521–2526.
- Gallese, V, and Goldman, A (1998). "Mirror neurons and the simulation theory of mind-reading." *Trends Cogn. Sci.* **2**, 493–501.
- Gavrila, DM (1997). "The visual analysis of human movements: A survey." *Comput. Vis. Image Underst.* **73**, 82–98.
- Gomi, H, and Kawato, M (1996). "Equilibrium-point control hypothesis

- examined by measured arm stiffness during multijoint movement.” *Science* **272**, 117–220.
- Haruno, M, Wolpert, DM, and Kawato, M (2001). “Mosaic model for sensorimotor learning and control.” *Neural Comput.* **13**, 2201–2220.
- Ijspeert, A, Nakanishi, J, and Schaal, S (2003). “Learning attractor landscapes for learning motor primitives.” In *Advances in Neural Information Processing Systems 15*, Becker, S, Thrun, S, and Obermayer, K, (eds), pp. 1547–1554, MIT Press, Cambridge MA.
- Inamura, T, Iwaki, T, Tanie, H, and Nakamura, Y (2004). “Embodied symbol emergence based on mimesis theory.” *Int. J. Robot. Res.* **23**, 363–377.
- Itti, L, Dhavale, N, and Pighin, F (2003). “Realistic avatar eye and head animation using a neurobiological model of visual attention.” *Proceedings of the SPIE 48th Annual International Symposium on Optical Science and Technology*, 64–78.
- Itti, L, and Koch, C (2000a). “A comparison of feature combination strategies for saliency-based visual attention systems.” *Proceedings of the SPIE Human Vision and Electronic Imaging IV (HVEI99)*, 3644, 473–482.
- Itti, L, and Koch, C (2000b). “A saliency-based search mechanism for overt and covert shifts of visual attention.” *Vision Res.* **40**, 1489–1506.
- Itti, L, and Koch, C (2001). “Computational modeling of visual attention.” *Nat. Rev. Neurosci.* **2**, 194–203.
- Jenkins, OC (2003). “Data-driven derivation of skills for autonomous humanoid agents.” Department of Computer Science, University of Southern California, Los Angeles.
- Johansson, G (1973). “Visual perception of biological motion and a model for its analysis.” *Percept. Psychophys.* **14**, 201–211.
- Kajita, S, and Tani, K (1996). “Experimental study of biped dynamic walking.” *IEEE Control Syst. Mag.* **16**, 13–19.
- Kanda, T, Ishiguro, H, Imai, M, and Ono, T (2004). “Development and evaluation of interactive humanoid robots.” *Proc. IEEE* **92**, 1839–1850.
- Kawato, M (1999). “Internal models for motor control and trajectory planning.” *Curr. Opin. Neurobiol.* **9**, 718–727.
- Khatib, O (1987). “A united approach to motion and force control of robot manipulators: The operational space formulation.” *Int. J. Robot. Res.* **31**, 43–53.
- Khatib, O, Brock, O, Chang, K, Conti, F, Ruspini, D, and Sentis, L (2002). “Robotics and interactive simulation.” *Commun. ACM* **X**, 46–51.
- Khatib, O, Sentis, L, Park, J, and Warrent, J (2004). “Whole-body dynamic behavior and control of human-like robots.” *Int. J. Humanoid Robot.* **1**, 1–15.
- Koch, C, and Ullman, S (1985). “Shifts in selective visual attention: Towards the underlying neural circuitry.” *Hum. Neurobiol.* **4**, 219–227.
- Kozima, H, and Ito, A (1998). An attention-based approach to symbol acquisition. In, pp. 852–856.
- Loeb, GE (2001). “Learning from the spinal cord.” *J. Physiol.* **533**, 111–117.
- Mataric, M (1998). “Behavior-based robotics as a tool for synthesis of artificial behavior and analysis of natural behavior.” *Trends Cogn. Sci.* **2**, 82–86.
- Mataric, M (2000). “Getting humanoids to move and imitate.” *IEEE Intell. Syst.* **15**, 18–24.
- McGeer, T (1990). “Passive dynamic walking.” *Int. J. Robot. Res.* **9**, 633–643.
- Meltzoff, AN, and Moore, MK (1994). “Imitation, memory, and the representation of persons.” *Infant Behav. Dev.* **17**, 83–99.
- Menzel, P, and D’Alusio, F (2000). *Robosapiens: Evolution of a New Species*, MIT Press, Cambridge, MA.
- Miall, RC, and Wolpert, DM (1996). “Forward models for physiological motor control.” *Neural Networks* **9**, 1265–1285.
- Miyamoto, H, and Kawato, M (1998). “A tennis serve and upswing learning robot based on bi-directional theory.” *Neural Networks* **11**, 1331–1344.
- Miyamoto, H, Schaal, S, Gandolfo, F, Koike, Y, Osu, R, Nakano, E, Wada, Y, and Kawato, M (1996). “A Kendama learning robot based on bi-directional theory.” *Neural Networks* **9**, 1281–1302.
- Nakamura, Y (1991). *Advanced Robotics: Redundancy and Optimization*, Addison-Wesley, Reading, MA.
- Nakanishi, J, Mistry, M, and Schaal, S (submitted). “Inverse dynamics control with floating base and constraints.” *International Conference on Robotics and Automation (ICRA2007)*.
- Narendra, KS, and Annaswamy, AM (2005). *Stable Adaptive Systems*, Dover, Mineola, NY.
- Navalpakkam, V, and Itti, L (2005). “Modeling the influence of task on attention.” *Vision Res.* **45**, 205–231.
- Ng, AY, and Russell, S (2000). “Algorithms for inverse reinforcement learning.” *Proceedings of the Seventeenth International Conference on Machine Learning (ICML 2000)*, Stanford, CA.
- Osu, R, Kamimura, N, Iwasaki, H, Nakano, E, Harris, CM, Wada, Y, and Kawato, M (2004). “Optimal impedance control for task achievement in the presence of signal-dependent noise.” *J. Neurophysiol.* **92**, 1199.
- Panerai, F, Metta, G, and Sandini, G (2003). “Learning visual stabilization reflexes in robots with moving eyes.” *Neurocomputing* **48**, 16.
- Parker, LE, Schneider, FE, and Schultz, AC (2005). “Multi-robot systems: From swarms to intelligent automata.” Vol. III, *Proceedings from the 2005 International Workshop on Multi-Robot Systems*, Springer, Dordrecht.
- Peters, J, Mistry, M, Udwadia, FE, Cory, R, Nakanishi, J, and Schaal, S (2005a). “A unifying framework for the control of robotics systems.” *IEEE International Conference on Intelligent Robots and Systems (IROS 2005)*, Edmonton, Alberta, Canada, Aug. 2–6, 1824–1831.
- Peters, J, Mistry, M, Udwadia, FE, and Schaal, S (2005b). “A new methodology for robot control design.” *Fifth ASME International Conference on Multibody Systems, Nonlinear Dynamics, and Control (MSNDC 2005)*, Long Beach, CA, Sept. 24–28.
- Peters, J, and Schaal, S (2006). “Learning operational space control.” In *Robotics: Science and Systems (RSS 2006)*, Burgard, W, Sukhatme, GS, and Schaal, S, (eds), MIT Press, Cambridge, MA.
- Peters, J, Vijayakumar, S, and Schaal, S (2003a). “Reinforcement learning for humanoid robotics.” *Humanoids 2003, Third IEEE-RAS International Conference on Humanoid Robots*, Karlsruhe, Germany, Sept. 29–30.
- Peters, J, Vijayakumar, S, and Schaal, S (2003b). “Scaling reinforcement learning paradigms for motor learning.” *Proceedings of the 10th Joint Symposium on Neural Computation (JSNC 2003)*, Irvine, CA, May.
- Peters, J, Vijayakumar, S, and Schaal, S (2005c). “Natural Actor-Critic.” Gama, J, Camacho, R, Brazdil, P, Jorge, A, and Torgo, L (eds), *Proceedings of the 16th European Conference on Machine Learning (ECML 2005)*, vol 3720, Springer, Porto, Portugal, Oct. 3–7, pp. 280–291.
- Piaget, J (1951). *Play, dreams, and imitation in childhood*, Norton, New York.
- Righetti, L, and Ijspeert, A (2006). “Design methodologies for central pattern generators: An application to crawling humanoids.” *Proceedings of Robotics: Science and Systems*, MIT Press, Philadelphia.
- Rizzolatti, G, and Arbib, MA (1998). “Language Within Our Grasp.” *Trends Neurosci.* **21**, 188–194.
- Robins, B, Dautenhahn, K, Boekhorst, R, and Billard, A (2005). “Robotic assistants in therapy and education of children with autism: Can a small humanoid robot help encourage social interaction skills?” *Universal Access in the Information Society* **4**.
- Russell, SJ, and Norvig, P (1995). *Artificial intelligence: A Modern Approach*, Prentice-Hall, Englewood Cliffs, NJ.
- Sabes, PN (2000). “The planning and control of reaching movements.” *Curr. Opin. Neurobiol.* **10**, 740–746.
- Saltzman, EL (1979). “Levels of sensorimotor representation.” *J. Math. Psychol.* **20**, 91–163.
- Scassellati, B (1998). “Imitation and mechanisms of joint attention: A developmental structure for building social skills in a humanoid robot.” In *Computation for Metaphors, Analogy and Agents*, Nehaniv, C. L., pp. 176–195, Springer, New York.
- Scassellati, B (2002). “Theory of mind for a humanoid robot.” *Auton. Rob.* **12**.
- Schaal, S (1999). “Is imitation learning the route to humanoid robots?” *Trends Cogn. Sci.* **3**, 233–242.

- Schaal, S (2002) "Learning robot control." In *The Handbook of Brain Theory and Neural Networks*, 2nd ed., Arbib, M. A., MIT Press, Cambridge, M.A., pp. 983–987.
- Schaal, S, *Dynamic systems: Brain, body, and imitation*. In *Action to Language via the Mirror Neuron System* (in press).
- Schaal, S, and Atkeson, CG (1994). "Robot learning by nonparametric regression." *Proceedings of the International Conference on Intelligent Robots and Systems (IROS'94)*. Munich, Germany, pp. 478–485.
- Schaal, S, Atkeson, CG, and Botros, S (1992). "What should be learned?" *Proceedings of Seventh Yale Workshop on Adaptive and Learning Systems*, New Haven, CT, May 20–22, pp. 199–204.
- Schaal, S, Atkeson, CG, and Vijayakumar, S (2002). "Scalable techniques from nonparametric statistics for real-time robot learning." *Appl. Intell.* **17**, 49–60.
- Schaal, S, Ijspeert, A, and Billard, A (2003). "Computational approaches to motor learning by imitation." *Philos. Trans. R. Soc. London, Ser. B* **358**, 537–547.
- Schaal, S, Peters, J, Nakanishi, J, and Ijspeert, A. (2004). "Learning movement primitives." *International Symposium on Robotics Research (ISRR2003)*, Springer, Ciena, Italy.
- Schaal, S, and Schweighofer, N (2005). "Computational motor control in humans and robots." *Curr. Opin. Neurobiol.* **15**, 675–682.
- Schmidt, RA (1988). "Motor control and learning." *Human Kinetics*, Champaign, IL.
- Scholz, JP, and Schoner, G (1999). "The uncontrolled manifold concept: Identifying control variables for a functional task." *Exp. Brain Res.* **126**, 289–306.
- Schweighofer, N, and Doya, K (2003). "Meta-learning in reinforcement learning." *Neural Networks* **16**, 5–9.
- Schweighofer, N, Doya, K, and Kuroda, S (2004). "Cerebellar aminergic neuromodulation: towards a functional understanding." *Brain Res. Rev.* (in press).
- Sciavicco, L, and Siciliano, B (1996). *Modeling and control of robot manipulators*, MacGraw-Hill, New York.
- Sciavicco, L, and Siciliano, B (2000). *Modelling and Control of Robot Manipulators*, Springer, London.
- Sentis, L, and Khatib, O (2004). "Task-oriented control of humanoid robots through prioritization." *IEEE-RAS/RSJ International Conference on Humanoid Robots*, Santa Monica, CA, Nov. 10–12.
- Shadmehr, R, and Wise, SP (2005). "The computational neurobiology of reaching and pointing: A foundation for motor learning." MIT Press, Cambridge, Mass.
- Shibata, T, Vijayakumar, S, Conradt, J, and Schaal, S (2001). "Biomimetic oculomotor control." *Adaptive Behavior* **9**, 189–207.
- Slotine, JJ. E, and Li, W (1991). *Applied Nonlinear Control*, Prentice-Hall, Englewood Cliffs, NJ.
- Strogatz, SH (1994). "Nonlinear dynamics and chaos: With applications to physics, biology, chemistry, and engineering." Addison-Wesley, Reading, MA.
- Sutton, RS, and Barto, AG (1998). *Reinforcement Learning. An introduction*. MIT Press, Cambridge.
- Taub, E, and Wolf, SL (1997). "Constraint induced techniques to facilitate upper extremity use in stroke patients." *Topics in Stroke Rehabilitation* **3**, 38–61.
- Todorov, E (2005). "Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system." *Neural Comput.* **17**, 1084–1108.
- Todorov, E, and Jordan, MI (2002). "Optimal feedback control as a theory of motor coordination." *Nat. Neurosci.* **5**, 1226–1235.
- Tomasello, M, Savage-Rumbaugh, S, and Kruger, A C (1993). "Imitative learning of actions on objects by children, chimpanzees, and enculturated chimpanzees." *Chim. Oggi* **64**, 1688–1705.
- van der Lee, JH, Wagenaar, RC, Lankhorst, GJ, Vogelaar, TW, Deville, WL, and Bouter, LM (1999). "Forced-use of the upper extremity in chronic stroke patients: Results from a single-blind randomized clinical trial." *Stroke* **30**, 2369–2375.
- van Essen, DC, and Maunsell, JM. R (1983). "Hierarchical organization and functional streams in the visual cortex." *Trends Neurosci.* **6**, 370–375.
- Vijayakumar, S, Conradt, J, Shibata, T, and Schaal, S (2001). "Overt visual attention for a humanoid robot." *IEEE International Conference on Intelligent Robots and Systems (IROS 2001)*.
- Vijayakumar, S, D'Souza, A, and Schaal, S (2005). "Incremental online learning in high dimensions." *Neural Comput.* **17**, 1–33.
- Vijayakumar, S, and Schaal, S (2000). "Real time learning in humanoids: A challenge for scalability of online algorithms." *Humanoids 2000, First IEEE-RAS International Conference on Humanoid Robots*. CD-Proceedings, Cambridge, MA.
- Viviani, P, Baud-Bovy, G, and Redolfi, M (1997). "Perceiving and tracking kinesthetic stimuli: Further evidence of motor—perceptual interactions." *J. Exp. Psychol. Hum. Percept. Perform.* **23**, 1232–1252.
- Viviani, P, and Schneider, R (1991). "A developmental study of the relationship between geometry and kinematics in drawing movements." *J. Exp. Psychol. Hum. Percept. Perform.* **17**, 198–218.
- Wada, Y, and Kawato, M (1995). "A theory for cursive handwriting based on the minimization principle." *Biol. Cybern.* **73**, 3–13.
- Wiener, N (1948). *Cybernetics*, Wiley, New York.
- Wolf, SL, and Blanton, S (1999). "An application of upper-extremity constraint-induced movement therapy in a patient with subacute stroke." *Phys. Ther.* **79**, 847–953.
- Wolpert, DM, Doya, K, and Kawato, M (2003). "A unifying computational framework for motor control and social interaction." *Philos. Trans. R. Soc. London, Ser. B* **358**, 593–602.
- Wolpert, DM, Doya, K, and Kawato, M (submitted) "A unifying computational framework for motor control and social interaction." *Philos. Trans. R. Soc. London, Ser. B*.
- Wolpert, DM, and Kawato, M (1998). "Multiple paired forward and inverse models for motor control." *Neural Networks* **11**, 1317–1329.
- Zinn, M, Khatib, O, Roth, B, and Salisbury, JK (2004). "Playing it safe [human-friendly robots]." *Robot. Autom. Mag.* **11**, 12–21.