

# Principles of sensorimotor learning

Daniel M. Wolpert\*<sup>||</sup>, Jörn Diedrichsen<sup>#||</sup> and J. Randall Flanagan<sup>§||</sup>

**Abstract** | The exploits of Martina Navratilova and Roger Federer represent the pinnacle of motor learning. However, when considering the range and complexity of the processes that are involved in motor learning, even the mere mortals among us exhibit abilities that are impressive. We exercise these abilities when taking up new activities — whether it is snowboarding or ballroom dancing — but also engage in substantial motor learning on a daily basis as we adapt to changes in our environment, manipulate new objects and refine existing skills. Here we review recent research in human motor learning with an emphasis on the computational mechanisms that are involved.

## Optimal

A system is said to be optimal if it minimizes some cost function under given constraints.

## Saccade

A rapid movement of the eyes that changes fixation from one point to another.

## Visuomotor mapping

Typically, the relationship between the hand's actual and visual locations that can be altered using devices (such as a prism) or virtual reality to examine visuomotor learning.

\*Department of Engineering, University of Cambridge, Trumpington Street, Cambridge CB2 1PZ, UK.

<sup>||</sup>Institute of Cognitive Neuroscience, University College London, London WC1N 3AR, UK.

<sup>§</sup>Department of Psychology and Centre for Neuroscience Studies, Queen's University, Kingston, Ontario K7L 3N6, Canada.

<sup>||</sup>All authors contributed equally to this work  
Correspondence to D.M.W.  
e-mail: wolpert@eng.cam.ac.uk

doi:10.1038/nrn3112  
Published online  
27 October 2011

Humans show a remarkable capacity to learn a variety of motor skills, ranging from tying shoelaces to hitting a tennis ball. Learning such skills involves a number of interacting elements. First, there are different task components that must be learned for skilled performance, including efficient gathering of task-relevant sensory information, decision making and selection of strategies, and the implementation of both predictive and reactive control mechanisms. Second, there are different learning processes that apply to these components, which specify how errors and rewards drive learning. Finally, learning is strongly determined by the neural representations of motor memory that influence how we assign credit during learning and how learning generalizes to novel situations. In recent years, there have been substantial advances in our understanding of the computations that underlie these three elements. Here, we primarily focus on empirical and computational studies of the learning of sensorimotor behaviours rather than on studies of the neural circuits that underlie this behaviour.

## Components of motor learning

**Information extraction.** Skilled performance requires the effective and efficient gathering and processing of sensory information relevant to an action. This is an active process because what we see, hear and touch is influenced by our movements. For example, the motor system controls the eyes' sensory stream by orienting the fovea to points of interest within the visual scene. Studies have shown that eye movements can be driven both in a bottom-up, task-independent manner based on low-level features of the visual scene<sup>1</sup> (for example, towards moving high-contrast objects) as well as in a top-down, task-dependent manner<sup>2</sup>. In some laboratory-based visual search tasks it has been shown that the eye movements are at least in qualitative agreement with an

optimal extraction of task-relevant information<sup>3</sup>. That is, during visual search for a target among distractors, people choose to saccade to the location that will minimize their uncertainty over possible target locations. This suggests that the motor system is involved in active learning, choosing where to sample the sensory input in a way that is most informative to the task at hand. Studies of eye movements in motor tasks have also shown that gaze behaviour is highly task-specific. Although relatively few studies have examined gaze behaviour during motor learning, in real-world tasks it has been shown that eye movements can betray the difference between skilled and amateur performers (for a review see REF. 4). For example, a cricket batsman will make a predictive saccade to the place where he expects a bowled ball to hit the ground, wait for it to bounce, and use a pursuit eye movement to follow the ball's trajectory after the bounce<sup>5</sup>. A shorter latency for this first saccade distinguishes expert from amateur batsmen. When learning to control a cursor to hit targets under a novel visuomotor mapping, eye movements change across learning stages, with gaze reactively chasing the cursor in early learning and simply fixating the target after extensive practice<sup>6</sup>. This suggests that different information is required (and extracted) during different stages of learning.

After the motor system has been used to sample the sensory world, it is still critical to selectively extract task-relevant information, as attentional and processing resources are limited. This includes filtering the incoming sensory information based on the current action (for example, REFS 7,8). One example of how the brain filters out task-irrelevant information in purely perceptual tasks is the phenomenon of inattention blindness, in which people fail to notice prominent stimuli in the visual scene that are irrelevant to the task that they are performing<sup>9</sup>. Similarly, in motor tasks subjects are often

only aware of large sensory input changes that have a bearing on the task at the precise time of the change, and are unaware of such changes otherwise<sup>10</sup>. One possible interpretation of these results is that information may be extracted or processed only at the moment when it is needed to drive action.

Motor learning itself can also push the limits of what our perceptual system can do. For example, expert video game players develop an extraordinary ability to extract information and spread their attention over a wide spatial frame without any apparent decrease in attentional performance<sup>11</sup>. In addition, recent studies have shown that motor learning can change basic sensory processing. For example, adapting speech motor commands to compensate for forces that are applied to the jaw by a robotic manipulandum results in a change in the perceptual classification of speech sounds<sup>12</sup>.

Sensory streams are temporally delayed and tend to be corrupted by appreciable amounts of noise<sup>13</sup>. Given the stream of sensory input, there are at least three computations that can improve the accuracy of the sensory information and that can be understood within the framework of Bayesian inference. First, multiple streams of sensory information, within and across modalities (for example, visual and tactile inputs), can be optimally combined to achieve estimates that reduce the effects of noise (for a review see REF. 14). Interestingly, this integration process can take into account the properties of external objects, such as tools, so that the visuo-haptic integration is optimal even when the tactile input comes through a hand-held tool<sup>15</sup>. Second, by learning the statistical distribution of possible states of the world — that is, different possible configurations or scenarios, termed the prior within Bayesian inference — the estimate can be further refined (for a review see REF. 16). Lastly, by combining these processes with internal models of the body that map the motor commands (as signalled through the efference copy) into the expected sensory inputs, Bayesian inference can be used to estimate the evolving state of our body and the world (for example, REF. 17). Such an estimator is termed a Kalman filter and aims to optimally estimate the state, given sensory feedback, efference copy and knowledge of the dynamics and properties of sensory and motor noise.

**Decisions and strategies.** Most motor tasks involve a sequence of decision-making processes that determine, based on information that is extracted during the unfolding task, when to make the next movement and which movement to make<sup>18,19</sup>. The skill of a footballer, for example, is not only determined by the precision with which he can pass the ball but also by the speed with which he can make the correct decision on where to pass it. Decision making that is related to movement selection has been extensively studied in the context of saccadic eye movements in which participants must decide whether to look left or right based on the motion of dots on a screen<sup>18</sup>. This task has been modelled by considering how the accumulated sensory evidence, termed the decision variable, supports one or the other alternative. This single decision variable represents the integral of the

noisy evidence over time. Choice accuracy and reaction time are then explained by a bounded drift-diffusion model in which the decision variable reaches a positive or negative bound. It has been generally assumed that once the bound has been reached, the decision-making process is terminated. However, because there are substantial delays in the sensorimotor system, at the point of movement initiation there is sensory information in the processing pipeline that was not used to initiate the decision but could still be used to revise a decision. A recent study of reach target selection showed that this sensory information is processed and can even lead to subjects changing their mind mid-movement, usually to correct an error but sometimes also to spoil a good start<sup>20</sup>.

As researchers who are interested in sensorimotor control have broadened the scope of the tasks under study, the distinctions among sensorimotor, perceptual and cognitive components of the task have become blurred. For example, some studies have investigated how the motor system performs reward optimization. In many explicit cognitive tasks, people make suboptimal judgments when faced with a set of decisions with uncertain outcomes<sup>21</sup>. By contrast, when confronted with the motor variants of the same tasks, people often exhibit close to optimal decisions. For example, when pointing to target configurations that have different reward and penalty regions, it has been shown that people are able to choose their average pointing location to minimize the loss that accrues through the variability of pointing<sup>22</sup>. However, in some motor tasks subjects can show a tendency to be risk averse or risk seeking; that is, accepting a lower average reward if the variance of the reward is decreased or increased, respectively<sup>23–25</sup>.

Although most studies of decisions and strategies have focused on an individual, more recent studies have examined two-person (dyadic) sensorimotor interactions (BOX 1).

**Classes of control.** In general, optimizing motor performance is achieved through three classes of control: predictive or feedforward control, which is critical given the feedback delays in the sensorimotor system; reactive control, which involves the use of sensory inputs to update ongoing motor commands; and biomechanical control, which involves modulating the compliance of the limb. In general, all three of these control processes are adaptable and can contribute to motor learning.

As a result of time delays associated with receptor transduction, neural conduction, central processing and muscle activation, skilled action often relies on predictive control. For example, when lifting an object, people scale lifting force applied by the fingertips in anticipation of object weight. Such predictions, which are essential for smooth and dexterous manipulation<sup>26</sup>, require a system that can effectively simulate the behaviour of our body and environment — the so-called forward model. Using a copy of the motor command (that is, efference copy), the forward model predicts the sensory consequences. These predictions may also support a number of other movement-related functions including state estimation, likelihood estimation, sensory filtering,

## Noise

Random or unpredictable fluctuations and disturbances of neural, neuromuscular or environmental origin.

## Bayesian inference

A method of statistical inference in which observations are used to calculate or update the probability distribution of hidden variables.

## Visuo-haptic integration

The process that combines visual information (for example, the visual size of an object) and haptic information (for example, the felt size of a grasped object) into a single percept (for example, its size).

## Efference copy

A copy of the outgoing (efferent) motor command that can be used in conjunction with a forward model to predict the sensory consequences of action.


## Dynamics

The relationship between force and motion that can be altered using robotic interfaces to study the learning of novel dynamics.

## Forward model

A neural simulator that predicts (in the causal — and hence, forward — direction) the sensory consequences of an action given the current state and efference copy of the motor command.

Box 1 | Multi-person sensorimotor learning



Although most decision-making tasks involve a single subject, sensorimotor research has started to examine interaction between people and has even shown that two people can be as good as, if not better, than one in joint motor tasks<sup>116,117</sup>. Closing the sensorimotor loop between two people also allows an examination of strategy in tasks that are either combative<sup>118</sup> or cooperative<sup>119</sup>. Such decision making is typically examined within the framework of game theory. For example, in classical prisoner's dilemma, two prisoners each choose to cooperate (claim the other person is innocent) or defect (claim the other person is guilty). If both cooperate, they each receive a short sentence (3 years), whereas if both defect they each receive a moderate sentence (7 years), and if one cooperates and the other defects, the defector is freed and the cooperator receives a lengthy sentence (10 years). The globally optimal solution in which the players benefit the most is for both players to cooperate. However, if one of the players decides to defect, the defector reduces their sentence at the expense of the other player. In such a non-cooperative setting, the Nash solution — which minimizes each player's maximum possible punishment — is for both players to defect. When people have to make decisions based on a set of rules such as these, they are typically sub-optimal.

Recently, such theoretic problems have been examined in the motor domain by translating years-in-prison into movement effort. In the motor version of this game, each player makes a reaching movement from a starting position to a target bar (see the figure, part a) and can choose any path between the bars. A robot simulates a stiff spring that resists the subjects' movements and the spring constant for each subject depends on the lateral position of both players. The lateral position of each subject's hand is mapped from fully cooperate at one extreme to fully defect at the other. The spring constant for each player matches the typical prisoner's dilemma payoff at the extremes (see the figure, part b) and is linearly interpolated between these extremes. When a single subject controls both robots with their two arms (not shown) the dominant strategy is cooperative (see the figure, part c), whereas in the two-player game the Nash solution becomes dominant. Therefore, in contrast to the cognitive version of the game, in such two-player motor games, subjects rapidly develop near-optimal game-theoretic solutions — that is, the players adopt Nash equilibrium solutions in which they choose actions so that neither has anything to gain by changing only his or her strategy<sup>118</sup>. However, there are other game-theoretic tasks in which there are coordination patterns that are beneficial to both players (for example, when walking down a corridor; as long as each person moves in a different direction they can easily pass each other). In such cooperative games, subjects are able to find close-to-optimal coordination patterns<sup>119</sup>. Data from REF. 118.

agency attribution and mental simulation<sup>27</sup>. Prediction is supported by learned correlations, or priors. For example, when lifting, people use information about material and size of an object to predict its weight.

Skilled sensorimotor behaviour also requires the prediction of the sensory consequences of motor commands. For example, when lifting objects, the sensorimotor system predicts the sensory events that are associated with object lift-off — for example, the tactile afferents. If a mismatch between predicted and actual sensory information is detected, the system can launch appropriate, task-protective corrective actions and can also update the knowledge of object weight to improve

future actions. Thus, through the prediction of sensory consequences, there is an intimate relationship between predictive and reactive control mechanisms<sup>28</sup>.

Fast reactive feedback loops, such as the monosynaptic stretch reflex, can rapidly drive motor responses but cannot be easily modified even by extended experience. By contrast, longer loop reflexes, which involve supraspinal mechanisms, can be modified in a task-dependent manner and may also be tuned through learning<sup>29</sup>. For example, in arm movements, both the sign and size of the early component of the long-latency response to a force pulse perturbation depend on whether the pulse moves the hand towards or away from the intended

target<sup>30,31</sup>. Such intelligent, context-specific responses are consistent with the theoretical framework of optimal feedback control (for a review see REF. 32), which suggests that the central nervous system sets up feedback controllers that continuously convert sensory inputs into motor outputs, and that these are optimally tuned to the goals of the task by trading off energy consumption with accuracy constraints. An elegant demonstration of such flexibility is when one hand is perturbed in a task that requires a bimanual reaching movement<sup>33</sup>. If each hand controls its own cursor, only the perturbed hand shows an appropriate response. However, when the two hands control a single cursor, located at the spatial average of the two hands, then the perturbation of one hand results in appropriate responses in both the perturbed hand and the other hand — that is, the corrective response is shared between the hands.

An important feature of the optimal feedback control model is the concept of minimum intervention<sup>34</sup>; that is, setting up feedback controllers that only correct for variation that is deleterious to the task goal. Corrections of task-irrelevant errors are not only wasteful but they can also generate task-relevant errors. The minimum intervention principle has now been demonstrated in a number of tasks including the seemingly simple task of generating a target force with the tip of the index finger<sup>35</sup>. The control of this task can be characterized within a seven-dimensional space representing the seven muscles that regulate index finger force. The variability in this space can be partitioned into a task-relevant component that modulates force in the target direction and a task-irrelevant component that does not. During this task, task-irrelevant variability is consistently larger than task-relevant variability, suggesting that at the muscle level there is a preferential control of task-relevant dimensions. An important question for future research is whether this reduction in variability arises through the feedback control process alone or also from pre-wired synergies that reflect the structure of the feedforward command.

A third form of control can be exerted by specifying the biomechanical properties of the body and tools with which we interact. For example, by varying the activations of the muscles in the arm it is possible to control the stiffness at the hand or the tip of a hand-held tool. Not only can people scale their overall hand stiffness, or compliance, but they can also shape the pattern of stiffness either by varying muscle activations or the posture of the arm<sup>36,37</sup>. By modulating stiffness, the motor system can exercise control over the immediate response to external perturbations. Although stiffness can be used to deal with some perturbations, it is limited in its flexibility and, because it often requires co-contraction of opposing muscles, it can be an effortful solution to maintaining stability.

Most action tasks involve a combination of all three control mechanisms, with the contribution of each depending on the nature of the task. In terms of the interaction between stiffness and predictive control, it has been proposed that when errors are large, stiffness increases but as the predictive component begins to learn, and errors

reduce, stiffness decreases<sup>38</sup>. Recently, it has been proposed that such impedance control can be brought within an optimal control framework by formulating impedance control as the optimal response in the face of uncertainty about the dynamics of the body and environment<sup>39</sup>.

### Processes of motor learning

The previous section discussed how learning can occur at different levels of the motor hierarchy. To understand how these changes are implemented in neurons, we need theories of the processes and representations<sup>40</sup> through which learning is achieved. The processes of motor learning can be distinguished by the type of information that the motor system uses as a learning signal. Although different sensory modalities, such as vision, proprioception and touch can all play an important part in motor learning, we focus here on the nature of the information, independent of modality, that is used during learning.

**Error-based learning.** When a movement is made, the sensorimotor system can sense the movement's outcome and compare this to the desired or predicted outcome. The information contained in such sensory prediction errors not only tells the system that it missed the goal but also specifies the particular way in which the target was missed. To be able to use this information, the nervous system needs to estimate the gradient of the error with respect to each component of the motor command — that is, whether the error will go up or down as a component is increased or decreased. For example, when a thrown dart lands to the right of its intended goal, the system can adjust the motor command for the next throw by changing the orientation of the upper body, by adjusting the movement of the wrist or arm, or through any combination of these components. If the system knew the true gradient, all of the components could be adjusted to reduce error. However, because the gradient can be only estimated with some noise, the same error can lead to very different adjustments, each of which may reduce the error.

Error-based learning is the driving force behind many well-studied adaptation paradigms, including prism adaptation<sup>41</sup>, saccade adaptation<sup>42</sup>, reaching in force fields<sup>43,44</sup>, visuomotor adaptation<sup>45</sup> and grip force adaptation<sup>46</sup>. It can also drive motor learning and movement corrections in the absence of external perturbations<sup>47</sup>. A common feature across these different task domains is that the system can — and will — learn from an error on a single trial. Thus, adaptation is observable even when all perturbations are random and the subject is told not to adapt<sup>48–50</sup>.


There is extensive evidence that fast trial-by-trial error-based learning relies on the cerebellum. Patients with cerebellar lesions show substantial impairment in fast adaptation across many task domains<sup>51–55</sup>. Although transcranial magnetic stimulation of neocortical regions does not modify the initial learning in adaptation tasks<sup>56–58</sup>, adaptation can be sped up through transcranial direct current stimulation of the cerebellar cortex<sup>59</sup>. However, given that the neocortex and the cerebellum form a closed loop, they must both be involved in error-based learning. How the cerebellum and cortex interact

**Optimal feedback control**  
Optimality that is applied to setting up time-varying feedback controllers to drive a movement so as to minimize a function that is typically a combination of accuracy and effort.

**Impedance control**  
Impedance refers to the force produced by the limb to resist an externally induced motion (or deviation from desired motion). Impedance control changes this biomechanical behaviour of the limb by changing the configuration or stiffness through muscular co-contraction.

**Force fields**  
A type of dynamic motor learning in which forces are applied to the hand by a robotic manipulandum and in which the force direction and magnitude depends on the state of the hand (for example, its position and velocity), allowing the perturbation to be plotted as a force field.





**Figure 1 | Error-based learning in redundant systems. a** | A simple redundant task in which a target has to be reached with a stick using two effectors; rotations of the upper and lower arms that together contribute to a combined outcome (see REF. 134). **b** | Many combinations of the two rotations (the result of motor commands) will on average produce the correct solution (the set of combinations that are on average correct is known as the solution manifold). The error signal indicates whether the last movement overshot or undershot the target. For error-based learning to occur, the system needs to assign the error to a combination of the two effectors — that is, follow the error-based learning gradient (shown by red arrows). However, to find a less variable or less effortful 'optimal' solution (shown by a green circle) along the solution manifold (direction shown by a dashed arrow) different learning mechanisms, such as reinforcement learning, are needed. Data from REF. 134.

during error-based learning and where different types of adaptation are stored, remains an open question.

In many learning situations, the gradients of the error with respect to changes in the motor command are familiar. For example, when reaching under a displacing prism, although the visually displayed hand is shifted relative to the true hand's position, a movement of the hand to the left still moves the viewed hand in the same direction (to the left). By contrast, there are tasks in which the sensory error at each point in time is known but the gradient back to motor commands has to be learned<sup>6,60,61</sup>. For example, in a recent study the 19 joint angles in the human hand were mapped onto the motion of a two-dimensional cursor on a screen<sup>60,62</sup>. Successful performance in rapidly moving the cursor between targets requires subjects to learn — initially through undirected search — the mapping between hand configurations and cursor motion as well as the gradient relating cursor errors to hand configuration changes.

In summary, during error-based learning the system exploits a directional (signed) error signal and follows an internal estimate of the gradient in this direction. Thus, error-based learning serves to keep behaviours well calibrated and corrects for any systematic biases.

**Reinforcement learning.** Error-based learning can reduce the average error to zero, but once this is achieved it does not provide a mechanism to systematically improve performance further. For example, consider the simple task of using two degrees of freedom, rotations of the upper and lower arm, to touch a target with a hand-held stick (FIG. 1a). Such a task is redundant in that many combinations of two rotations can achieve the goal, leading to a solution manifold (FIG. 1b). Error-based learning

provides a gradient to move the system onto the line, but once the error is nulled on average, error-based learning cannot improve the solution further. To achieve a reduction in, for example, the variability of the errors, other learning mechanisms are needed to move the system to the optimal location on the solution manifold (FIG. 1b). A reduction in the variability for a given movement speed can be considered the hallmark of skill learning<sup>63</sup>. One possible candidate signal that could drive such learning is information about the relative success and failure of the movement. In contrast to a signed error signal, reinforcement signals such as success or failure are inherently unsigned, and therefore do not give information about the direction of required behavioural change<sup>64</sup>. Thus, the motor system needs to explore different possibilities to gradually improve its motor commands. Like error-based learning, reinforcement learning can also be used to guide learning towards the solution manifold, but as the signal (the reward) provides less information than in error-based learning (the vector of errors) such learning tends to be slow.

In situations in which a complex sequence of actions needs to take place to achieve a goal and the outcome or reward is far removed from the action (for example, learning the movements required to make a playground swing go higher), error-based learning cannot easily be applied and reinforcement learning techniques can be used to assign credit or blame, back in time, to actions that led to success or failure. A classic example of such a reinforcement learning task is to balance a rigid pendulum that is attached, to pivot freely on the side of a cart. Forces can be applied to the cart with the aim of swinging the pendulum from a vertically downwards to an upwards configuration and maintaining it in a vertical position in the presence of noise on motor and sensory signals. Simply providing an error at each point in time, such as the angular deviation from vertical, does not lead to the task being solved because the pendulum often has to go the wrong way initially to build up enough momentum to reach the vertical position. There have been some exciting advances in machine learning techniques for such reinforcement learning problems that achieve learning speeds that approach, if not exceed, human motor learning. For example, fast learning in the cart and pendulum task can be achieved by learning a probabilistic model of the input–output relationship so that it is possible to predict not only the expected response of the system to an action but also the confidence in all possible outcomes<sup>65</sup>. Although reinforcement learning in human motor tasks has not yet received much attention, recent evidence indicates that reaching movements and saccades can be changed through explicit rewards<sup>66,67</sup>. For skill learning, reward signals may also promote the subsequent consolidation of a motor behaviour<sup>68,69</sup>. However, in general we have yet to develop a full understanding of what constitutes a rewarding signal for the motor system.

A possible neural correlate of reward signals are dopaminergic projections from the ventral tegmental area to the primary motor cortex<sup>70</sup>. Disruption of these projections leads to profound deficits in the acquisition

#### Solution manifold

The set of solutions that can each, on average (perhaps owing to noise), solve a task.

Box 2 | **Observational learning**

An important source of information in the development of motor skills is the observation of others. Since the discovery of mirror neurons that fire when performing and observing action tasks<sup>120</sup>, many studies have provided evidence that watching another person perform an action engages sensorimotor representations of the observed action<sup>121</sup>. This coupling can even be seen at the behavioural level; when watching an actor manipulate objects, the observer's gaze behaviour closely corresponds to that of the actor<sup>122,123</sup>. It is well established that people can learn high-level information about what movements to make, and in what sequence, by observing actions (for example, REF. 124). However, more recent work has shown that people can also learn how to compensate for movement perturbations through action observation. Specifically, watching a video of arm movements that are initially perturbed and then slowly adapt, facilitates adaptation when the observer subsequently reaches with the same load<sup>125</sup>. This type of observational learning may involve learning from prediction errors. Specifically, the observer, like the actor, may generate predictions about movement outcomes, compare these predictions to actual outcomes and use the error to update an internal model of the perturbing load. One possible way in which the observer could generate such predictions would be to covertly simulate the motor commands of the observed action and pass these commands (similar to an efference copy) through a forward model of the controlled system<sup>126,127</sup>. Our ability to interpret the actions of others can depend on our skill set. For example, professional basketball players are more accurate at judging the success of an observed basket shot<sup>128</sup>. However, the question of whether simulation is used to predict movement outcomes remains a matter of debate<sup>128–130</sup>.

The neural correlates of motor learning through action observation have recently been examined. It has been shown that a network that is engaged in processing self-generated reach errors<sup>49</sup> also seems to be involved in observing reach errors<sup>131</sup>. Specifically, when watching a video of an actor reaching with a novel hand-held load, activity in left intraparietal sulcus, left dorsal premotor cortex and right cerebellar cortex was modulated by the amplitude of observed kinematic errors. It has been suggested that observational learning involves processing two forms of prediction errors: errors that are associated with predicting the action chosen by the observed actor and errors that are associated with predicting the outcome of the action<sup>132</sup>. It was found that activity in the dorsolateral and the ventromedial prefrontal cortices, respectively, corresponded to these two learning signals.

of forelimb reaching movements in the rat<sup>71</sup>. Although reinforcement learning must play an important part in the acquisition of skillful movements, its mechanisms are poorly understood compared with error-based learning.

**Use-dependent learning.** Use-dependent learning refers to the phenomenon that the state of the motor system can change through the pure repetition of movements, even if no outcome information is available. For example, the repeated execution of thumb abduction movements biases the direction of the movements elicited by transcranial magnetic stimulation over the thumb area of motor cortex<sup>72</sup>. This learning mechanism can be disrupted through GABA agonists<sup>73</sup>.

A recent study has shown that the repetition of a reaching movement to a target reduces the variability of such movements, and induces a bias towards this trained direction when reaching to neighbouring targets<sup>74</sup>. Similar influences of use-dependent learning towards the last executed movement solutions were found in a study that used visuomotor rotations to influence movement directions<sup>69</sup>. A third study showed that use-dependent learning can occur in parallel with error-based learning<sup>75</sup>. Initially, participants made reaching movements to a horizontal target region located in front of them and tended to reach straight ahead to the

centre of the target. Reaching movements were then constrained by a robot to move in a straight line to a location on the right side of the target. Subjects generated leftward forces to resist the robot to no avail. When the robot was subsequently turned off, participants showed an after-effect — they initially reached to the left of centre. After a few reaches, this error-based after-effect wore off and reaching became biased towards the right until, eventually, reaches became centred again. The bias to the right was interpreted as an effect of use-dependent learning associated with the constrained reaches. These results suggest that error-based and use-dependent learning can occur in parallel and that, at least in this task, use-dependent learning wears off more slowly than error-based learning. An important question is whether such use-dependent learning is modulated by rewards, and if so, how.

The three forms of learning that are described above usually involve a subject interacting with the environment so as to experience their own errors, successes or failures, whereas recent work has examined the contribution of observing others to motor learning (BOX 2).

**Representations in motor learning**

Sensorimotor learning involves learning new mappings between motor and sensory variables. Such transformations are termed internal models, as they represent features of the body or the environment, such as the way in which a hand-held racquet responds to force and torques, or the way in which prism glasses change the visuo-motor alignment. Numerous factors can change these mappings, such as muscle fatigue or a change in object weight, and successful performance requires adaptation to these factors.

The information that is obtained during a single movement, however, is often too sparse or too noisy to unambiguously determine the source of the error. Therefore, the information does not adequately specify the way in which the motor commands should be updated, leading to the so-called inverse problem<sup>76</sup>. To resolve this issue, the system does not start from a blank slate. Instead, it uses representations that reflect the internal assumptions about the task structure and that constrain the way in which the system is updated in response to errors. Such representations can be conceptualized in two ways, either as a mechanistic or a normative model.

Mechanistic models specify the representations and learning algorithms directly. In this framework, representations are often considered to be based on motor primitives (the neural building blocks out of which new motor memories are formed). Normative models suggest that the nervous system optimally adapts when faced with an error. To determine this optimal adaptation, the normative model must specify two key features of the world. First, how different factors, such as tools or levels of fatigue, influence the motor system — the so-called generative model. Second, how these factors are likely to vary over both space and time — that is the prior distribution. The structure of the generative model and the prior distribution together determine

**After-effect**

The deviations of a system from pre-perturbation behaviour after learning when the perturbation is first removed.

how the motor system should attribute an error to the underlying causes and, therefore, how it should adapt. Next, we describe how both of these viewpoints explain how errors are credited both spatially and temporally, and how experience may change these representations of task structure.

**Motor primitives.** An important idea about the way in which internal models are represented is the idea of motor primitives. Motor primitives can be thought of as neural control modules that can be flexibly combined to generate a large repertoire of behaviours. For example, a primitive might represent the temporal profile of a particular muscle activity. The overall motor output will be the sum of all primitives, weighted by the level of the activation of each module<sup>77</sup>. The makeup of the population of such primitives then determines which structural constraints are imposed on learning. For example, a behaviour for which the motor system has many primitives will be easy to learn, whereas a behaviour that cannot be approximated by any set of primitives would be impossible to learn. This principle is elegantly demonstrated by a study in which participants' arm movements were perturbed with forces that depended either on the velocity or on the position of the hand<sup>78</sup>. Confronted with such perturbations, participants produced compensatory forces that included both a velocity- and a position-dependent component in the same direction. This pattern of adaptation could be modelled using motor primitives. For this task, the relevant elements span the different combinations of velocity- and position-dependent forces (FIG. 2a). The authors assumed that the motor system has a biased distribution of such primitives that favours perturbations in which the velocity- and position-dependent components point in the same direction (FIG. 2b). This intrinsic correlation may reflect the tuning of peripheral sensors such as muscle spindles<sup>78</sup> or more central mechanisms that capture the statistics of natural loads. The model successfully explains the observed biases during the learning of pure velocity- or position-dependent force fields and also predicts which combinations of forces are easy and which are hard to learn (FIG. 2b).

Motor primitives also determine the way in which learning generalizes. Current models of adaptation seek to explain how an error on one trial changes the behaviour on the next, thereby modelling the learning process. This approach is based on state–space models in which, typically, the internal state represents the estimate of the perturbation<sup>44,48,79–81</sup>. The estimate is updated after each trial based on the error in the previous trial. To account for generalization of learning — for example, across different movement directions — each direction is associated with its own state representing the estimate of the perturbation for a movement in that direction. A generalization function specifies how an error that is experienced during a movement in one direction affects the states that are associated with other directions. Such generalization functions tend to be relatively narrow Gaussian-like functions so that learning in one direction has a decaying effect for movements that diverge from

this direction. The shape of the generalization can be used as an indicator of the underlying motor primitives that form the motor memory.

**Credit assignment.** According to normative models, the process of motor learning can be understood as a credit assignment problem: the question of how to attribute an error signal to the underlying causes. For example, if a tennis player starts hitting shots into the net on the serve, the problem could be that the ball was not thrown high enough, was hit too early, that the racquet strings are loose or that he or she is fatigued. If the racquet dynamics have changed, the player would do well to learn these dynamics and remember them for the next time that they use this particular racquet. Conversely, if the player is simply tired, the necessary adjustments should only be temporary but should be applied even if the racquet is changed at that moment. Recent work has examined such credit assignment in redundant systems (BOX 3). Two types of credit assignment can be distinguished: contextual and temporal credit assignment.

Many studies have shown that following adaptation of a reaching movement to loads that are applied to the hand through the handle of a robotic interface, large after-effects are seen when the load is suddenly turned off (for example, REF. 43). However, these after-effects are substantially reduced, but not eliminated, when participants release the robotic handle (hence removing the load) before making a movement<sup>82–85</sup>. This result suggests that most of the perturbation was credited to the context (the grasped handle, which is external to the body), with a smaller portion credited to other sources, including internal sources such as changes in the dynamics of the arm. The proportion of the perturbation that is attributed to the robotic handle reduces when forces are slowly introduced over a number of trials so that participants are not aware of the perturbation<sup>84</sup>. These results have been formalized in a Bayesian model in which the relative allocation of credit to internal or external sources depends on both a prior shaped by previous experience as well as on the source that is most consistent with the perturbation<sup>86</sup>. A Bayesian model has also been used to account for the assignment of errors to either sensory or motor sources during motor learning<sup>87</sup>. This model confirmed a surprising prediction, that adaptation of reaching movement to perturbing forces elicits a change in the mapping between vision and proprioception (for example, REFS 88,89).

When we learn new dynamics or kinematics, we must also be able to link this learning to appropriate contextual cues such as objects, tasks or environments. As mentioned above, grasping and releasing an object (for example, a robot handle) can provide a powerful contextual cue for switching the internal model of the object dynamics<sup>90</sup>. Without such cues, it is difficult to learn novel opposing loads when these are applied to the hand<sup>91,92</sup>. Although some arbitrary contextual cues, such as colour, allow some slow learning of the two tasks<sup>93</sup>, behaviourally relevant contexts can lead to rapid context-dependent learning<sup>85</sup>. For example, when opposing loads are linked to bimanual and unimanual contexts, by

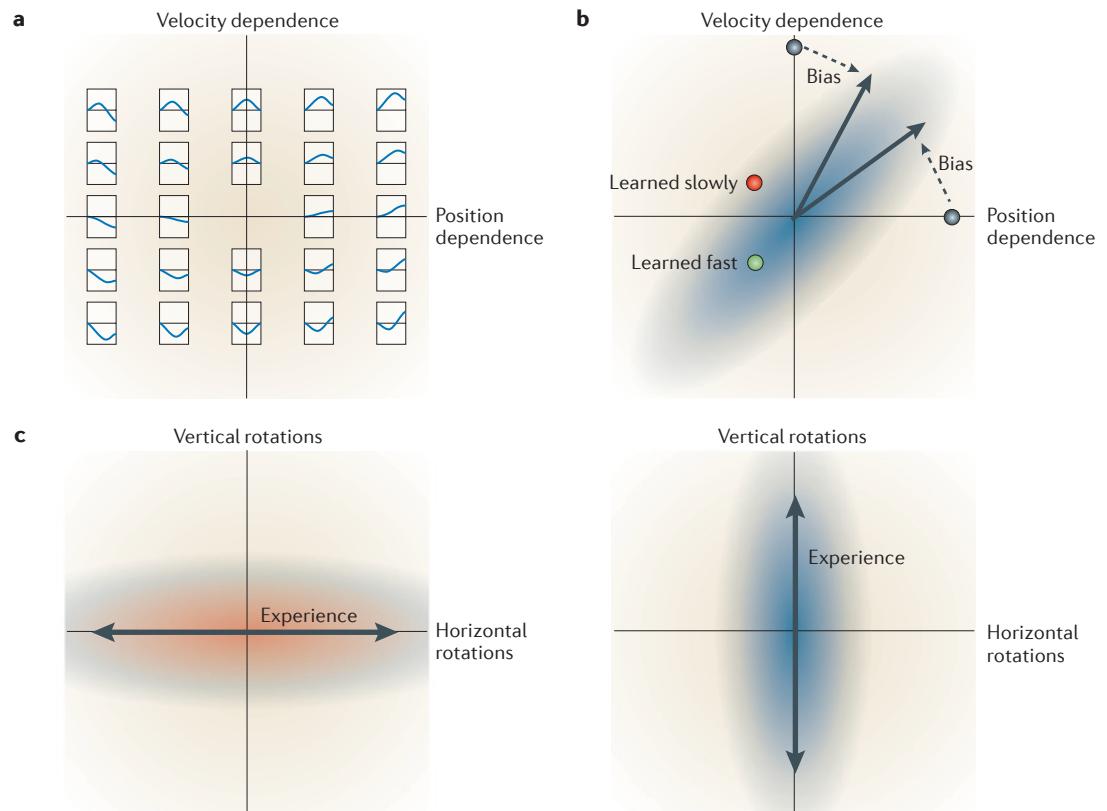
#### Kinematics

This refers to the relationship between positional variables, such as joint angles and hand position.

having participants simultaneously move the other arm for one of the loads, the interference is reduced<sup>94</sup>. Even within bimanual movements, the relative motion directions of the two hands act as a strong contextual cue to learning<sup>95</sup>. Similarly, the motor system can learn separate dynamics for movements in which a single object is jointly grasped with both hands, and movements in which each hand moves independently<sup>96</sup>. Several recent models have been developed to account for the reduction in interference in the presence of contextual cues<sup>80,97</sup>. These models propose multiple, overlapping internal representations that can be selectively engaged by each movement context. Although the field is starting to compile a list of the cues that can or cannot serve as contextual switches, a general underlying principle of what determines this striking divide is still elusive.

Errors during motor learning can also be assigned to processes that act on different timescales. A recent

model<sup>79</sup> based on this idea of temporal credit assignment can account for a number of phenomena that are related to reach movement adaptation to perturbations as well as saccadic adaptation<sup>98</sup>. The key idea is that there are fast and slow learning processes acting in parallel (FIG. 3a). Each process receives the same error and incorporates a proportion of the error into its current estimate of the underlying perturbation. In addition, each state decays passively over trials. The fast and slow processes differ by their rates of learning and unlearning (or decay). This dual-rate model is able to account for several features of motor learning, such as the double exponential rate of learning, savings in relearning and the spontaneous recovery of previously learned memories (FIG. 3b). Using techniques in which errors are removed so that the fast process decays quickly, it is possible to identify the state of the slow process and thereby establish evidence for the dual-rate model<sup>79</sup>.




**Figure 2 | Motor primitives and structural learning. a** | The motor system may have primitives for the dependence of a force magnitude through the course of a reaching movement. Each primitive represents a time course of force production for the duration of the movement. The final output is the sum of the primitives weighted by their activation. Possible primitives are either purely position- or velocity-dependent (primitives on x and y axes, respectively) or represent a combination of the two force components (off axis). **b** | The motor system may have a prior that favours perturbations that combine position- and velocity-dependent forces in the same direction. This prior can be represented as a non-isotropic distribution of motor primitives with more primitives on the positive diagonal. The prior leads to faster learning of perturbations that lie along the preferred direction (for example, the green disk) compared to perturbations that lie off the diagonal (for example, the red disk). The prior also leads to a bias of learning towards the main axis of the distribution when pure position- or velocity-dependent perturbations (shown by black circles) are experienced<sup>78</sup>. **c** | Structural learning can be achieved by changing the prior distribution of primitives through experience. For example, in a visuomotor rotation learning experiment, two groups of participants were either exposed to random horizontal perturbations (left part; shown by the double-ended arrow) or to random vertical perturbations (right part; shown by the double-ended arrow). After experience, adaptation to the matching perturbation type was accelerated, suggesting that the primitives became aligned with the axis<sup>107</sup>.

**Savings**

This refers to the phenomenon that relearning of a perturbation or skill for a second time is faster than initial learning.



## Box 3 | Credit assignment in redundant systems



A particular form of the credit assignment problem occurs in redundant systems, in which the actions of multiple effectors contribute to the outcome. In this situation there is a fundamental ambiguity as to which particular body part or muscle caused the error. One particular assignment has been studied in a bimanual task (see the figure, part a) in which subjects control a visual cursor that is located at the spatial average of the positions of the two hands<sup>133</sup>. When a visual perturbation such as a rotation is applied to the cursor path (see the figure, part a; black path) the error can be credited to either the right or left hand, or any combination of the two hands. The paths taken by the left and right hands (see the figure, part a; blue and red paths, respectively) can be used to quantify the extent of correction taken up by the two hands. Examination of the correction asymmetry, which reflects the proportion of the correction taken up by the right hand, across subjects (see the figure, part b) showed that right-handed participants corrected errors more with their left hands (correction asymmetric  $<0.5$ ), presumably because the left hand is more likely to make errors. Across participants, the hand that was most involved in the corrective response within a movement also showed the biggest adaptive change in the next movement (adaptation symmetry reflects the proportion of the total adaptation for which the right hand is responsible), suggesting that participants assigned errors in a unified manner for correction and adaptation. In a similar manner to structural learning, this credit assignment can be modulated by pre-exposing one of the hands to a high level of random perturbations during unimanual movements. This causes this hand to exhibit a greater corrective response when exposed to the bimanual redundant situation. Data from REF. 133.

Recently, models with additional learning processes have been proposed. For example, one model proposes a single fast process combined with multiple slow processes that are tuned for different learning contexts<sup>80</sup>. One advantage of having multi-rate learning processes is that these processes can parallel the temporal variations in the causes of sensorimotor errors. For example, some disturbances — such as muscle fatigue — arise and disappear on a short timescale, whereas others — such as muscle damage — tend to be long-lasting. The credit assignment problem can then be solved using Bayesian inference on a generative model that contains processes acting on these timescales<sup>99</sup>. In agreement with such a model, learning has been shown to be sensitive to the temporal statistics of the perturbation, with experience in a rapidly changing environment leading to motor memories that decay faster than motor memories resulting from experience in an environment with more gradual changes<sup>100</sup>. However, we have yet to determine whether the different timescales in motor learning are implemented by a system that contains a distribution of possible timescales, or by a finite set of qualitatively distinguishable neural and behavioural systems.

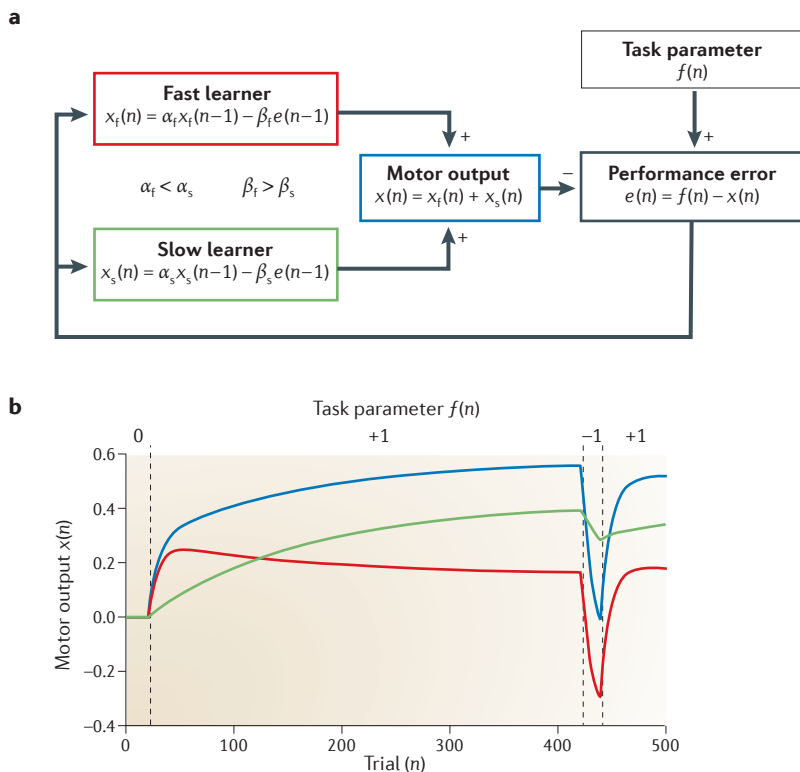
**Declarative memory**  
Memories that can be consciously recalled, such as facts and events.

**Procedural memory**  
Unconscious memories of skills and how to do things, such as being able to walk downstairs.

It has been suggested that the initial part of motor learning is more cognitively driven than later learning, and recent studies have started to explore this issue in more detail. For example, it has been shown that spatial working memory is particularly important for the early but not for the late stage of visuomotor learning<sup>101,102</sup>. In addition, a recent study suggests that the fast but not the slow learning process may share resources with the declarative memory system, leading to the fast process being disrupted by a task that engages declarative memory<sup>103</sup>. The fast and slow processes may be related to the declarative and procedural memory components that are observed in sequence learning tasks such as the serial reaction time (SRT) task. In the SRT task, disrupting the declarative component immediately following learning leads to slow off-line skill improvement or consolidation<sup>104</sup>. This suggests that the declarative component interacts with, and in this case hinders, the procedural component.

**Structural learning.** We have discussed how representations (motor primitives or priors underlying credit assignment) change how we learn the parameters of a task. However, in general there are at least three levels of representation that are relevant: the structure of the task, its parameters and the relevant state. The structure represents the relevant inputs and outputs of the system and the functional form of the equations that relate them. For example, when we learn to play tennis or squash we have to identify the task-relevant inputs and outputs — such as arm motor commands and racquet head motion — as well as the mapping between them, which depends on the geometry and dynamics of the racquet. Learning the structure in one task, such as tennis, can be beneficial for tasks that share a similar structure, such as other racquet sports. What differs between tennis, squash and badminton are the parameters of the structures, such as the racquet length, head size and weight. In principle, if the structure is known then the parameters of the system can be quickly identified, allowing rapid learning. However, if the structure is unknown, or one's estimate of the structure is incorrect, learning will be impeded. Finally, estimating task-relevant state information, such as the racquet head's position and orientation is facilitated if the structure and parameters of the task are estimated correctly, but can be impaired or biased otherwise.

An example from the robotics field illustrates one way in which structural learning can take place<sup>105</sup>. A spider-like robot knew only that it had a set of eight actuators, each of which controlled a single joint angle in its body, but did not know the structure of its own body; that is, how many legs it had, how many segments each leg had or the geometry of its body and legs. Its only sensory input was a tilt sensor mounted on the body. The robot took random actions and considered possible kinematic structures, honing in on those that could best explain the joint-angles and tilt dataset. It then chose new actions that aimed to distinguish between possible structures. The robot was quickly able to learn its own structure.



**Figure 3 | A dual-rate learning process.** **a** | Both the state of the fast learner ( $x_f$ ) and slow learner ( $x_s$ ) decay over trials and incorporate part of the performance error into their state. The fast learner learns more rapidly (indicated by a higher learning rate,  $\beta_f > \beta_s$ ) and also ‘unlearns’ more rapidly (indicated by a lower retention factor,  $\alpha_f < \alpha_s$ ), compared to the slow learner. The motor output is the combination of both the fast and slow learners’ outputs. The performance error is the difference between this output and the task parameter setting,  $f(n)$  (for example, the magnitude of a force field). **b** | Simulations of learning in the dual rate model. When the task parameter is set to +1 the total motor output (shown in blue) increases to reduce the error. Initially, the fast learner (shown in red) contributes most to the learning but with time the slow learner (shown in green) takes over and the fast learner’s contribution decays. This leads to the typical double exponential learning curve that is seen in many learning studies. When the task parameter is set briefly to -1, the total motor output quickly falls to zero as the fast learner adapts. However, although the total motor output now has the same value as before learning, the states of the fast and slow learners are not back at baseline. When the original task parameter setting of +1 is reinstated, the overall learning is faster than the initial learning (compare the rise in the blue curves on the first and second occurrence of task parameter +1) as the slow learner has retained much of their learning, thereby demonstrating savings (for more details of this model see REF. 79).

Recent studies have shown that new structures can be learned by exposing participants to a randomly varying set of tasks that share a common structure but vary in their parameter settings<sup>106–109</sup>. For example, after being exposed to horizontal (FIG. 2c) or vertical (FIG. 2d) visuomotor rotations in a three-dimensional reaching task, participants adapted more rapidly to new tasks that share the same structure. In this case, the new perturbation fell within their prior assumption about the distribution of possible perturbations, and they therefore only needed to adjust the key parameters that were appropriate for that structure. Whether we make optimal actions to facilitate the learning of new structures, and the identification of previously learned structures and parameters, is an open question<sup>110</sup>.

Many studies of motor learning have applied loads to the hand that have novel and unusual structures. One common example is the use of rotary viscous force fields that perturb the arm perpendicular to the direction of the movement with a magnitude that scales with hand speed (for example, REFS 91, 111). In these situations, learning tends to be slow because the structure is new and because it is difficult to separate the structure from the parameters. Learning may only become faster if the learner experiences more than a single example of the structure with different parameters; for example, racquets with different lengths or rotary viscous loads with different viscosities. In some cases, the structure of the task is so novel and complex that an initial exploratory period is observed during which subjects must discover the structure of the task before they can apply error-based learning and show clear improvements in performance<sup>6,112</sup>. By contrast, most of the learning that we do in everyday motor tasks involves rapid learning of the parameters of familiar structures. For example, most of the new objects that we pick up are inertial loads and we must simply learn the mass. Cues combined with priors can be used to identify the structure (for example, familiar inertial or elastic loads) and we can then quickly estimate the parameters (for example, mass or elasticity) of the task<sup>28,113,114</sup>. It remains unclear whether the learning of structures and parameters relies on fundamentally different neural operations or whether they share the same neural machinery.

**Conclusions and future directions**

We have reviewed three key topics that are related to sensorimotor learning: what has to be learned, how it is learned and how knowledge developed during learning is represented. In reviewing these topics, we have tried to identify a number of specific research questions that remain unanswered. Here, we highlight three overarching challenges for the future of the field.

*From laboratory learning to real-world learning.* We now have a detailed understanding of the learning and control of a narrow range of tasks, including simple reaching tasks in which visuomotor and dynamics perturbations are applied. Although these tasks are amenable to analysis and modelling, they do not capture the full complexity of real-world motor control and it is not clear whether the learning models that are developed will generalize to tasks such as tying shoelaces or learning to skateboard. The study of sensorimotor control is fundamentally difficult because it deals with a dynamic, real-time control system that turns sensations and memory into action and vice versa. Given this complexity, it is understandable that the field has focused on a limited number of simplified tasks. However, expanding the range of tasks may help us deal with new challenges.

*From sensorimotor control to robotics and brain-machine interfaces.* Progress in sensorimotor control research is reflected in the successful implementation of learning and control models in robotic devices. Although robots have faster and less noisy sensors and actuators, and can perform rapid decision making and control, in general,

current robots lag far behind human performance. So far, relatively few principles from the study of biological sensorimotor control have been influential in the robotics community. However, one notable example is the idea of incorporating adaptive passive properties into robot actuators (for example, adjustable springs) to enable the sort of impedance control of limb motion that humans exploit.

There has been substantial interest in trying to extract information from signals that are recorded from the brain to control external devices. The goal is to develop devices that will allow patients with neural impairments, including spinal cord injury, as well as amputees to effect movement. At present, such systems do not fully close the sensorimotor loop; although the patient can see the robotic interface and therefore guide it visually, effective tactile feedback, which may allow finer manipulation abilities, has yet to be developed.

**From models to neuronal implementations.** Although substantial progress has been made in computational sensorimotor control, the field has been less successful in linking computational models to neurobiological models of control. For example, despite the plethora

of experimental studies and extensive computational modelling, the fundamental computations performed by — for example — the motor cortex or the cerebellum remain elusive. The hope is that as models of sensorimotor control and learning become more precise, the chances of linking these models to the underlying neural implementations will increase. For example, model-based imaging studies (for example, REF. 49) and new theoretical concepts in neurophysiology (for example, REF. 115) are beginning to link models and neural data. Studying how these computations are implemented in the nervous system will also provide constraints for computational theories.

In conclusion, over the past 10 years, tremendous progress has been made in our understanding of the computational aspects of motor learning. The exciting challenges ahead are to understand the learning of real-world tasks and the neural implementation of the underlying processes. The success of this research field will be measured by whether the theories can inform behaviour training regimes in sports and rehabilitation, inspire new developments in the field of robots and contribute to new treatments for neurological movement disorders.

1. Itti, L. & Koch, C. Computational modelling of visual attention. *Nature Rev. Neurosci.* **2**, 194–203 (2001).
2. Hayhoe, M. & Ballard, D. Eye movements in natural behavior. *Trends Cogn. Sci.* **9**, 188–194 (2005).
3. Najemnik, J. & Geisler, W. S. Optimal eye movement strategies in visual search. *Nature* **434**, 387–391 (2005).  
**The authors derive an ideal Bayesian observer to select the eye movement that gains the most information about target location in a cluttered environment. By examining humans in this task they show that they are nearly optimal compared to this model, suggesting that eye movements are chosen to maximize information about target location.**
4. Land, M. & Tatler, B. *Looking and Acting: Vision and Eye Movements in Natural Behaviour* (Oxford Univ. Press, New York, 2009).
5. Land, M. & McLeod, P. From eye movements to actions: how batsmen hit the ball. *Nature Neurosci.* **3**, 1340–1345 (2000).
6. Sailer, U., Flanagan, J. R. & Johansson, R. S. Eye-hand coordination during learning of a novel visuomotor task. *J. Neurosci.* **25**, 8833–8842 (2005).
7. Seki, K., Perlmutter, S. & Fetz, E. Sensory input to primate spinal cord is presynaptically inhibited during voluntary movement. *Nature Neurosci.* **6**, 1309–1316 (2003).
8. Bays, P. M. & Wolpert, D. M. Computational principles of sensorimotor control that minimize uncertainty and variability. *J. Physiol.* **578**, 387–396 (2007).
9. Most, S. B., Scholl, B. J., Clifford, E. R. & Simons, D. J. What you see is what you set: sustained inattentional blindness and the capture of awareness. *Psychol. Rev.* **112**, 217–242 (2005).
10. Triesch, J., Ballard, D. H., Hayhoe, M. M. & Sullivan, B. T. What you see is what you need. *J. Vis.* **3**, 86–94 (2003).
11. Green, C. S. & Bavelier, D. Action video game modifies visual selective attention. *Nature* **423**, 534–537 (2003).
12. Nasir, S. M. & Ostry, D. J. Auditory plasticity and speech motor learning. *Proc. Natl Acad. Sci. USA* **106**, 20470–20475 (2009).
13. Faisal, A. A., Selen, L. P. J. & Wolpert, D. M. Noise in the nervous system. *Nature Rev. Neurosci.* **9**, 292–303 (2008).
14. Ernst, M. & Bulthoff, H. Merging the senses into a robust percept. *Trends Cogn. Sci.* **8**, 162–169 (2004).
15. Takahashi, C., Diedrichsen, J. & Watt, S. J. Integration of vision and haptics during tool use. *J. Vis.* **9**, 3 (2009).
16. Kording, K. & Wolpert, D. M. Bayesian decision theory in sensorimotor control. *Trends Cogn. Sci.* **10**, 319–326 (2006).
17. Vaziri, S., Diedrichsen, J. & Shadmehr, R. Why does the brain predict sensory consequences of oculomotor commands? Optimal integration of the predicted and the actual sensory feedback. *J. Neurosci.* **26**, 4188–4197 (2006).
18. Gold, J. I. & Shadlen, M. N. The neural basis of decision making. *Annu. Rev. Neurosci.* **30**, 535–574 (2007).
19. Cisek, P. & Kalaska, J. F. Neural mechanisms for interacting with a world full of action choices. *Annu. Rev. Neurosci.* **33**, 269–298 (2010).
20. Resulaj, A., Kiani, R., Wolpert, D. M. & Shadlen, M. N. Changes of mind in decision-making. *Nature* **461**, 263–266 (2009).
21. Kahneman, D. & Tversky, A. *Choices, Values, and Frames* (Cambridge Univ. Press, Cambridge, UK, 2000).
22. Trommershäuser, J., Maloney, L. T. & Landy, M. S. Decision making, movement planning and statistical decision theory. *Trends Cogn. Sci.* **12**, 291–297 (2008).
23. Nagengast, A. J., Braun, D. A. & Wolpert, D. M. Risk-sensitive optimal feedback control accounts for sensorimotor behavior under uncertainty. *PLoS Comput. Biol.* **6**, e1000857 (2010).
24. Braun, D. A., Nagengast, A. J. & Wolpert, D. M. Risk-sensitivity in sensorimotor control. *Front. Hum. Neurosci.* **5**, 1 (2011).
25. Nagengast, A. J., Braun, D. A. & Wolpert, D. M. Risk-sensitivity and the mean-variance trade-off: decision making in sensorimotor control. *Proc. Biol. Sci.* **278**, 2325–2332 (2011).
26. Johansson, R. S. & Flanagan, J. R. Coding and use of tactile signals from the fingertips in object manipulation tasks. *Nature Rev. Neurosci.* **10**, 345–359 (2009).
27. Wolpert, D. M. & Flanagan, J. R. Motor prediction. *Curr. Biol.* **11**, R729–R732 (2001).
28. Flanagan, J. R., Bowman, M. C. & Johansson, R. S. Control strategies in object manipulation tasks. *Curr. Opin. Neurobiol.* **16**, 650–659 (2006).
29. Wagner, M. J. & Smith, M. A. Shared internal models for feedforward and feedback control. *J. Neurosci.* **28**, 10663–10673 (2008).
30. Pruszynski, J. A., Kurtzer, I. & Scott, S. H. Rapid motor responses are appropriately tuned to the metrics of a visuospacial task. *J. Neurophysiol.* **100**, 224–238 (2008).
31. Pruszynski, J. A., Kurtzer, I., Lillicrap, T. P. & Scott, S. H. Temporal evolution of “automatic gain-scaling”. *J. Neurophysiol.* **102**, 992–1003 (2009).
32. Todorov, E. Optimality principles in sensorimotor control. *Nature Neurosci.* **7**, 907–915 (2004).
33. Diedrichsen, J. Optimal task-dependent changes of bimanual feedback control and adaptation. *Curr. Biol.* **17**, 1675–1679 (2007).
34. Todorov, E. & Jordan, M. I. Optimal feedback control as a theory of motor coordination. *Nature Neurosci.* **5**, 1226–1235 (2002).
35. Valero-Cuevas, F. J., Venkadesan, M. & Todorov, E. Structured variability of muscle activations supports the minimal intervention principle of motor control. *J. Neurophysiol.* **102**, 59–68 (2009).
36. Burdet, E., Osu, R., Franklin, D., Milner, T. & Kawato, M. The central nervous system stabilizes unstable dynamics by learning optimal impedance. *Nature* **414**, 446–449 (2001).
37. Franklin, D. W. *et al.* Endpoint stiffness of the arm is directionally tuned to instability in the environment. *J. Neurosci.* **27**, 7705–7716 (2007).
38. Franklin, D. W. *et al.* CNS learns stable, accurate, and efficient movements using a simple algorithm. *J. Neurosci.* **28**, 11165–11175 (2008).  
**A new model of motor learning in which stability, accuracy and efficiency are simultaneously optimized by specifying how feedforward commands to individual muscles are adjusted based on error. This model is the first to account for the temporal evolution of both net force and impedance control during learning.**
39. Mitrovic, D., Klanke, S., Osu, R., Kawato, M. & Vijayakumar, S. A computational model of limb impedance control based on principles of internal model uncertainty. *PLoS ONE* **5**, e13601 (2010).
40. Marr, D. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information* (Freeman, New York, 1982).
41. Martin, T., Keating, J., Goodkin, H., Bastian, A. & Thach, W. Throwing while looking through prisms. I. Focal olivocerebellar lesions impair adaptation. *Brain* **119**, 1183–1198 (1996).
42. Pelisson, D., Alahyane, N., Panouillères, M. & Tiliakete, C. Sensorimotor adaptation of saccadic eye movements. *Neurosci. Biobehav. Rev.* **34**, 1103–1120 (2010).
43. Shadmehr, R. & Mussa-Ivaldi, F. A. Adaptive representation of dynamics during learning of a motor task. *J. Neurosci.* **14**, 3208–3224 (1994).



44. Thoroughman, K. A. & Shadmehr, R. Learning of action through adaptive combination of motor primitives. *Nature* **407**, 742–747 (2000).
45. Krakauer, J., Pine, Z., Ghilardi, M. F. & Ghez, C. Learning of visuomotor transformations for vectorial planning of reaching trajectories. *J. Neurosci.* **20**, 8916–8924 (2000).
46. Flanagan, J. R. & Wing, A. M. The role of internal models in motion planning and control: evidence from grip force adjustments during movements of hand-held loads. *J. Neurosci.* **17**, 1519–1528 (1997).
47. van Beers, R. J. Motor learning is optimally tuned to the properties of motor noise. *Neuron* **63**, 406–417 (2009).
48. Donchin, O., Francis, J. T. & Shadmehr, R. Quantifying generalization from trial-by-trial behavior of adaptive systems that learn with basis functions: theory and experiments in human motor control. *J. Neurosci.* **23**, 9032–9045 (2003).
49. Diedrichsen, J., Hashambhoy, Y., Rane, T. & Shadmehr, R. Neural correlates of reach errors. *J. Neurosci.* **25**, 9919–9931 (2005).
50. Srimal, R., Diedrichsen, J., Rykkin, E. B. & Curtis, C. E. Obligatory adaptation of saccade gains. *J. Neurophysiol.* **99**, 1554–1558 (2008).
51. Diedrichsen, J., Verstynen, T., Lehman, S. L. & Ivry, R. B. Cerebellar involvement in anticipating the consequences of self-produced actions during bimanual movements. *J. Neurophysiol.* **93**, 801–812 (2005).
52. Smith, M. A. & Shadmehr, R. Intact ability to learn internal models of arm dynamics in Huntington's disease but not cerebellar degeneration. *J. Neurophysiol.* **93**, 2809–2821 (2005).
53. Morton, S. M. & Bastian, A. J. Cerebellar contributions to locomotor adaptations during splitbelt treadmill walking. *J. Neurosci.* **26**, 9107–9116 (2006).
54. Tseng, Y.-W., Diedrichsen, J., Krakauer, J. W., Shadmehr, R. & Bastian, M. S. Sensory prediction errors drive cerebellum-dependent adaptation of reaching. *J. Neurophysiol.* **98**, 54–62 (2007).
55. Golla, H. *et al.* Reduced saccadic resilience and impaired saccadic adaptation due to cerebellar disease. *Eur. J. Neurosci.* **27**, 132–144 (2008).
56. Baraduc, P., Lang, N., Rothwell, J. C. & Wolpert, D. M. Consolidation of dynamic motor learning is not disrupted by rTMS of primary motor cortex. *Curr. Biol.* **14**, 252–256 (2004).
57. Della-Maggiore, V., Malfait, N., Ostry, D. J. & Paus, T. Stimulation of the posterior parietal cortex interferes with arm trajectory adjustments during the learning of new dynamics. *J. Neurosci.* **24**, 9971–9976 (2004).
58. Hadipour-Niktarash, A., Lee, C. K., Desmond, J. E. & Shadmehr, R. Impairment of retention but not acquisition of a visuomotor skill through time-dependent disruption of primary motor cortex. *J. Neurosci.* **27**, 13413–13419 (2007).
59. Galea, J. M., Vazquez, A., Pasricha, N., Orban de Vivry, J.-J. & Celnik, P. Dissociating the roles of the cerebellum and motor cortex during adaptive learning: the motor cortex retains what the cerebellum learns. *Cereb. Cortex* **21**, 1761–1770 (2011). **This study showed that nodal transcranial direct current stimulation (tDCS) stimulation of the cerebellum increases the rate of error-based learning during adaptation of reaching movements under a visuomotor rotation. Anodal tDCS of primary motor cortex does not change the rate of learning, but makes the adaptation more resistant to washout.**
60. Mosier, K. M., Scheidt, R. A., Acosta, S. & Mussa-Ivaldi, F. A. Remapping hand movements in a novel geometrical environment. *J. Neurophysiol.* **94**, 4362–4372 (2005).
61. Johansson, R. S. *et al.* How a lateralized brain supports symmetrical bimanual tasks. *PLoS Biol.* **4**, e158 (2006).
62. Liu, X., Mosier, K. M., Mussa-Ivaldi, F. A., Casadio, M. & Scheidt, R. A. Reorganization of finger coordination patterns during adaptation to rotation and scaling of a newly learned sensorimotor transformation. *J. Neurophysiol.* **105**, 454–473 (2011).
63. Reis, J. *et al.* Noninvasive cortical stimulation enhances motor skill acquisition over multiple days through an effect on consolidation. *Proc. Natl Acad. Sci. USA* **106**, 1590–1595 (2009).
64. Sutton, R. S. & Barto, A. G. *Reinforcement Learning* (The MIT Press, Cambridge, Massachusetts, 1998).
65. Deisenroth, M., Rasmussen, C. & Peters, J. Gaussian process dynamic programming. *Neurocomputing* **72**, 1508–1524 (2009).
66. Izawa, J. & Shadmehr, R. Learning from sensory and reward prediction errors during motor adaptation. *PLoS Comput. Biol.* **7**, e1002012 (2011).
67. Madelain, L., Paeye, C. & Wallman, J. Modification of saccadic gain by reinforcement. *J. Neurophysiol.* **106**, 219–232 (2011).
68. Abe, M. *et al.* Reward improves long-term retention of a motor memory through induction of offline memory gains. *Curr. Biol.* **21**, 557–562 (2011).
69. Huang, V. S., Haith, A., Mazzoni, P. & Krakauer, J. W. Rethinking motor learning and savings in adaptation paradigms: model-free memory for successful actions combines with internal models. *Neuron* **70**, 787–801 (2011). **This paper suggests that biases in movement direction and savings in relearning during visuomotor adaptation can be explained by two different use-dependent learning mechanisms. In both cases it is the repetition of the actual movement, rather than an error-signal, that underlies the learning changes.**
70. Luft, A. R. & Schwarz, S. Dopaminergic signals in primary motor cortex. *Int. J. Dev. Neurosci.* **27**, 415–421 (2009).
71. Hosp, J. A., Pektanovic, A., Rioult-Pedotti, M. S. & Luft, A. R. Dopaminergic projections from midbrain to primary motor cortex mediate motor skill learning. *J. Neurosci.* **31**, 2481–2487 (2011). **This study shows that dopaminergic neurons from the ventral tegmental area (VTA) in the rat project to primary motor cortex. Lesions of the VTA lead to profound deficits in learning of a reaching task, which could be partially reversed through levodopa administration.**
72. Classen, J., Liepert, J., Wise, S. P., Hallett, M. & Cohen, L. G. Rapid plasticity of human cortical movement representation induced by practice. *J. Neurophysiol.* **79**, 1117–1123 (1998).
73. Bütefisch, C. M. *et al.* Mechanisms of use-dependent plasticity in the human motor cortex. *Proc. Natl Acad. Sci. USA* **97**, 3661–3665 (2000).
74. Verstynen, T. & Sabes, P. N. How each movement changes the next: an experimental and theoretical study of fast adaptive priors in reaching. *J. Neurosci.* **31**, 10050–10059 (2011). **Repeating an arm movement towards the same target decreases the variability of this movement while inducing directional biases for movements to neighbouring targets. This paper offers both a mechanistic (network model) and normative explanation (Bayesian model) for this use-dependent learning effect.**
75. Diedrichsen, J., White, O., Newman, D. & Lally, N. Use-dependent and error-based learning of motor behaviors. *J. Neurosci.* **30**, 5159–5166 (2010).
76. Jordan, M. & Rumelhart, D. Forward models: supervised learning with a distal teacher. *Cogn. Sci.* **16**, 307–354 (1992).
77. d'Avella, A., Portone, A., Fernandez, L. & Lacquaniti, F. Control of fast-reaching movements by muscle synergy combinations. *J. Neurosci.* **26**, 7791–7810 (2006).
78. Sing, G. C., Joiner, W. M., Nanayakkara, T., Braynov, J. B. & Smith, M. A. Primitives for motor adaptation reflect correlated neural tuning to position and velocity. *Neuron* **64**, 575–589 (2009). **This study shows that when subjects are exposed to a novel force field that depends only on position or on the speed of the hand, the initial adaptation is biased towards an interpretation that the force field depends on both position and speed. The interpretation of these results is that there is a strong prior that forces experienced by the hand will depend on both position and velocity in a correlated manner, suggesting that distribution of motor primitives are biased for this correlation.**
79. Smith, M. A., Ghazizadeh, A. & Shadmehr, R. Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS Biol.* **4**, e179 (2006). **This study shows that two distinct processes underlie motor adaptation, one that learns quickly but retains information poorly and one that learns slowly but retains information well. This two-process learning system can account for a range of empirical data that a single process system cannot.**
80. Lee, J.-Y. & Schweighofer, N. Dual adaptation supports a parallel architecture of motor memory. *J. Neurosci.* **29**, 10396–10404 (2009).
81. Tanaka, H., Sejnowski, T. J. & Krakauer, J. W. Adaptation to visuomotor rotation through interaction between posterior parietal and motor cortical areas. *J. Neurophysiol.* **102**, 2921–2932 (2009).
82. Lackner, J. R. & DiZio, P. Motor control and learning in altered dynamic environments. *Curr. Opin. Neurobiol.* **15**, 653–659 (2005).
83. Cothros, N., Wong, J. & Gribble, P. Are there distinct neural representations of object and limb dynamics? *Exp. Brain Res.* **173**, 689–697 (2006).
84. Kluzik, J., Diedrichsen, J., Shadmehr, R. & Bastian, A. Reach adaptation: what determines whether we learn an internal model of the tool or adapt the model of our arm? *J. Neurophysiol.* **100**, 1455–1464 (2008).
85. Cothros, N., Wong, J. & Gribble, P. L. Visual cues signaling object grasp reduce interference in motor learning. *J. Neurophysiol.* **102**, 2112–2120 (2009).
86. Berniker, M. & Kording, K. Estimating the sources of motor errors for adaptation and generalization. *Nature Neurosci.* **11**, 1454–1461 (2008).
87. Haith, A., Jackson, C., Miall, C. & Vijayakumar, S. Unifying the sensory and motor components of sensorimotor adaptation. *Adv. Neural Inf. Process. Syst.* **21**, 593–600 (2009).
88. Cressman, E. K. & Henriques, D. Y. P. Sensory recalibration of hand position following visuomotor adaptation. *J. Neurophysiol.* **102**, 3505–3518 (2009).
89. Ostry, D. J., Darainy, M., Mattar, A. A. G., Wong, J. & Gribble, P. L. Somatosensory plasticity and motor learning. *J. Neurosci.* **30**, 5384–5393 (2010).
90. Bursztyn, L. L. C. D., Ganesh, G., Imamizu, H., Kawato, M. & Flanagan, J. R. Neural correlates of internal-model loading. *Curr. Biol.* **16**, 2440–2445 (2006).
91. Calthness, G. *et al.* Failure to consolidate the consolidation theory of learning for sensorimotor adaptation tasks. *J. Neurosci.* **24**, 8662–8671 (2004).
92. Gupta, R. & Ashe, J. Lack of adaptation to random conflicting force fields of variable magnitude. *J. Neurophysiol.* **97**, 738–745 (2007).
93. Krouchev, N. I. & Kalaska, J. F. Context-dependent anticipation of different task dynamics: rapid recall of appropriate motor skills using visual cues. *J. Neurophysiol.* **89**, 1165–1175 (2003).
94. Nozaki, D., Kurtzer, I. & Scott, S. H. Limited transfer of learning between unimanual and bimanual skills within the same limb. *Nature Neurosci.* **9**, 1364–1366 (2006).
95. Howard, I. S., Ingram, J. N. & Wolpert, D. M. Context-dependent partitioning of motor learning in bimanual movements. *J. Neurophysiol.* **104**, 2082–2091 (2010).
96. Howard, I. S., Ingram, J. N. & Wolpert, D. M. Composition and decomposition in bimanual dynamic learning. *J. Neurosci.* **28**, 10531–10540 (2008).
97. Nozaki, D. & Scott, S. H. Multi-compartment model can explain partial transfer of learning within the same limb between unimanual and bimanual reaching. *Exp. Brain Res.* **194**, 451–463 (2009).
98. Kojima, Y., Iwamoto, Y. & Yoshida, K. Memory of learning facilitates saccadic adaptation in the monkey. *J. Neurosci.* **24**, 7531–7539 (2004).
99. Kording, K. P., Tenenbaum, J. B. & Shadmehr, R. The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nature Neurosci.* **10**, 779–786 (2007).
100. Huang, V. S. & Shadmehr, R. Persistence of motor memories reflects statistics of the learning event. *J. Neurophysiol.* **102**, 931–940 (2009).
101. Anguera, J., Reuter-Lorenz, P., Willingham, D. & Seidler, R. Contributions of spatial working memory to visuomotor learning. *J. Cogn. Neurosci.* **22**, 1917–1930 (2009).
102. Fernandez-Ruiz, J., Wong, W., Armstrong, I. T. & Flanagan, J. R. Relation between reaction time and reach errors during visuomotor adaptation. *Behav. Brain Res.* **219**, 8–14 (2011).
103. Keisler, A. & Shadmehr, R. A shared resource between declarative memory and motor memory. *J. Neurosci.* **30**, 14817–14823 (2010).
104. Brown, R. M. & Robertson, E. M. Inducing motor skill improvements with a declarative task. *Nature Neurosci.* **10**, 148–149 (2007).
105. Bongard, J., Zykov, V. & Lipson, H. Resilient machines through continuous self-modeling. *Science* **314**, 1118–1121 (2006).



106. Braun, D. A., Aertsen, A., Wolpert, D. M. & Mehring, C. Learning optimal adaptation strategies in unpredictable motor tasks. *J. Neurosci.* **29**, 6472–6478 (2009).
107. Braun, D. A., Aertsen, A., Wolpert, D. M. & Mehring, C. Motor task variation induces structural learning. *Curr. Biol.* **19**, 352–357 (2009).  
**This paper shows that after experiencing multiple sensorimotor transformations that conform to a structure (for example, the set of visuomotor rotations), interference between opposite visuomotor rotations is substantially reduced. These results show that subjects can extract the structure across a set of tasks and use this to facilitate learning of new tasks provided that they conform to the learned structure.**
108. Braun, D. A., Mehring, C. & Wolpert, D. M. Structure learning in action. *Behav. Brain Res.* **206**, 157–165 (2010).
109. Braun, D. A., Waldert, S., Aertsen, A., Wolpert, D. M. & Mehring, C. Structure learning in a sensorimotor association task. *PLoS ONE* **5**, e8973 (2010).
110. Huang, V. S., Shadmehr, R. & Diedrichsen, J. Active learning: learning a motor skill without a coach. *J. Neurophysiol.* **100**, 879–887 (2008).
111. Brashers-Krug, T., Shadmehr, R. & Bizzi, E. Consolidation in human motor memory. *Nature* **382**, 252–255 (1996).
112. Brooks, V., Hilperath, F., Brooks, M., Ross, H. & Freund, H. Learning “what” and “how” in a human motor task. *Learn. Mem.* **2**, 225–242 (1995).
113. Ahmed, A. A., Wolpert, D. M. & Flanagan, J. R. Flexible representations of dynamics are used in object manipulation. *Curr. Biol.* **18**, 763–768 (2008).
114. Ingram, J. N., Howard, I. S., Flanagan, J. R. & Wolpert, D. M. Multiple grasp-specific representations of tool dynamics mediate skillful manipulation. *Curr. Biol.* **20**, 618–623 (2010).
115. Churchland, M. M., Afshar, A. & Shenoy, K. V. A central source of movement variability. *Neuron* **52**, 1085–1096 (2006).
116. Burstedt, M., Edin, B. & Johansson, R. S. Coordination of fingertip forces during human manipulation can emerge from independent neural networks controlling each engaged digit. *Exp. Brain Res.* **117**, 67–79 (1997).
117. Reed, K. *et al.* Haptically linked dyads: are two motor-control systems better than one? *Psychol. Sci.* **17**, 365–366 (2006).
118. Braun, D. A., Ortega, P. A. & Wolpert, D. M. Nash equilibria in multi-agent motor interactions. *PLoS Comput. Biol.* **5**, e1000468 (2009).
119. Braun, D. A., Ortega, P. A. & Wolpert, D. M. Motor coordination: when two have to act as one. *Exp. Brain Res.* **221**, 631–641 (2011).
120. Rizzolatti, G. & Luppino, G. The cortical motor system. *Neuron* **31**, 889–901 (2001).
121. Rizzolatti, G. & Craighero, L. The mirror-neuron system. *Annu. Rev. Neurosci.* **27**, 169–192 (2004).
122. Flanagan, J. R. & Johansson, R. S. Action plans used in action observation. *Nature* **424**, 769–771 (2003).
123. Rotman, G., Troje, N. F., Johansson, R. S. & Flanagan, J. R. Eye movements when observing predictable and unpredictable actions. *J. Neurophysiol.* **96**, 1358–1369 (2006).
124. Heyes, C. M. & Foster, C. L. Motor learning by observation: evidence from a serial reaction time task. *Q. J. Exp. Physiol.* **55**, 593–607 (2002).
125. Mattar, A. A. G. & Gribble, P. L. Motor learning by observing. *Neuron* **46**, 153–160 (2005).
126. Wolpert, D. M., Doya, K. & Kawato, M. A unifying computational framework for motor control and social interaction. *Phil. Trans. R. Soc. Lond. B* **358**, 593–602 (2003).
127. Oztog, E., Wolpert, D. M. & Kawato, M. Mental state inference using visual control parameters. *Cogn. Brain Res.* **22**, 129–151 (2005).
128. Aglioti, S. M., Cesari, P., Romani, M. & Urgesi, C. Action anticipation and motor resonance in elite basketball players. *Nature Neurosci.* **11**, 1109–1116 (2008).
129. Brass, M., Schmitt, R. M., Spengler, S. & Gergely, G. Investigating action understanding: inferential processes versus action simulation. *Curr. Biol.* **17**, 2117–2121 (2007).
130. Hesse, M. D., Sparing, R. & Fink, G. R. End or means—the “what” and “how” of observed intentional actions. *J. Cogn. Neurosci.* **21**, 776–790 (2009).
131. Malfait, N. *et al.* fMRI activation during observation of others’ reach errors. *J. Cogn. Neurosci.* **22**, 1493–1503 (2010).
132. Burke, C. J., Tobler, P. N., Baddeley, M. & Schultz, W. Neural mechanisms of observational learning. *Proc. Natl Acad. Sci. USA* **107**, 14431–14436 (2010).
133. White, O. & Diedrichsen, J. Responsibility assignment in redundant systems. *Curr. Biol.* **20**, 1290–1295 (2010).  
**This study shows that when participants control a single cursor that is located at the spatial average of the two hands, the left hand corrects more for cursor errors that are induced by a visuomotor rotation, although the right hand corrects more efficiently during unimanual movements. This indicates that the motor system assigns the error during redundant movement to the more likely source of the error.**
134. Vetter, P., Flash, T. & Wolpert, D. M. Planning movements in a simple redundant task. *Curr. Biol.* **12**, 488–491 (2002).

#### Acknowledgements

We thank the Wellcome Trust, the Canadian Institutes of Health Research and the Human Frontiers Science Programme for support. J.D. is supported by a Scholar award from the James S. McDonnell foundation.

#### Competing interests statement

The authors declare no competing financial interests.

#### FURTHER INFORMATION

Daniel M. Wolpert’s homepage: [www.wolpertlab.com](http://www.wolpertlab.com)

Jörn Diedrichsen’s homepage: [www.icn.ucl.ac.uk/](http://www.icn.ucl.ac.uk/)

[motorcontrol](#)

J. Randall Flanagan’s homepage: [www.flanaganlab.com](http://www.flanaganlab.com)

ALL LINKS ARE ACTIVE IN THE ONLINE PDF