

Advances in Experimental Medicine and Biology 782

Michael J. Richardson
Michael A. Riley
Kevin Shockley *Editors*

Progress in Motor Control

Neural, Computational and
Dynamic Approaches

 Springer

ADVANCES IN EXPERIMENTAL MEDICINE AND BIOLOGY

Volume 782

Editorial Board:

IRUN R. COHEN, *The Weizmann Institute of Science*

ABEL LAJTHA, *N. S. Kline Institute for Psychiatric Research*

JOHN D. LAMBRIS, *University of Pennsylvania*

RODOLFO PAOLETTI, *University of Milan*

For further volumes:

<http://www.springer.com/series/5584>

Michael J. Richardson • Michael A. Riley
Kevin Shockley
Editors

Progress in Motor Control

Neural, Computational and Dynamic
Approaches

 Springer

Editors

Michael J. Richardson
University of Cincinnati
Cincinnati, Ohio
USA

Michael A. Riley
University of Cincinnati
Cincinnati, Ohio
USA

Kevin Shockley
University of Cincinnati
Cincinnati, Ohio
USA

ISSN 0065-2598

ISBN 978-1-4614-5464-9

ISBN 978-1-4614-5465-6 (eBook)

DOI 10.1007/978-1-4614-5465-6

Springer Dordrecht Heidelberg London New York

Library of Congress Control Number: 2012954565

© Springer Science+Business Media New York 2013

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed. Exempted from this legal reservation are brief excerpts in connection with reviews or scholarly analysis or material supplied specifically for the purpose of being entered and executed on a computer system, for exclusive use by the purchaser of the work. Duplication of this publication or parts thereof is permitted only under the provisions of the Copyright Law of the Publisher's location, in its current version, and permission for use must always be obtained from Springer. Permissions for use may be obtained through RightsLink at the Copyright Clearance Center. Violations are liable to prosecution under the respective Copyright Law.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Preface

This volume is the most recent installment of the *Progress in Motor Control* series. It contains contributions based on presentations by invited speakers at the Progress in Motor Control VIII meeting held in Cincinnati, OH, USA in July, 2011. Progress in Motor Control is the official scientific meeting of the International Society of Motor Control (ISMC). There were 23 invited presentations at the meeting, which was organized into eight themed symposia and included a special ISMC Past President's Address by Michael Turvey, along with 137 poster presentations.

The Progress in Motor Control VIII meeting, and consequently this volume, were meant to provide a broad perspective on the latest research on motor control in humans and other species. The invited talks at the meeting addressed topics such as neural regeneration, the mirror neuron system, movement disorders, dynamical systems models and analyses, cortical representation and control of movement, spinal circuitry for movement control, neuromechanics, motor learning, computational modeling, and interactions between cognitive and motor processes. Neuroscience, psychology, physiology, kinesiology, biomechanics, engineering, neurology, physics and applied mathematics are among the disciplines represented by the chapters and their authors. The chapters also reflect a broad range of approaches and theoretical points of view, including neural, computational, and dynamical systems perspectives.

This diversity of perspectives and approaches, while certainly not exhaustive or even fully representative, provides a flavor of the complex and multi-faceted nature of motor coordination and control. While it is clear that much progress has been made—fueled in part, hopefully, by the eight Progress in Motor Control meetings to date and the publications associated with them—it is nonetheless apparent that a thorough and complete understanding of motor control is not yet within our grasp. It will require a sustained effort to achieve this understanding, and continued efforts to synthesize the results of studies that are accruing at what seems to be an exponentially increasing rate. We hope that this volume contributes to these important goals in at least some small way.

We would like to acknowledge the extremely valuable help of Jamie Miller and the University Conferencing staff who helped us plan and execute Progress in Motor Control VIII. Thanks are also due to the graduate students from the Perceptual-Motor Dynamics Laboratory at the Center for Cognition, Action, and Perception in

the University of Cincinnati Psychology Department who helped make the meeting run smoothly—Dilip Athreya, Laura Bachus, Scott Bonnette, Tehran Davis, Nikita Kuznetsov, MaryLauren Malone, Michael Tolston, Julie Weast-Knapp, and Eli White. The meeting was supported financially by the National Institute of Neurological Disorders and Stroke (grant number 1R13NS073205-1). We also appreciate the input from past and present ISMC officers and previous Progress in Motor Control organizers. Finally, we would like to thank Arthur Smilios at Springer for his encouragement and assistance with putting together this volume.

Cincinnati, Ohio, USA
April 2012

Michael J. Richardson
Michael A. Riley
Kevin Shockley

Contents

1 Model-Based and Model-Free Mechanisms of Human Motor Learning	1
Adrian M. Haith and John W. Krakauer	
2 The Molecular Basis of Experience-Dependent Motor System Development	23
Robert G. Kalb, Weiguo Zhou and Lei Zhang	
3 Neurocognitive Mechanisms of Error-Based Motor Learning	39
Rachael D. Seidler, Youngbin Kwak, Brett W. Fling and Jessica A. Bernard	
4 Plasticity in the Motor Network Following Primary Motor Cortex Lesion	61
Numa Dancause	
5 The Mirror System in Monkeys and Humans and its Possible Motor-Based Functions	87
Leonardo Fogassi and Luciano Simone	
6 A Molecular Basis for Intrinsic Muscle Properties: Implications for Motor Control	111
Kiisa C. Nishikawa, Jenna A. Monroy, Krysta L. Powers, Leslie A. Gilmore, Theodore A. Uyeno and Stan L. Lindstedt	
7 Theoretical and Methodological Issues in Serial Correlation Analysis	127
Didier Delignières and Vivien Marmelat	
8 On the Control of Unstable Objects: The Dynamics of Human Stick Balancing	149
Ramesh Balasubramaniam	
9 Intermittent Motor Control: The “drift-and-act” Hypothesis	169
John G. Milton	
Index	195

Contributors

Ramesh Balasubramaniam Sensorimotor Neuroscience Laboratory, McMaster University, Hamilton, ON L8S 4K1, Canada
e-mail: ramesh@mcmaster.ca

Jessica A. Bernard Department of Psychology, University of Michigan, 401 Washtenaw Avenue, Ann Arbor, MI 48109-2214, USA
e-mail: jessbern@umich.edu

Numa Dancause Groupe de Recherche sur le Système Nerveux Central (GRSNC), Département de Physiologie, Pavillon Paul-G-Desmarais, Université de Montréal, 2960, Chemin de la Tour, bureau 4138, Montréal, Québec H3T 1J4, Canada
e-mail: numa.dancause@umontreal.ca

Didier Delignières EA 2991 Movement to Health - Euromov, UFR STAPS, University Montpellier 1, 700 avenue du Pic Saint Loup, 34090 Montpellier, France
e-mail: didier.delignieres@univ-montp1.fr

Brett W. Fling School of Kinesiology, University of Michigan, 401 Washtenaw Avenue, Ann Arbor, MI 48109-2214, USA
e-mail: bfling@hs.uci.edu

Leonardo Fogassi Dipartimento di Neuroscienze e Istituto Italiano di Tecnologia (RTM), Università di Parma, Parma, Italy

Dipartimento di Psicologia, Università di Parma, Parma, Italy
e-mail: leonardo.fogassi@unipr.it

Leslie A. Gilmore Department of Biological Sciences, Northern Arizona University, Flagstaff, AZ 86011-5640, USA

Adrian M. Haith Department of Neurology, Johns Hopkins University, Baltimore, MD 21205, USA
e-mail: adrian.haith@jhu.edu

John W. Krakauer Departments of Neurology and Neuroscience, Johns Hopkins University, Baltimore, MD 21205, USA
e-mail: jkrakau1@jhmi.edu

Robert G. Kalb Department of Neurology, Children's Hospital of Philadelphia, Perelman School of Medicine, University of Pennsylvania, 3615 Civic Center Boulevard, Philadelphia, PA 19104, USA

Research Institute and Division of Neurology, Department of Pediatrics, Children's Hospital of Philadelphia, Perelman School of Medicine, University of Pennsylvania, 3615 Civic Center Boulevard, Philadelphia, PA 19104, USA
e-mail: kalb@email.chop.edu

Youngbin Kwak Neuroscience Program, University of Michigan, 401 Washtenaw Avenue, Ann Arbor, MI 48109-2214, USA

Center for Cognitive Neuroscience, Duke University, Durham, NC 27708, USA
e-mail: youngbin.kwak@duke.edu

Stan L. Lindstedt Department of Biological Sciences, Northern Arizona University, Flagstaff, AZ 86011-5640, USA

Vivien Marmelat EA 2991 Movement to Health - Euromov, UFR STAPS, University Montpellier 1, 700 avenue du Pic Saint Loup, 34090 Montpellier, France

John G. Milton W. M. Keck Science Department, The Claremont Colleges, Claremont, CA 91711, USA
e-mail: jmilton@jsd.claremont.edu

Jenna A. Monroy Department of Biological Sciences, Northern Arizona University, Flagstaff, AZ 86011-5640, USA

Kiisa C. Nishikawa Department of Biological Sciences, Northern Arizona University, Flagstaff, AZ 86011-5640, USA
e-mail: Kiisa.Nishikawa@nau.edu

Krysta L. Powers Department of Biological Sciences, Northern Arizona University, Flagstaff, AZ 86011-5640, USA

Rachael D. Seidler Department of Psychology and School of Kinesiology, University of Michigan, 401 Washtenaw Avenue, Ann Arbor, MI 48109-2214, USA
e-mail: rseidler@umich.edu

Luciano Simone Dipartimento di Psicologia, Università di Parma, Parma, Italy

Theodore A. Uyeno Department of Biology, Valdosta State University, Valdosta, GA 31698-0015, USA

Lei Zhang Research Institute and Division of Neurology, Department of Pediatrics, Children's Hospital of Philadelphia, Perelman School of Medicine, University of Pennsylvania, 3615 Civic Center Boulevard, Philadelphia, PA 19104, USA
e-mail: ZhangL@email.chop.edu

Weiguo Zhou Research Institute and Division of Neurology, Department of Pediatrics, Children's Hospital of Philadelphia, Perelman School of Medicine, University of Pennsylvania, 3615 Civic Center Boulevard, Philadelphia, PA 19104, USA
e-mail: weiguo Zhou@gmail.com

Chapter 1

Model-Based and Model-Free Mechanisms of Human Motor Learning

Adrian M. Haith and John W. Krakauer

Introduction

In laboratory settings, motor learning has typically been studied in the context of adaptation paradigms in which subjects must learn to compensate for a systematic perturbation—either some manipulation of visual feedback (Krakauer et al. 2000) or a change in the dynamics of the motor apparatus, e.g., a force applied to the hand (Shadmehr and Mussa-Ivaldi 1994), Coriolis forces induced by rotation of the body (Lackner and Dizio 1994), or an inertial load attached to the arm (Krakauer et al. 1999). What is typically observed in these tasks is a monotonic improvement in performance that is initially rapid, and then slows to an asymptote close to initial baseline levels of performance. The progress of learning is well described by exponential fits, implying that the amount of improvement on each trial is proportional to the error (Thoroughman and Shadmehr 2000; Donchin et al. 2003). This kind of fast, trial-by-trial reduction in systematic errors is typically referred to as *adaptation*. The term adaptation has been used in some cases to imply a particular mechanism of learning; however, we will adhere to a behavioral definition (as a gradual reduction in error following an abrupt change in conditions) and describe potential underlying learning mechanisms in more computational terms. As we will argue, learning in adaptation paradigms is likely predominantly mediated by a specific learning mechanism that is based on changing an internal forward model.

Not all motor learning falls under our behavioral definition of adaptation. Often one learns to synthesize entirely novel movements even when there is no perturbation, e.g., learning to swing a golf club, hit a tennis serve, balance a pole, or drive a car. Although this kind of learning corresponds more closely to everyday usage of

A. M. Haith (✉)

Department of Neurology, Johns Hopkins University, Baltimore, MD 21205, USA
e-mail: adrian.haith@jhu.edu

J. W. Krakauer

Departments of Neurology and Neuroscience, Johns Hopkins University,
Baltimore, MD 21205, USA
e-mail: jkrakau1@jhmi.edu

the term “motor learning,” it has hardly been studied in laboratory settings. The few exceptions typically involve learning to manipulate an unfamiliar, possibly complex virtual object (Carmena et al. 2003; Mosier et al. 2005; Nagengast et al. 2009; Sternad et al. 2011). In these kinds of tasks, subjects progress from initial incompetence to a high degree of proficiency, even approaching theoretically optimal behavior. However, performance improvements are far slower than in adaptation paradigms: while tens of trials are usually enough to reach asymptote after a systematic perturbation is introduced, performance in these more complex tasks continues to improve over hundreds of trials and even across days. This slow improvement is not entirely due to the unfamiliarity of the task. Even in much more simple tasks that involve maneuvering a cursor along a constrained path (Shmuelof et al. 2012) or through a series of via points (Reis et al. 2009), overall variability in task performance reduces substantially over days of practice, even though subjects immediately exhibit near-perfect performance at slow speeds. It appears that a qualitatively different kind of learning may be occurring in these tasks—one that is not reliant on compensating for the highly salient errors that are present in adaptation settings, but instead is associated with incrementally improving the quality of one’s movements with practice. We define this long-term reduction in movement variability as *skill learning*. It is not currently clear whether adaptation, skill, and learning to control external objects draw upon identical, overlapping, or entirely different neural mechanisms.

In this review, we argue for the existence of two distinct mechanisms underlying motor learning: (1) a model-based system in which improvements in motor performance occur indirectly, guided by an internal forward model of the environment which is updated based on prediction errors, and (2) a direct, model-free system in which learning occurs directly at the level of the controller and is driven by reinforcement of successful actions. These distinct learning systems are each suited to different tasks and as such are complementary to one another. Model-based processes are likely to predominate in adaptation paradigms, and model-free processes predominate in skill tasks. However, we argue that both can contribute to learning in any given task.

Theory: Model-Based and Model-Free Approaches to Learning Control Policies

We adopt a general definition of motor learning as the process of improvement in execution of a task according to some chosen measure of performance such as increased chance of success or decreased effort (or potentially a combination of the two). Formally, we describe the state of learning in terms of a control policy π mapping current states, stimuli and time to motor commands u_t ,

$$u_t = \pi(x_t, s_t, t).$$

This general framework can encompass multiple levels of description. A control policy could describe selection of a single action per trial or describe an ongoing

stream of motor commands in continuous time according to the instantaneous state. The motor commands u_t could model a high-level decision such as which direction to move the hand or a low-level decision such as which muscles to activate and when. The stimulus s_t would typically correspond to an observed target location and the state x_t would reflect the state of the motor plant. Any systematic, experience-driven change in this control policy can be described as motor learning. The quality of each potential control policy can be quantified in terms of the expected outcome value, i.e., the average performance that would be expected to be obtained when following that control policy for a given task. In studying motor learning, we study the process whereby individuals use experience to improve their control policy.

The optimal policy will depend on two specific things: (1) the structure of the task, i.e., which states are associated with valuable or successful outcomes and what costs may be associated with different states or actions and (2) the dynamics of the motor apparatus and environment, i.e., how do motor commands affect the state. In most motor control paradigms we would generally expect that the structure of the task is unambiguous; however, in general it may be that neither the task structure nor the dynamics dynamics is known precisely.

This general framework and the problem of determining suitable actions in an uncertain environment based on ongoing observations is precisely the theoretical problem studied, at a more abstract level, in the field of *reinforcement learning* (Sutton and Barto 1998). At the heart of reinforcement learning theory is the notion of the *value function* $V(x_t, t)$ which reflects, for a given control policy, the total future reward that can be expected to be gained given the current state and time. The goal of reinforcement learning is to determine the optimal value function—from which the optimal policy follows straightforwardly.

Different variants of reinforcement learning differ in exactly what kind of value function is represented and how this value function is updated based on experience. In particular, two distinct computational strategies have emerged for using experience to update estimated values and thereby determine optimal control policies. The first approach is to use experience to build models of the dynamics of the motor apparatus and environment and the structure of the task, and compute the value function based on these models (Fig. 1.1a). This approach is termed *model-based learning*. Note that model-based learning of this kind is very different from what most people understand intuitively by the term ‘reinforcement learning.’

A second approach, which accords with most people’s informal or colloquial use of the term ‘reinforcement learning,’ is to learn the value function directly through a process of trial and error—explore the space of potential actions in each state and keep track of which states and actions lead, either directly or indirectly, to successful outcomes (Fig. 1.1b). This approach is often termed *model-free*, in contrast to model-based approaches. Other learning strategies are clearly possible besides the model-free and model-based approaches described here. However, these represent the most common approaches.

While model-free strategies clearly work and can in certain cases be shown theoretically to be guaranteed to converge upon optimal behavior (Sutton and Barto 1998), learning by trial and error is typically very slow in terms of the number of attempts necessary before a good policy can be acquired, even in relatively simple




Fig. 1.1 Comparison of model-based and model-free strategies for updating a control policy based on experience. **a** Model-based learning schematic. Changes to the control policy are brought about indirectly through first updating a forward model of the motor apparatus and environment based on sensory prediction errors, then using this knowledge to calculate an appropriate controller for the current task. **b** Model-free learning schematic. The control policy is updated directly based on reward prediction errors

environments. Model-based learning, by contrast, makes the best possible use of all observations. Any information acquired about the outcome of a particular action is retained and can influence planning of future movements, regardless of whether that action led to success or not. Model-based methods also allow more principled generalization. If the reward structure of the task changes in a known way (e.g., the target moves to a new location), an appropriate new control policy can be computed based on the model of the dynamics that was built in the context of the previous task.

The major disadvantage with a model-based approach is that although the value of any state/action pair can in principle be computed exactly, it can be prohibitively computationally intensive to do so. Existing methods for computing the optimal policy typically involve either dynamic programming—a backward iteration through time to exhaustively compare all possible paths to the target and identify the best ones, or some iterative sequence of approximations to the value function or policy that converge upon a local optimum. Details of these methods are beyond the scope of this chapter (but for excellent introductions see Bertsekas 1996; Sutton and Barto 1998; Todorov 2007).

The complexity associated with computing optimal value functions and policies need not preclude biological systems from utilizing some form of model-based control. In certain very simple scenarios, it can be trivial to compute the optimal policy given a specification of the task and plant e.g., if the action on a given trial is simply the aiming direction for a particular movement, then a model-based solution to a rotational perturbation simply amounts to subtracting an estimate of the rotation angle from an observed target angle. The feasibility of the model-based approach therefore largely depends on the nature of the task. Even if the computations are simple, however, errors may still arise from accumulation of noise that inevitably accompanies computations in biological systems (McGuire and Sabes 2009).

Model-free approaches, by contrast, require only relatively trivial computations because experiences lead directly to changes in the controller. Unlike a model-based approach, there is no intermediate forward model representation and no calculation required to transform a forward model into a control policy. In the long-run, model-free approaches tend to deliver superior performance on a particular task because they do not rely so heavily on noisy computations each and every time a movement must be made. The disadvantage is that the scope of the learned control policy is restricted to the task performed during learning. Even if the reward structure of the task changes in a known way, one must start from scratch (or at least from some previous but incorrect control policy). This is in sharp contrast to the flexibility offered by model-based learning.

In summary, if one wants to learn a good control policy in an uncertain environment, a model-based learning strategy is, in principle, the most powerful and flexible approach but requires unwieldy computations. Direct, model-free approaches rely only on simple computations but can require far more training (exploration) before they lead to a competitive policy. What learning strategy do animals use when placed in a situation where they must learn what to do? The contrasting ways in which model-based and model-free learning mechanisms should be expected to generalize to novel scenarios can act as hallmarks that potentially allow us dissociate an animal's learning strategy based on observing its behavior.

Model-Based and Model-Free Learning in Operant Conditioning Paradigms

In a situation where an animal must learn what actions will lead to reward, such as a rat navigating through a maze to find food, it seems that animals adopt both model-based and model-free learning mechanisms in parallel (Daw et al. 2005). Although any given control policy could be arrived at by either model-based or model-free strategies, these two modes of control can be dissociated by changing the reward structure of the task. In rodents this is typically achieved by stimulus revaluation. For instance, imagine examining the behavior of thirsty rats in a maze that they had learned while they were hungry and seeking food. Under a model-free approach, the thirsty rat will have no way of knowing how to obtain water and will likely either behave like a naïve rat, or rely on the same policy that led to reward while hungry. A model-based approach, by contrast, will enable the rat to flexibly change its behavior immediately in line with its new objective of finding water instead of food (provided, of course that it had previously explored the maze sufficiently to have found the location of the water). In practice, rewards are typically revalued either by satiating the animal prior to the task or, more drastically, pairing a familiar food with a strongly aversive stimulus (e.g., poison).

Behavior in such devaluation paradigms has been studied extensively, leading to a classical division between *goal-directed* behavior, in which animals are sensitive to reward devaluation, and *habitual* behavior in which they are not (Killcross and

Coutureau 2003; Balleine and O’Doherty 2010; van der Meer and Redish 2011). Behavior tends to be goal-directed early in learning but becomes more habitual later on (Balleine and O’Doherty 2010). These differences in behavior can be interpreted in terms of reinforcement learning: goal-directed behavior can be understood as model-based, while habitual behavior is model-free (Daw et al. 2005; Dayan 2009). The transition from goal-directed to habitual with experience can even be explained as an evolving, intelligent trade-off between the advantages of each strategy.

Remarkably, these alternative model-based and model-free strategies are neurally dissociable. Lesions to distinct regions of prefrontal cortex can isolate one pattern of behavior or another in hungry rats (Balleine and Dickinson 1998; Killcross and Coutureau 2003). Sequential decision-making tasks in humans have revealed that their behavior can similarly be decomposed into model-based and model-free components (Fermin et al. 2010; Gläscher et al. 2010), while fMRI reveals that these components have distinct underlying neural substrates (Gläscher et al. 2010).

The kinds of control tasks that we are primarily concerned with in this review are quite different from the problem that a rat faces in a maze. In decision-making tasks it is the high-level choice of which path to follow at a junction that is of interest. The low-level movements that register this decision are considered incidental. In motor control, however, it is precisely these low-level movements that are of interest. Critically, control of movements can be cast within the same broad theoretical framework used to describe decision-making. The only differences are that movements of the eyes and limbs occupy a space of potential states and actions that is continuous and potentially high-dimensional, and decisions must be made in continuous time. Nevertheless, the same considerations for solving the general problem apply as in more discrete domains. In particular, both model-based and model-free learning strategies are possible and have similar advantages and disadvantages as in discrete domains. We will argue that, as in the case of rodent decision-making, both strategies are employed by the motor system for continuous control of movement. The underlying neurophysiology may, however, be quite different for the motor system as compared with the discrete action selection paradigms studied in rodents.

Model-Based Motor Learning

Forward models—neural networks which generate predictions about future states of the motor system given a current state and an outgoing motor command—have long been posited to be utilized by the motor system (Wolpert and Miall 1996). Model-based learning has become a dominant framework for understanding human motor learning, with arguably the majority of theories of motor learning assuming a model-based perspective (Shadmehr and Krakauer 2008; Shadmehr et al. 2010). The proposed advantages of maintaining a forward model are twofold: (1) A forward model allows for faster and more precise estimation of the state of the body and/or environment, and (2) Forward models may aid in planning future movements by directing changes in the controller itself, i.e. they may participate in model-based control. While (1) has by now become a relatively uncontroversial claim, (2) is much more difficult to establish.

Before assessing the case for model-based learning in the motor system, we briefly disambiguate model-based learning from learning involving *inverse models*. The simplest kind of controller considered in motor learning theories is a static mapping from a desired outcome to a single action. Such controllers have been referred to as *inverse models* since they are the direct inverse of the forward model. However, inverse models are not really “models” in the true sense of the word—they do not provide an internal representation of any process occurring in the outside world. It is more accurate to think of inverse models as simple control policies. An inverse model control policy can be arrived at in a model-based manner by first learning a forward model and then inverting it (Jordan 1992). Alternatively, changes to an inverse model could be driven directly by task errors (Thoroughman and Shadmehr 2000). We would not describe such learning as model-based, however, since the learning occurs directly at the level of the controller rather than via a forward model representation of the task or plant. Learning of this kind is only really feasible in simple, single-time step scenarios.

Theories based on the notion of inverse models are fairly limited in scope. More generally, motor control is described in terms of time-dependent feedback control policies (Todorov and Jordan 2002). In this context, there is no way to directly update the control policy based on performance errors. By contrast, model-based learning is a very general approach to obtaining a good control policy that is applicable to any problem that can be framed as a Markov decision process. The only limitation to model-based learning is being able to gather enough information to build the model.

Nothing is presently known about the neural computations that underlie the translation of knowledge about the environment in the form of a forward model into a control policy. However, even though the potential mechanisms underlying model-based control processes are poorly understood, this understanding is not necessary to establish whether or not it occurs. Here, we focus on reviewing the evidence at the behavioral level for the existence of forward models and their involvement in motor learning.

The Cerebellum and Forward Models

The cerebellum has long been implicated in motor control and coordination and has emerged as the most likely neural substrate of putative internal models (Bastian 2006; Shadmehr and Krakauer 2008; Wolpert et al. 1998). Patients with hereditary cerebellar ataxia or lesions to the cerebellum have general difficulties in coordinating movement and are grossly impaired in adaptation tasks (Martin et al. 1996; Maschke et al. 2004; Smith and Shadmehr 2005; Tseng et al. 2007; Synofzik et al. 2008; Rabe et al. 2009; Criscimagna-Hemminger et al. 2010; Donchin et al. 2011). There are many potential roles for the cerebellum in learning that might give rise to such an adaptation deficit in cerebellar ataxia. The cerebellum may, for instance, compute an inverse model that directly maps desired outcomes to actions (Medina 2011). We argue here, however, that the adaptation deficit following cerebellar damage stems from an inability to learn forward models.

Neurophysiological recordings from the cerebellum show that Purkinje cell simple spike activity reflects the kinematics of movement, and not the motor commands required to achieve the kinematics (Pasalar et al. 2006). This finding clearly demonstrates that the output of the cerebellum is not directly related to motor output, as would be predicted if the cerebellum were computing an inverse model or otherwise contributing directly to control. Furthermore, Purkinje cell activity during movement precedes the actual kinematic state of the limb (Roitman et al. 2005). So this activity in the cerebellum does not simply reflect a reporting of sensory feedback—instead it appears that the cerebellum implements an internal forward model that predicts the kinematic or sensory consequences of motor commands before that information actually becomes available from the periphery.

Numerous studies have argued from a behavioral standpoint that an estimate of state from a forward model underlies state-specific feedback corrections during movement (Ariff et al. 2002; Chen-Harris et al. 2008; Wagner and Smith 2008; Munuera et al. 2009). This process appears to be cerebellar-dependent (Miall et al. 2007; Xu-Wilson et al. 2009). Together with the above-mentioned neurophysiological findings, these studies make a strong case that the cerebellum generates predictions about future motor states on the basis of outgoing motor commands, and that these predicted states are made available to an already-learned feedback controller that guides ongoing execution of a movement. While this constitutes model-based control of sorts (Mehta and Schaal 2002), in this article we are more interested in the question of whether a forward model brings about changes in the controller, rather than influencing control only through estimates of state. Nevertheless, if forward models exist and can be used to guide online feedback control, it perhaps makes it more likely that the same forward models might participate in planning feedforward control.

Evidence for Forward Model Involvement in Feedforward Control

An often-cited instance where predictions of a forward model are claimed to influence feedforward control (as opposed to only feedback control) is in compensating for the consequences of one effector's actions on another—for instance stabilizing one hand holding a load while removing that load with the other hand or increasing grip force on an object to prevent slippage when accelerating it upward. It has often been argued that such *anticipatory control* is possible because of a forward model that predicts adverse consequences of an upcoming action before it has happened, enabling an appropriate compensation to be planned and executed concomitantly (Wolpert and Miall 1996; Flanagan and Wing 1997; Wolpert et al. 2011). Although the use of a forward model could, in principle, enable this kind of anticipatory control, coordination *per se* is no proof of the existence of forward models. Anticipatory control is simply a feature of a good control policy and there is no way of knowing how this controller may have been arrived at simply by observing it in action. Good coordination could have been learned via model-free mechanisms through trial and error.

Studies of anticipatory control in cerebellar ataxia patients offer some clues as to the nature of anticipatory control. Interestingly, cerebellar ataxic patients demonstrate intact coordination in manual unloading tasks (Diedrichsen et al. 2005) and exhibit intact modulation of grip force with varying load forces (although baseline grip forces are abnormal) (Rost et al. 2005), suggesting that forward models are not at all a prerequisite for performing coordinated movement. Cerebellar patients do, however, show impairment in learning novel anticipatory adjustments (Nowak et al. 2004; Diedrichsen et al. 2005). This suggests that initial acquisition of anticipatory control is facilitated by a forward model that can predict the consequences of the actions of one effector on the goals of another but, with prolonged practice, coordinated control eventually becomes independent of the forward model. To put it another way, there may be a transition from model-based to model-free mechanisms.

The notion of model-based learning implies that improvements in performance are driven by errors in the prediction of a forward model. One plausible alternative to this idea is that adaptation is driven by the feedback corrections one makes to correct errors, rather than by the errors themselves (Kawato and Gomi 1992). This does not appear to be the case for reaching movements, however: learning rates in adaptation tasks are identical whether or not feedback corrections are allowed during movement (Tseng et al. 2007). Similarly, corrective saccades do not appear to be necessary to adapt saccade amplitude (Wallman and Fuchs 1998).

Although not driven by corrective movements, adaptation may not necessarily be driven by prediction errors of a forward model. If control is mediated by an inverse model, changes to a control policy could be driven directly by task errors, without any need for a forward model. In most cases, task errors and prediction errors are closely aligned. In certain cases, however, performance errors and prediction errors can be dissociated. For instance, saccades to visual targets usually tend to fall slightly short of the target, but this shortfall does not lead to an increase in saccadic gain as one would expect if it were induced through a target jump. In fact, if the target is surreptitiously jumped mid-saccade such that the eye lands perfectly on the target every time, then saccadic gain actually begins to decrease despite the absence of performance errors (Wong and Shelhamer 2011). Indeed it is even possible for adaptation to occur in the opposite direction from a task error. This provides compelling evidence that prediction errors and not task errors are what drive motor adaptation.

A similar, even more striking result can be found for reaching movements. In a study by Mazzoni and Krakauer (2006) (Fig. 1.2a), subjects were exposed to a 45° rotation of visual feedback but were also provided an explicit strategy to counter the rotation: simply aim to an adjacent target deliberately spaced at a 45° separation from the true target. Initially, subjects were able to flawlessly implement the strategy and hit the target. However, performance rapidly began to drift away from the target in the direction of the perturbation despite the fact that the task was being performed without errors. It therefore does not seem to be task error *per se* that drives adaptation, but discrepancies between predicted and observed behavior. Interestingly, this drift effect does not persist indefinitely—after prolonged exposure, subjects begin to reduce their errors again, suggesting that there is some component of learning that

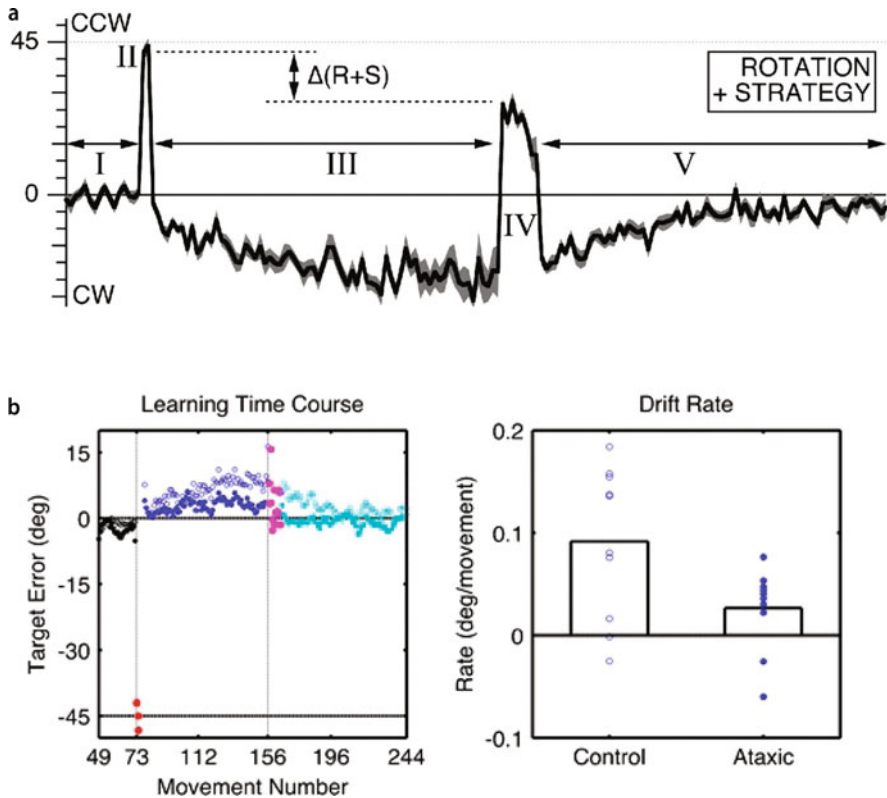


Fig. 1.2 Motor learning is driven by sensory prediction errors. **a** Healthy subjects that are provided with an explicit strategy to counter a 45° rotation initially counter the perturbation successfully, but performance immediately drifts in the direction opposite the rotation. (Reproduced from Mazzoni and Krakauer 2006). **b** This drift is attenuated in patients with cerebellar ataxia (note that rotation direction is opposite as compared with panel a). (Reproduced from Taylor et al. 2010)

acts to close task errors rather than prediction errors (Taylor and Ivry 2011). When patients with cerebellar ataxia are given an explicit strategy, they are able to successfully maintain performance without undergoing any drift in performance (Taylor et al. 2010) (Fig. 1.2b). Thus the adaptation deficit in cerebellar ataxic patients is due to a reduced sensitivity to prediction errors not task errors.

The idea that adaptation is mediated by changes in predictions about the consequences of one's actions can be tested more directly through paradigms that ask subjects to estimate where they perceived their hand to have moved during a reach. Although such assays inevitably contaminate forward model-based predictions with actual visual and proprioceptive sensory experiences, a number of interesting results have been obtained using this approach. Following exposure to rotated visual feedback, healthy subjects undergo a corresponding change in their perceived hand path during movement (Synofzik et al. 2006). Cerebellar ataxic patients show no such

perceptual changes (Synofzik et al. 2008; Izawa et al. 2011). These results support the idea that changes in a forward model, which presumably lead to the changes in predicted hand position, are a prerequisite for adaptation.

In summary, adaptation is driven by prediction errors and not by task errors or online motor responses to correct those errors. Exposure to rotated visual feedback leads to a shift in perceived hand location during movement. In patients with cerebellar ataxia, sensory prediction errors do not result in changes in feedforward control in future trials and do not lead to changes in perceived hand position. We believe that the most parsimonious explanation for all of these results is that the cerebellum computes an internal forward model that predicts the consequences of motor commands and that this forward model influences feedforward control of future movements.

Generalization of Learning Across Tasks

A final thread of evidence that has been cited in support of model-based control frameworks concerns generalization. Human subjects exhibit a high degree of generalization of learned compensation for a perturbation to a new movement (Shadmehr and Mussa-Ivaldi 1994; Krakauer et al. 2000). While this generalization is consistent with the idea of model-based control, it is important to bear in mind that model-free learning will also be expected to exhibit some degree of generalization—only in this case the generalization will be of a learned control policy, rather than of an internal model. The amount of generalization across states will be entirely determined by the underlying representation. There is no specific reason why one should expect model-based learning to generalize more broadly across states than model-free. However, subjects trained on a visuomotor rotation with full vector error (presumably engaging primarily model-based mechanisms) do generalize more broadly than subjects who learned to compensate the same perturbation but were given only binary feedback about the success or failure of their movements (presumably relying on model-free learning) (Izawa and Shadmehr 2011).

A more concrete dissociation between model-based and model-free learning mechanisms is the extent to which learning should transfer across tasks within the same workspace—for instance tracking a cursor along a curved path versus making point-to-point reaches. This form of generalization across tasks is directly analogous to the reward devaluation protocols that dissociate model-based from model-free action selection processes in rodents (Daw et al. 2005)—in both cases the reward structure of the task is altered but the consequences of actions remain the same. A number of studies have examined generalization of learning from a redundant task, in which the perturbation is task-irrelevant, to a nonredundant task. For example, Schaefer et al. (2012) had subjects make reaching movements to a point anywhere on a circular arc while imposing a rotation of visual feedback. This rotation did not compromise task performance, since subjects still easily landed on the arc as required. The rotation, however, did lead to sensory prediction errors. In subsequent catch trials toward a single target, subjects showed significant aftereffects, supporting the idea that learning was driven by sensory prediction errors rather than by task