Chapter 1
From Motor Learning to Interaction Learning in Robots

Olivier Sigaud and Jan Peters

Abstract The number of advanced robot systems has been increasing in recent years yielding a large variety of versatile designs with many degrees of freedom. These robots have the potential of being applicable in uncertain tasks outside well-structured industrial settings. However, the complexity of both systems and tasks is often beyond the reach of classical robot programming methods. As a result, a more autonomous solution for robot task acquisition is needed where robots adaptively adjust their behaviour to the encountered situations and required tasks.
Learning approaches pose one of the most appealing ways to achieve this goal. However, while learning approaches are of high importance for robotics, we cannot simply use off-the-shelf methods from the machine learning community as these usually do not scale into the domains of robotics due to excessive computational cost as well as a lack of scalability. Instead, domain appropriate approaches are needed. In this book, we focus on several core domains of robot learning. For accurate task execution, we need motor learning capabilities. For fast learning of the motor tasks, imitation learning offers the most promising approach. Self improvement requires reinforcement learning approaches that scale into the domain of complex robots. Finally, for efficient interaction of humans with robot systems, we will need a form of interaction learning. This chapter provides a general introduction to these issues and briefly presents the contributions of the subsequent chapters to the corresponding research topics.

Olivier Sigaud
Institut des Systèmes Intelligents et de Robotique - CNRS UMR 7222
Université Pierre et Marie Curie
Pyramide Tour 55 - Boîte Courrier 173, 4 Place Jussieu, 75252 Paris CEDEX 5, France
e-mail: Olivier.Sigaud@upmc.fr

Jan Peters
Max Planck Institute for Biological Cybernetics, Dept. Schölkopf
Spemannstrasse 38, 72076 Tübingen, Germany
e-mail: mail@jan-peters.net

Key words: motor learning, interaction learning, imitation learning, reinforcement learning, robotics

1.1 Introduction

Robot learning has reached an unprecedented amount of interest in recent years. However, as the robotics domain is of particular complexity for learning approaches, it has become quite demanding for students and young researchers to get started in this area. Furthermore, due to the current high speed of development it is often hard for scientists from other areas to follow the developments. For this reason, the need for an easy entrance into this field has become strong while it is not yet time for a robot learning textbook. The idea of this book arose from these considerations.

This chapter serves two purposes: firstly, it allows us to quickly familiarize the reader with the background. In Section 1.1.1, we give an overview on the importance of robot learning approaches at this moment. In Section 1.1.2, we discuss essential background on motor, imitation and interaction learning and recommend the reader to briefly survey this part before diving into the respective chapters. Secondly, we give in Section 1.2 a brief overview on the chapters included in this book, before concluding in Section 1.3 on the necessity of a more integrated effort.

1.1.1 The Need for Robot Learning Approaches

At the beginning of the 21st century, robotics research is experiencing large changes in its aims and objectives. In most of the previous century, the majority of all operational robot were performing the same manufacturing task again and again in extremely structured environments such as automobile factories. Often it was easier and cheaper to build a new factory with new robots to accommodate a new car model than to reprogram an existing one. By contrast, robots are now "leaving" factory floors and start becoming part of the everyday life of average citizens. Vacuum cleaning robots have become the most sold robots to date with 4-5 million units shipped up to 2008, and programmable entertainment robots such as the RoboSapiens have become many children's favorite toy. This evolution raises the major challenge of "personalizing" the programming of our robots and making them compatible with human-inhabited environments. As a result, a variety of new issues arise that we will discuss below.

First, robots will often be in physical contact with people that are not specially trained to interact with them, thus they must be less dangerous. This concern first implies some mechanical requirements: robots must become lighter and their actuators must have inherent compliance properties as human muscles have. But in turn, these changes result in the necessity to think differently about their control loops. Either we stay with the same kind of actuator technology and we must use extremely low gains while yielding sufficient accuracy, or we come to completely new actuators like artificial muscles where the classical control knowledge is missing and learning techniques will play an important role. Finally, in any case they must never become unsafe in unforeseen situations. All these considerations result in the necessity to move from the previous standard way of thinking about robot control to new approaches that rely more on on-line, adaptive model identification and autonomous action selection properties.

Second, future robots need to be more versatile and more flexible when encountering some of the infinitely many potential situations that are part of our daily life. Despite the impressive results of human manual plan design and robot programming, these hand-crafted solutions are not likely to transfer to that large variety of different tasks and environmental states. Hence, it is becoming increasingly clear that a new approach is essential. To interact better with their environment, robots will need more and more sensing capabilities but also algorithms that can make use of this richer sensory information. Due to this increase of complexity both on the perception and action side, robots will need to *learn* the appropriate behavior in many situations. This challenge is becoming recognized in the general robotics community. It has resulted in supportive statements by well-known roboticists such as "*I have always said there would come a time for robot learning — and that time is now*" by O. Khatib (at Stanford, October 2006) and "*robot learning has become the most important challenge for robotics*" by J. Hollerbach (at the NIPS Workshop on Robotics Challenges for Machine Learning, December 2007).

Third, apart from the security aspects of human robot interaction that we have already outlined, the fact that robots will be in interaction with people must also be taken into account in their control and learning architecture. Thinking of a robot as interacting with people has a lot of consequences in the design of their control architecture, as will be investigated in subsequent parts of this book. As we will see, robots can be instructed from their human user how to perform a new task by imitation, or by physical or language-based interaction during task execution, etc.

All these changes on the way to consider robots and their control comes along with another major evolution about the hardware platforms available. In the last twenty years, a huge technological effort has resulted in the design of more complex, more efficient and more flexible platforms with the challenges above in mind. In particular, these last years have seen the emergence of humanoid robots of diverse size, capabilities and price as well as a variety of bimanual mobile robotics platforms. In these platforms, all the challenges listed above are essential problems that always need to be addressed.

This large shift in robotics objectives has resulted in an increasing visibility of the corresponding lines of research. In the last few years, we have seen an increasing amount of robot learning publications both at top machine learning (such as NIPS, ICML and ECML) conferences and mainstream robotics conferences (particularly at R:SS, ICRA and IROS). The number of learning tracks has been increasing the IEEE multi-track conferences ICRA and IROS and there have been at least 12 workshops on robot learning in 2007–2009. This development has resulted in numerous special issues in excellent robotics international journals such as the *International*

Journal of Robotics Research, *Autonomous Robots*, the *International Journal of Humanoid Robots* and the *IEEE Robotics & Automation Magazine*. Recently, it even gave rise to the creation of an IEEE Technical Committee on Robot Learning.

All these considerations led us to consider that this is the good time to publish a book about Motor learning, Imitation and Interaction Learning in Robots. In the next section, we will highlight the relationships between the corresponding different subfields before giving an overview of the contributions to the book.

1.1.2 Motor learning, Imitation learning and Interaction Learning

Making humanoid robots perform movements of the same agility as human movements is an aim difficult to achieve even for the simplest tasks. Although a lot of impressive results have been obtained based on pure hand-coding of the behaviour, this approach seems too costly and, probably, even too difficult if we ever want humanoid robots and mobile manipulators to leave research labs or factories and enter human homes.

The most immediate alternative to this manual programming is *imitation learning*, also known as *learning from demonstration* or *programming by showing*. This approach is relatively well-developed and has resulted in a variety of excellent results in previous work. It includes several different teaching approaches. Some researchers record human motion in the context of a task with motion capture tools and transfer the motion on the robot, which implies solving the *correspondence problem* resulting from the differences between the mechanics of a human being and a robot system (e.g., already mapping the human arm kinematics on a non-anthropomorphic robot arm is a difficult problem). Hence, it is often easier to employ teleoperating interfaces for teaching, or even use the robot by itself as a haptic device. Furthermore, different approaches are employed in order to recover a policy; while the approaches discussed in this book directly mimic the observed behaviors, there is an alternative stream of research that employs an *inverse reinforcement learning* approach which rather attempts to recover the intent of the teacher by modeling his cost function and, subsequently, derives the policy that is optimal with respect to the cost-to-go (Abbeel and Ng, 2004; Coates et al, 2008; Ratliff et al, 2009). While several chapters in this book are dedicated to imitation learning and should yield a good start in this area, we want to point out to the reader that several important groups in imitation learning are not covered and we urge the reader to study (Atkeson and Schaal, 1997; Schaal, 1999; Ijspeert et al, 2003; Abbeel and Ng, 2004; Calinon et al, 2007; Coates et al, 2008; Ratliff et al, 2009).

Nevertheless, learning from demonstration does not suffice as the behavior of the robot will be restricted to the behaviors that have been demonstrated, even if some generalization mechanisms can slightly remediate that situation, allowing for instance adaptation to slightly changing contexts. To go beyond an initial imitation, we need the robots to adapt online so that they can react to new situations. There exist a few situations where such an adaptation can be achieved purely by supervised learning, e.g., when the functional relationship can be directly observed as in inverse dynamics model learning and a relearning after a change of the dynamics due to an external load or a failure is straightforward. However, the majority of all problems require some kind of self-improvement, e.g., we need to adapt elementary movements to an unforeseen situation, improve a policy learned from a demonstration for better performance, learn new combinations of movements or even simply learn an inverse model of a redundant system. Addressing these problems is often formalized in the *reinforcement learning* framework, which mimics the way animals and humans improve their behaviour by trial-and-error in unforeseen situations. A key issue in reinforcement learning is exploration: since we do not know in advance which behaviour will give rise to high outcome and which will not, we have to try various different actions in order to come up with an efficient strategy.

Taken as a whole, learning of general motor capabilities, at the control level or at the behavioral strategy level, is a very hard problem. This involves the exploration of a huge space of possibilities where a lot of standard algorithmic steps boil down to hard optimization problems in a continuous domain. Given the difficulty of the exploration problem in that domain, the combination of reinforcement learning with imitation learning has been shown to be fruitful. Here, imitation provides an efficient way to initialize policies so that the explorative policy can focus on the behaviors or strategies that have a high probability of being efficient.

Finally, the third robot learning topic that we cover in this book is interaction learning. Interaction learning allows the robot to discover strategies for actively using the contact with human in its proximity. It often shares tools and methods with imitation learning, since both approaches have to take the presence of the human around the robot into account. In fact, imitation learning is a particular case of interaction learning in the sense that imitation is a particular type of interaction. However, interaction learning is not restricted to reproduce the observed behavior of a human. Instead, interacting means getting jointly involved in a common activity both taking the behavior of the respective other into account. Thus, interaction can be physical, when the human user actually exerts some force onto the robot or, conversely, when the robot does so to the human user. It can also be communicative, either through diverse modalities of language or through communicative gestures. It can finally be purely implicit, when the robot and the user try to adapt their behaviour to the other without any direct communication, just through observing. Interaction learning provides a challenging context for motor learning in general. Human motor behavior is often difficult to predict and, thus, interaction may require learning non-stationary models of the dynamics of the coupling between humans and robots.

Last, but not least, the study of human motor behavior requires a deep understanding of the connections between motor, imitation and interaction learning. For example, neurophysiological studies of human subjects suggest that motor learning processes and more cognitive learning and developmental processes have much in common, particularly when it comes to interaction with other beings. After the much celebrated discovery of the so called ”mirror neurons” relating motor learning to imitation and language acquisition, several neurophysiological studies have revealed that brain areas generally considered as motor, such as the cerebellum, or dedicated to action selection, like the basal ganglia, are in fact employed in more general cognitive functions such as learning tool use, imitation, language and so forth (Doya, 1999). Taken together, these facts advocate for a hierarchical understanding of the brain architecture where motor learning and interaction learning are tightly coupled processes at the root of cognition (Wolpert et al, 2003; Demiris and Khadhouri, 2006). These topics are highly relevant for robot learning as the human motor system is still the best prototype for us to study in order to obtain new and better algorithms.

1.2 Overview of the book

From the previous section, it has become apparent that *Motor, Imitation* and *Interaction Learning* are highly dependent on each other and complementary. The purpose of the book is to provide a state-of-the-art view of these different subfields within the same volume so as to cast the basis for an improved dialog between them. We have divided the book into three parts, but there is a strong overlap between the topics covered by these parts.

1.2.1 Biologically inspired Models for Learning in Robots

A common view in most learning approaches to robotics is that humans exhibit all the properties we want from a robot system in terms of adaptivity, learning capabilities, compliance, versatility, imitation and interaction capabilities etc. Hence, it might be a good idea to be inspired by their functionality and, as a result, a lot of robot learning approaches are bio-inspired in some sense. More precisely, in this book we can distinguish three different sources of inspiration in this line of thinking.

The first one has to do with trying to implement robot controllers on a representation that is as similar as possible to the neural substrate that one can find in the human motor system. The complexity of the computational models resulting from this line of thinking raises the problem of their validation. Here, robotics plays a prominent role as a tool to evaluate the capability of these wide scope models to account for the phenomena they address. In this book, two chapters, (Duff et al, 2009) and (Lagarde et al, 2009), are following this line of thinking. The first one proposes a biologically based cognitive architecture called Distributed Adaptive Control to explain the neuronal organization of adaptive goal oriented behavior. The second one, based on neural field models, is interested in low level, basic imitation mechanisms present early in the newborn babies, showing how the different proprioceptive signals used in the examples can be seen as bootstrap mechanisms for more complex interactions.

A second line of inspiration consists in trying to reproduce the learning properties of the human motor system as observed from outside, building models that rely on computational principles that may explain these observed properties. The work of Mitrovic et al (2009) illustrates this approach. It proposes an efficient implementation of a model of motor adaptation based on well accepted computational principles of human motor learning, using optimal feedback control methods that require a model of the dynamics of the plant in a context where this model is learned. The work of Herbort et al (2009) shares similar goals, but the authors address slightly different motor learning phenomena, with a particular focus on motor preparation. The authors propose an implementation of their system based on artificial neural networks on a simulated complex robot, perfectly illustrating the highly cross-disciplinary nature of this domain.

Finally, a third line of inspiration takes its sources in developmental psychology, giving rise to the so called developmental robotics or epigenetic robotics (Lungarella et al, 2004). In some sense, the work already discussed by Lagarde et al (2009) can also fall into this category. Moreover, the work in (Oudeyer et al, 2009) is a prominent example of this line of thinking, investigating how a model of activity selection based on curiosity can give rise to the capability to tackle more and more difficult tasks within a life-long learning paradigm.

1.2.2 Learning Models and Policies for Motor Control

The differences between papers about biologically inspired models for learning in robots and the one that fall into this technical part is often small. A lot of work about learning models and policies for motor control is also inspired by biological considerations but does not attempt to provide an explanation for biological behavior. As there are important differences between the mechanics of the human musculo-skeletal system and the mechanical design of robots, severe limitations are imposed on the degree of similarity between natural and artificial controllers. Indeed, for instance, the human musculo-skeletal system has the amazing ability to control both stiffness and position of each joint independently from each other due to co-contraction. By contrast, nearly all humanoid robots to date are having a single motor per joint and, thus, offer either position access (e.g., cheap RoboSapiens designs), setting desired velocities (e.g., the Fujitsu Hoap, iCub and many others) or are torque controlled (e.g., the SARCOS humanoids). This makes robots controllers unable to make profit of the nice properties of the muscles that human people use in practice and this drives robotics control towards control principles that may differ a lot from those observed in human movement. In such a context when the standard engineering knowledge is not well developed, the contribution from (Fumagalli et al, 2009) compares two learning techniques, namely Least Squares Support Vector Machines and Neural Networks, on their capability to estimate the forces and torques measured by a single six-axis force/torque sensor placed along the kinematic chain of a humanoid robot arm.

Beyond these considerations, the chapters regrouped in this part fall in two categories. The first category is about learning models of the plant, either direct or inverse, at the kinematics, velocity kinematics and dynamics level. This kind of work, giving rise to motor adaptation capabilities, is one of the main mechanisms to obtain compliance and versatility in robots. The contribution from Salaün et al (2009) provides an overview of how learned kinematics and velocity kinematics models can be used within a feedback control loop in the Operational Space Control framework. Learning these models is a difficult self-supervised learning problem in large continuous state and action spaces, thus having an efficient learning method is crucial. A few learning techniques have emerged in the last years as particularly competitive to address this task. In particular, the most recent Locally Weighted Regression methods give rise to very fast implementations that scale well and are able to tackle large problems but suffer from the necessity to tune a lot of parameters, whereas methods based on Gaussian Processes, are computationally more intensive as the size of the problems grows but require less tuning. The chapter by Nguyen-Tuong et al (2009) proposes a local method based on Gaussian Processes that combines the good properties of both families of approaches. They are able to show that the model works well in the context of learning inverse dynamics.

The second category of contributions in this part is about finding the good computational framework to derive efficient controllers from learning principles. In that domain, an important approach is about the automatic tuning of motor primitives, that already provided convincing results (e.g., as Ijspeert et al (2003)). But whereas in these previous approaches primitives were based on open-loop control, the chapter by Kober et al (2009) provides an extension to the case where primitives incorporate perceptual coupling to external variables, giving rise to closed-loop policies.

Taking a very different view, the chapter by Toussaint and Goerick (2009) presents a bayesian formulation of classical control techniques based on task space to joint space mapping, that results in the possibility to consider motor execution and motor planning as a unified bayesian inference mechanism. The chapter highlights the interesting robustness properties of the resulting framework and highlights deep relationships with the optimal control framework that provides convincing computational principles for motor control (Todorov and Jordan, 2002; Todorov, 2004; Todorov and Li, 2005). Still in the same category but based on different principles, (Howard et al, 2009) is focused on learning from trajectories a controller able to realize a set of tasks subject to a set of unknown constraints. Finally, the contribution from Roberts et al (2009) describes a nice application of a model-free reinforcement learning-based control methodology, based on an optimized policy gradient algorithm, to the control of an experimental system dedicated to the study of flapping wing flight.

1.2.3 Imitation and interaction learning

This part starts with a chapter by Lopes et al (2009) which provides an overview of imitation learning methods in robots. It provides a presentation of some recent developments about imitation in biological systems, as well as a focus on robotics work that consider self-modelling and self-exploration as a fundamental part of the cognitive processes required for higher-level imitation.

The contribution (Chalodhorn and Rao, 2009) describes an approach based purely on human motion capture to achieve stable gait acquisition in a humanoid robot despite its complex mechanical structure. The chapter gives a good example of the theoretical difficulties and technical intricacies that must be faced in such kind of imitation learning approaches given the “correspondence problem” that must be solved between the human musculo-skeletal system with its many redundant degrees of freedom and robot systems with their different & well-defined kinematic structures. The chapter insists on dimensionality reduction techniques that can be used to simplify the resolution of the correspondence problem.

After a chapter focused on learning one particular motor primitive from imitation, the chapter by Kulić and Nakamura (2009) proposes a broader approach for autonomous and incremental organization of a set of such primitives learned by observation of human motion, within a life-long learning framework. The hierarchical organization makes it easier to recognize known primitives and to determine when adding a new primitive in the repertoire is necessary. The different motor primitives are represented by Hidden Markov Models or by Factorial Hidden Markov Models.

The contribution of Grollman and Jenkins (2009) is focused on the case where some task is decomposed into a set of subtasks. With a more critical standpoint than previous chapters, it examines the limits of a regression-based approach for learning a Finite State Machine controller from demonstration of a basic robot soccer goal-scoring task, based on an Aibo robot.

We already discussed in the first section of this chapter that there is a lot of potential in the combination of imitation learning (or learning from demonstration) and automatic improvement of control policies. The chapter from Argall (2009) falls into this category. It presents an approach for the refinement of mobile robot control policies, that incorporates human teacher feedback.

The chapter (Detry et al, 2009) describes a method that combines imitation learning with actual interaction with object to learn grasp affordances. More precisely, the work is about learning to grasp objects described by learned visual models from different sources of data. The focus is on the organization of the whole knowledge that an agent has about the grasping of an object, in order to facilitate reasoning on grasping solutions and their likelihood of success.

The last two chapters are more focused on interaction learning. First, the chapter from Hörnstein et al (2009) is about language acquisition in humanoid robots, based on interaction with a caregiver and using as few built in a priori knowledge or primitives as possible. The importance of motor learning in the language acquisition process is underlined. Second, the contribution from Lallee et al (2009) presents an outstanding integration effort towards language based interaction between a robot and a non-expert user in the context of a cooperation between them. The focus is put on the use of the Spoken Language Programming approach to facilitate the interaction.

1.3 Conclusion and perspectives

Robot learning is a young, fruitful and exciting field. It addresses problems that will become increasingly important for robotics as the platforms get more and more complex and the environment get less and less prepared or structured. The reader will find in this book research works stemming from different areas – statistical learning of models, reinforcement learning, imitation and interaction learning – that all contribute to the global endeavour of having more adaptive robots able to deal with more challenging settings, in particular those where interaction with humans is involved. In Section 1.1.2, we highlighted some ways in which diverse research efforts could be combined given the complementary subproblems they address. However, when closing this book, the reader will probably have the feeling that the different contributors are working within isolated frameworks and that a global coordination effort is still missing. Our view as editors is that finding frameworks giving rise to the possibility of such coordination is the next step in the field, and we hope this book will play its role towards this next step.

References

Abbeel P, Ng AY (2004) Apprenticeship learning via inverse reinforcement learning. In: Proceedings of the International Conference on Machine Learning

Argall BD (2009) Mobile robot motion control from demonstration and corrective feedback. In: From Motor to Interaction Learning in Robots, Springer, pp 430–449

Atkeson CG, Schaal S (1997) Robot learning from demonstration. In: Proceedings of the International Conferenec on Machine Learning, pp 12–20

Calinon S, Guenter F, Billard A (2007) On Learning, Representing and Generalizing a Task in a Humanoid Robot. IEEE Transactions on Systems, Man and Cybernetics, Part B 37(2):286–298

Chalodhorn R, Rao RPN (2009) Learning to imitate human actions through eigenposes. In: From Motor to Interaction Learning in Robots, Springer, pp 356–380

Coates A, Abbeel P, Ng AY (2008) Learning for control from multiple demonstrations. In: Proceedings of the International Conference on Machine Learning

Demiris Y, Khadhouri B (2006) Hierarchical attentive multiple models for execution and recognition of actions. Robotics and Autonomous Systems 54:361–369

Detry R, Baseski E, Popovi M, Touati Y, Krüger N, Kroemer O, Peters J, Piater J (2009) Learning continuous grasp affordances from sensorimotor interaction. In: From Motor to Interaction Learning in Robots, Springer, pp 450–465

Doya K (1999) What are the computations of the cerebellum, the basal ganglia, and the cerebral cortex? Neural Networks 12:961–974

Duff A, César R, Costa R, Marcos E, Luvizotto AL, Giovannucci A, Sanchez Fibla M, Bernardet U, Verschure PFMJ (2009) Distributed adaptive control: A proposal on the neuronal organization of adaptive goal oriented behavior. In: From Motor to Interaction Learning in Robots, Springer, pp 15–41

Fumagalli M, Gijsberts A, Ivaldi S, Jamone L, Metta G, Natale L, Nori F, Sandini G (2009) Learning how to exploit proximal force sensing: a comparison approach. In: From Motor to Interaction Learning in Robots, Springer, pp 149–169

Grollman DH, Jenkins OC (2009) Can we learn finite state machine robot controllers from interactive demonstration? In: From Motor to Interaction Learning in Robots, Springer, pp 405–429

Herbort O, Butz MV, Pedersen G (2009) The sure reach model for motor learning and control of a redundant arm: from modeling human behavior to applications in robots. In: From Motor to Interaction Learning in Robots, Springer, pp 83–104

Hörnstein J, Gustavsson L, Santos-Victor J, Lacerda F (2009) Multimodal language acquisition based on motor learning and interaction. In: From Motor to Interaction Learning in Robots, Springer, pp 466–489

Howard M, Klanke S, Gienger M, Goerick C, Vijayakumar S (2009) Methods for learning control policies from variable-constraint demonstrations. In: From Motor to Interaction Learning in Robots, Springer, pp 255–289

Ijspeert AJ, Nakanishi J, Schaal S (2003) Learning attractor landscapes for learning motor primitives. In: Advances in Neural Information Processing Systems (NIPS), pp 1523–1530

Kober J, Mohler B, Peters J (2009) Imitation and reinforcement learning for motor primitives with perceptual coupling. In: From Motor to Interaction Learning in Robots, Springer, pp 210–226

Kulić D, Nakamura Y (2009) Incremental learning of full body motion primitives. In: From Motor to Interaction Learning in Robots, Springer, pp 381–404

Lagarde M, Andry P, Gaussier P, Boucenna S, Hafemeister L (2009) Proprioception and imitation: on the road to agent individuation. In: From Motor to Interaction Learning in Robots, Springer, pp 42–62

Lallee S, Yoshida E, Mallet A, Nori F, Natale L, Metta G, Warneken F, Dominey PF (2009) Human-robot cooperation based on interaction learning. In: From Motor to Interaction Learning in Robots, Springer, pp 490–538

Lopes M, Melo F, Montesano L, Santos-Victor J (2009) Abstraction levels for robotic imitation: Overview and computational approaches. In: From Motor to Interaction Learning in Robots, Springer, pp 313–355

Lungarella M, Metta G, Pfeifer R, Sandini G (2004) Developmental robotics: a survey. Connection Science 0:1–40

Mitrovic D, Klanke S, Vijayakumar S (2009) Adaptive optimal feedback control with learned internal dynamics models. In: From Motor to Interaction Learning in Robots, Springer, pp 63–82

Nguyen-Tuong D, Seeger M, Peters J (2009) Real-time local gp model learning. In: From Motor to Interaction Learning in Robots, Springer, pp 195–209

Oudeyer PY, Baranes A, Kaplan F (2009) Intrinsically motivated exploration and active sensorimotor learning. In: From Motor to Interaction Learning in Robots, Springer, pp 105–146

Ratliff ND, Silver D, Bagnell JA (2009) Learning to search: Functional gradient techniques for imitation learning. Autonomous Robots 27(1):25–53

Roberts JW, Zhang J, Tedrake R (2009) Motor learning at intermediate reynolds number: Experiments with policy gradient on a heaving plate. In: From Motor to Interaction Learning in Robots, Springer, pp 290–310

Salaün C, Padois V, Sigaud O (2009) Learning forward models for the operational space control of redundant robots. In: From Motor to Interaction Learning in Robots, Springer, pp 170–194

Schaal S (1999) Is imitation learning the route to humanoid robots? Trends in Cognitive Sciences 6:233–242

Todorov E (2004) Optimality principles in sensorimotor control. Nature Neurosciences 7(9):907–915

Todorov E, Jordan MI (2002) Optimal feedback control as a theory of motor coordination. Nature Neurosciences 5(11):1226–1235

Todorov E, Li W (2005) A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems. In: Proceedings of the American Control Conference, pp 300–306

Toussaint M, Goerick C (2009) A bayesian view on motor control and planning. In: From Motor to Interaction Learning in Robots, Springer, pp 227–254

Wolpert DM, Doya K, Kawato M (2003) A unifying computational framework for motor control and social interaction. Philosophical Transactions of the Royal Society 358:593–602