

Computational principles of movement neuroscience

Daniel M. Wolpert¹ and Zoubin Ghahramani²

¹ Sobell Department of Neurophysiology, Institute of Neurology, Queen Square, University College London, London WC1N 3BG, UK

² Gatsby Computational Neuroscience Unit, Queen Square, University College London, London WC1N 3AR, UK

Correspondence should be addressed to D.M.W. (wolpert@hera.ucl.ac.uk)

Unifying principles of movement have emerged from the computational study of motor control. We review several of these principles and show how they apply to processes such as motor planning, control, estimation, prediction and learning. Our goal is to demonstrate how specific models emerging from the computational approach provide a theoretical framework for movement neuroscience.

The computational study of motor control is fundamentally concerned with the relationship between sensory signals and motor commands. The transformation from motor commands to their sensory consequences is governed by the physics of the environment, the musculoskeletal system and sensory receptors. The transformation from sensory signals to motor commands is determined by processes within the central nervous system (CNS). Much of the complexity of our behavior arises as a simple coupling of these two transformations.

Although the CNS is not required to represent the motor-to-sensory transformation, as this is implemented by the physical world, exciting computational and experimental developments are arising from the realization that the CNS internally represents this transformation. Systems that model aspects of this transformation are known as 'forward internal models' because they model the causal relationship between actions and their consequences. The primary role of these models is to predict the behavior of the body and world, so we use the terms 'predictors' and 'forward models' synonymously. Systems that implement the opposite transformations, from desired consequences to actions, are known as 'inverse internal models'. We use the term 'internal model' to emphasize that the CNS is modeling the sensorimotor system, and that these are not models of the CNS.

Bellman¹ pointed out that when representing such transformations, the storage and number of calculations required increases exponentially with the dimension of the sensory and motor arrays, a problem he termed "the curse of dimensionality". This is vividly illustrated if we consider the 600 or so muscles in the human body as being, for extreme simplicity, either contracted or relaxed. This leads to 2^{600} possible motor activations, more than the number of atoms in the universe. This clearly prohibits a simple look-up table from motor activations to sensory feedback and vice versa. Fortunately for control, a compact representation with far lower dimensionality than the full sensor and motor array can generally be extracted, known as the 'state' of the system. When taken together with fixed parameters of the system and the equations governing the physics, the state contains all the relevant time-varying information needed to predict or control the future of the system. For example, knowing the current position, velocity and spin of a thrown ball, that is the ball's state, allows predictions of its future path without the need to know the configurations of all the atoms in the ball.

In general, the state, for example the set of activations of groups of muscles (synergies) or the position and velocity of the hand, changes rapidly and continuously within a movement. However, other key parameters change discretely, like the identity of a manipulated object, or on a slower time-scale, like the mass of the limb. We refer to such discrete or slowly changing parameters as the 'context' of the movement. Our ability to generate accurate and appropriate motor behavior relies on tailoring our motor commands to the prevailing movement context.

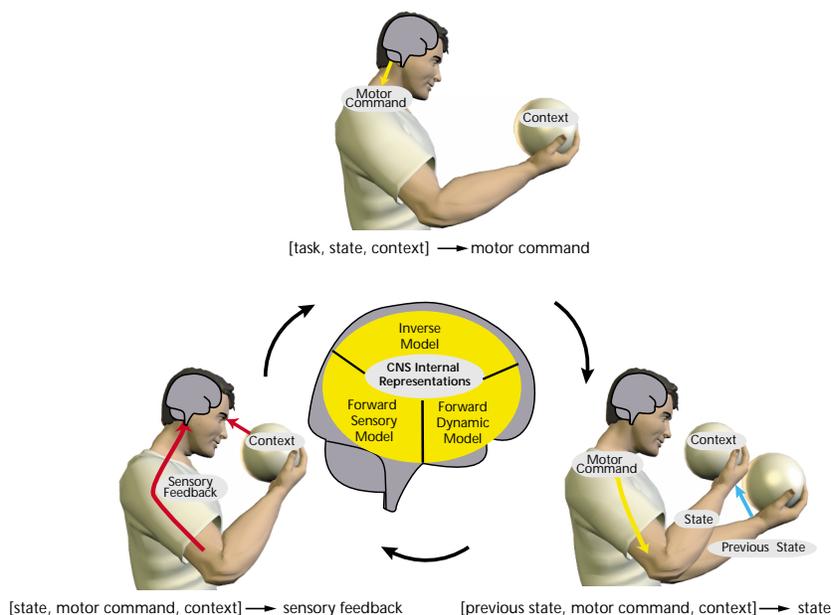
The sensorimotor loop (Fig. 1) can be divided into three stages, which govern the overall behavior of the sensorimotor system. The first stage specifies the motor command generated by the CNS given the state and a particular task (Fig. 1, top). The second stage determines how the state changes given the motor command (Fig. 1, right). The third closes the loop by specifying the sensory feedback given this new state (Fig. 1, left). These three stages are represented in the CNS as internal models, being the inverse model, forward dynamic model and forward sensory model respectively.

We will show how all the main themes of computational motor control, such as planning, control and learning, arise from considering how tasks determine behavior, motor commands are generated, states and contexts are estimated and predicted, and internal models are represented and learned. With the advent of virtual reality technologies and novel robotic interfaces, it has become possible, for the first time, to create sophisticated computer-controlled environments. Having such control over the physics of the world with which subjects interact has allowed detailed tests of computational models of planning, control and learning (for example, refs. 2–6).

Task: motor planning

Everyday tasks are generally specified at a high, often symbolic level, such as taking a drink of water from a glass. However, the motor system must eventually work at a detailed level, specifying muscle activations leading to joint rotations and the path of the hand in space. There is clearly a gap between the high-level task and low-level control. Indeed, almost any task can in principle be achieved in infinitely many different ways. Consider, for example, the number of ways, some sensible and some silly, in which you can bring a glass of water to your lips. Given all these possibilities, it is surprising that almost every study of how the

Fig. 1. The sensorimotor loop, showing motor command generation (top), state transition (right) and sensory feedback generation (left). Center, internal representation of these stages within the CNS.



motor system solves a given task shows highly stereotyped movement patterns, both between repetitions of a task and between individuals on the same task. In the same way that being able to rank different routes from New York to London allows us to select from those available, having a criterion with which to evaluate possible movements for a task would allow the CNS to select the best. Optimal control is an elegant framework for dealing with just such a selection problem and can, therefore, translate from high-level tasks into detailed motor programs. Specifically, a cost is specified as some function of the movement and task, and the movement with the lowest cost is executed. The challenge has been to try to reverse-engineer the cost function, that is, what is being optimized, from observed movement patterns and perturbation studies.

Optimal control models of movement have been proposed based on maximizing smoothness of the hand trajectory⁷ and of the torque commands⁸. Although these models successfully reproduce a range of empirical data, it is unclear why smoothness is important, and how it is measured by the CNS. Moreover, these models are limited to a single motor system such as the arm. Recently, a model has been proposed which provides a unifying cost for goal-directed eye and arm movements⁹. This model assumes that there is noise in the motor command and that the amount of noise scales with the motor command's magnitude. In the presence of such noise, the same sequence of intended motor commands if repeated many times will lead to a probability distribution over movements. Aspects of this distribution, such as the spread of positions or velocities of the hand at the end of the movement, can be controlled by modifying the sequence of motor commands. In this model, the task specifies how aspects of the distribution are penalized, and this forms the cost. For example, in a simple aiming movement, the task is to minimize the final error, as measured by the variance about the target. Figure 2 shows the consequences of two possible sequences of motor commands, one of which leads to higher endpoint variability (blue ellipsoid) than the other. The aim of the optimal control strategy is to minimize the volume of the ellipsoid, thereby being as accurate as possible. This model accurately predicts the trajectories of both saccadic eye movements and arm movements. Non-smooth movements require large motor commands, which generate increased noise; smoothness thereby leads to accuracy but is not a goal in its own right. The cost, movement error, is behaviorally relevant and simple for the CNS to measure. Climbing fibers, which may act as a training signal to the cerebellum, code for such reaching errors at the end of a movement¹⁰. Although the cost function specifies the optimal movement, how we approach this optimum for novel, unrehearsed movements is an open question.

Motor command: control

Several models of how motor commands are generated have been proposed. One of the first proposals was that the CNS specifies

spatial parameters and relies on the spring-like properties of muscles and reflex loops to move the limb^{11–13}. For example, in one form of this model, a set of muscle activations defines a stable equilibrium position of the hand in space. Movement is achieved via a succession of such equilibrium positions along a desired trajectory¹⁴. This model uses muscle and spinal cord properties as a feedback controller to pull the hand along the desired trajectory. Because of the system dynamics, the actual positions will not follow exactly the desired trajectory unless the stiffness (the force generated per unit displacement from equilibrium) of the muscle and reflex loops is high. An alternative proposal is that an inverse model is constructed to map the desired state at each point along the trajectory into motor commands¹⁵. These two approaches can be contrasted by considering the problem of moving a ball around a circular path by specifying the forces acting on it. For inverse model control, the equations of motion are

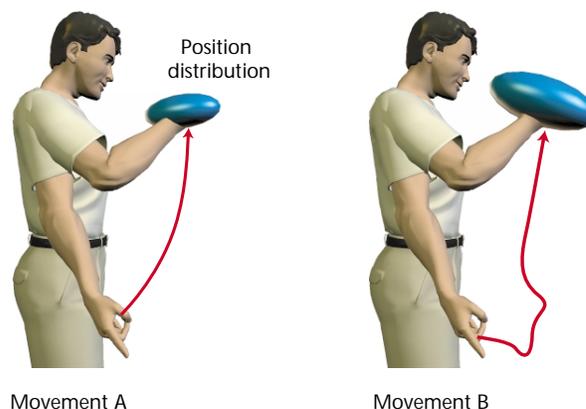


Fig. 2. Task optimization in the presence of signal-dependent noise (TOPS) model of Harris and Wolpert⁹. Average paths and expected final position distributions for two different motor sequences. Although the sequences bring the hand on average to the same final position, they have different final distributions because of noise in the motor commands. Movement A has smaller spread than B and therefore has lower cost than B. In general, the task determines the desired statistics of the movement, and the trajectory that optimizes the statistics is selected.

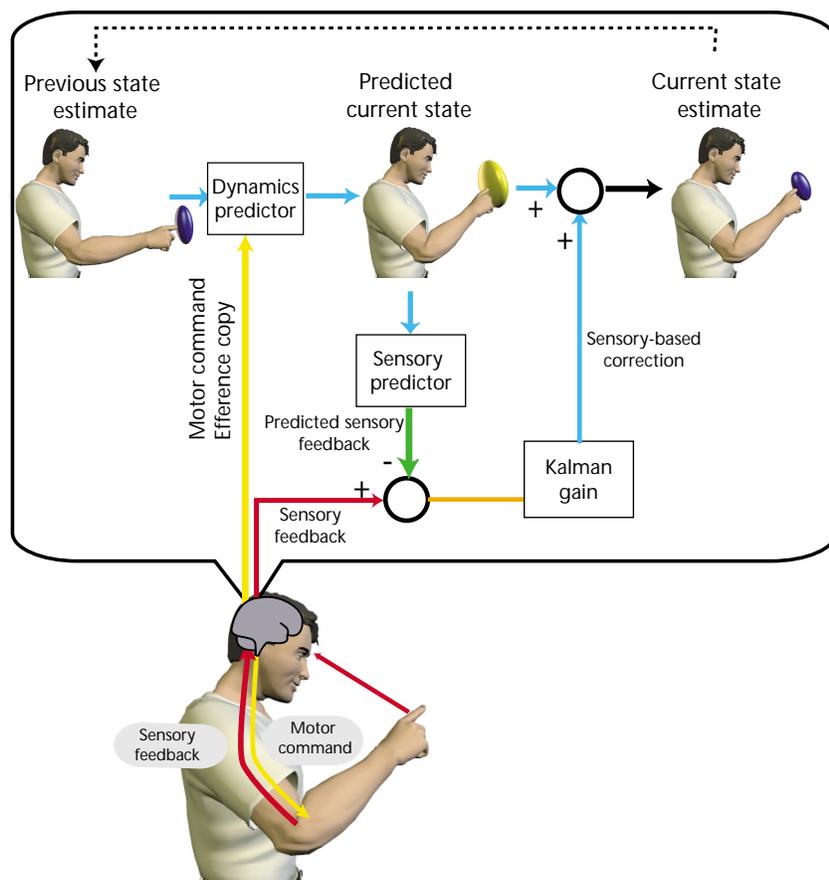


Fig. 3. A schematic of one step of a Kalman filter model recursively estimating the finger's location during a movement. The current state is constructed from the previous state estimate (top left), which represents the distribution of possible finger positions, shown as a cloud of uncertainty (blue). A copy of the motor command, that is, an efference copy, and a model of the dynamics allow the current state distribution to be predicted from this previous state. In general, the uncertainty is increased (yellow cloud). This new estimate is then refined by using it to predict the current sensory feedback. The error between this prediction and the actual sensory feedback is used to correct the current estimate. The Kalman gain changes this sensory error into state errors and also determines the relative reliance placed on the efference copy and sensory feedback. The final state estimate (top right) now has a reduced uncertainty (blue cloud). Sensory signals, red arrows; motor commands, yellow arrows; state signals, blue arrows. Although delays in sensory feedback must be compensated, they have been omitted from the diagram for clarity.

to the lowest point in the landscape, the arm will move in the force field to an equilibrium point. By summing valleys together, new landscapes, that is behaviors, can be constructed.

At a higher level, in motor cortex, there has been considerable debate about what is coded within individual neurons and pop-

ulations^{19–23}. Proposals include the direction of movement, velocity, acceleration, posture and joint torques. A recent model²⁴ proposes that the cortex activates muscle groups and that many of these conflicting views can be resolved by considering the relationship among cortical activity, muscle filtering properties and movement kinematics. Given the kinematics of observed movements, the model shows that motor cortical neurons would be expected to encode all of the above movement properties.

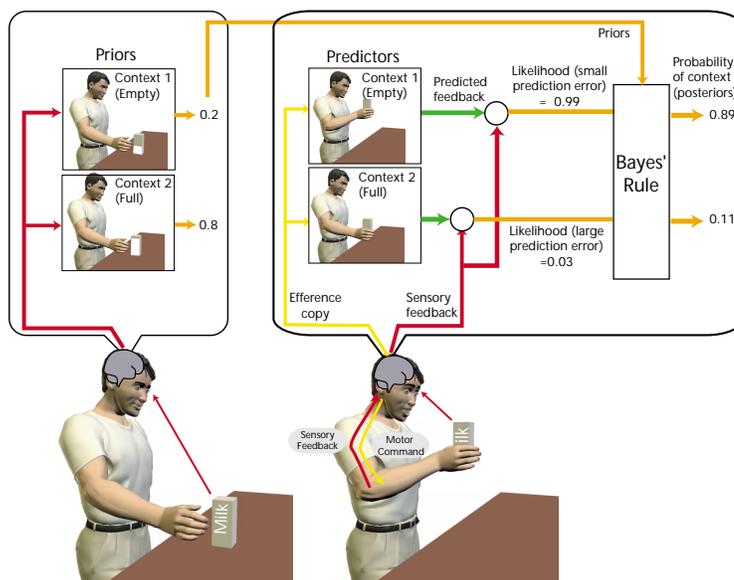
State: estimation and prediction

To model any of the three stages of the sensorimotor loop (Fig. 1), the CNS needs to know the current state, but it faces two problems. First, the transduction and transport of sensory signals to the CNS involve considerable delays. Second, the CNS must estimate the system's state from sensory signals that may be contaminated by noise and may only provide partial information about the state. For example, consider a tennis ball we have just hit. If we simply used the ball's retinal location to estimate its position, our estimate would be delayed by around 100 ms. A better estimate can be made by predicting where the ball actually is now using a forward model. Second, the ball's spin cannot be directly observed, but it can be estimated using sensory information integrated over time, that is the ball's seen path. This estimate can be improved by knowing how the ball was hit, that is the motor command, in conjunction with an internal forward model of the ball's dynamics. This combination, using sensory feedback and forward models to estimate the current state, is known as an 'observer', an example of which is the Kalman filter²⁵ (Fig. 3). The major objectives of the observer are to compensate for sensorimotor delays and to reduce the uncertainty in the state estimate that arises because of noise inherent in both sensory and

solved and the forces on the ball applied to generate the desired acceleration. For equilibrium point control, the ball is attached by a spring to a control point, which is simply moved around the circular path, with the ball following along behind. For the ball to travel along the circle, the stiffness of the spring, representing the muscles and reflex loops, must be high. If the stiffness is low, for example if a slinky were used, the path of the controlled point and ball would be very different. In contrast, with an inverse model, the arm can be controlled with low stiffness. Currently there is debate as to the gain (as measured by stiffness of the limb during movement), with one study⁵ suggesting that the stiffness is low. With low stiffness, equilibrium point control would require a complex equilibrium trajectory to achieve the simple movements observed, thereby making it a less attractive computational solution.

The final common pathway of the motor command is formed by about 200,000 alpha motor neurons. The question arises as to what the coding of this motor output is and how high-level objectives are represented in the motor command. In spinal cord, experiments support the computationally attractive idea that control is achieved by activating a few motor primitives or basis functions¹⁶. The idea is to simplify control by combining a small number of primitives, such as patterns of muscle activations, in different proportions rather than individually controlling each muscle. Such control reduces the effective dimensionality of the motor output. Spinal stimulation studies in frog suggest that the spinal cord may control movement through primitives that correspond to force fields in extrinsic space and can be combined to provide a 'grammar' with which to construct complex movements^{17,18}. Each force field primitive can be thought of as a differently shaped valley. Just as water flows down along some path

Fig. 4. A schematic of context estimation with just two contexts, that a milk carton is empty or full. Initially sensory information from vision is used to set the prior probabilities of the two possible contexts and, in this case, the carton appears more likely to be full. When the motor commands appropriate for a full carton are generated, an efference copy of the motor command is used to simulate the sensory consequences under the two possible contexts. The predictions based on an empty carton suggest a large amount of movement compared to the full-carton context. These predictions are compared with actual feedback. Because as the carton is, in fact, empty, the sensory feedback matches the predictions of the empty-carton context. This leads to a high likelihood for the empty carton and a low likelihood for the full carton. The likelihoods are combined with the priors using Bayes' rule to generate the final (posterior) probability of each context.



motor signals. For a linear system, the Kalman filter is the optimal observer in that it estimates the state with the least squared error. Such a model has been supported by empirical studies examining estimation of hand position^{3,26}, posture²⁷ and head orientation²⁸. Damage to parietal cortex can lead to an inability to maintain such state estimates²⁹.

Using the observer framework, it is a simple computational step from estimating the current state to predicting future states and sensory feedback. Such predictions have many potential benefits³⁰. State prediction, by estimating the outcome of an action before sensory feedback is available, can reduce the effect of feedback delays in sensorimotor loops. Such a system is thought to underlie skilled manipulation. For example, when an object held in the hand is accelerated, the fingers tighten their grip in anticipation to prevent the object slipping, a process that relies on prediction (for review, see ref. 31). State prediction can also be used in mental simulation of intended movements. Damage to parietal cortex can lead to an inability to mentally simulate movements with the affected hand³².

Sensory prediction can be derived from the state prediction and used to cancel the sensory effects of movement, that is, refference. By using such a system, it is possible to cancel out the effects of sensory changes induced by self-motion, thereby enhancing more relevant sensory information. Such a mechanism has been extensively studied in the electric fish, where it relies on a cerebellum-like structure (for example, ref. 33). In primates, neurophysiological studies³⁴ show predictive updating in parietal cortex, anticipating the retinal consequences of an eye movement. In man, predictive mechanisms are believed to underlie the observation that the same tactile stimulus, such as a tickle, is felt less intensely when it is self-applied. The reduced intensity of self-applied tactile stimuli critically depends upon on the precise spatiotemporal alignment between the predicted and actual sensory consequences of the movement³⁵. Similarly, sensory predictions provide a mechanism to determine whether a movement is self-produced, and hence predictable, or produced externally. A failure in this mechanism is proposed to underlie delusions of control, in which it appears to patients that their bodies are being moved by forces other than their own³⁶. Interestingly, damage to the left parietal cortex can lead to a relative inability to determine whether viewed movements are one's own or not³⁷.

Context: estimation

When we interact with objects with different physical characteristics, the context of our movement changes in a discrete

manner. Just as it is essential for the motor system to estimate the state, it must also estimate the changing context. One powerful formalism is the Bayesian approach, which can be used to estimate the probability of each context. The probability can be factored into two terms, the likelihood and the prior. The likelihood of a particular context is the probability of the current sensory feedback given that context. To estimate this likelihood, a sensory forward model of that context is used to predict the sensory feedback from the movement. The discrepancy between the predicted and actual sensory feedback is inversely related to the likelihood: the smaller the prediction error, the more likely the context. These computations can be carried out by a modular neural architecture in which multiple predictive models operate in parallel³⁸. Each is tuned to one context and estimates the relative likelihood of its context. This array of models therefore act as a set of hypothesis testers. The prior contains information about the structured way contexts change over time and how likely a context is before a movement. The likelihood and the prior can be optimally combined using Bayes' rule, which takes the product of these two probabilities and normalizes over all possible contexts, to generate a probability for each context. Figure 4 shows an example of picking up what appears to be a full milk carton, which is in reality empty. The predictive models correct on-line for erroneous priors that initially weighted the output of the controller for a full carton more than that for an empty carton. Bayes' rule allows a quick correction to the appropriate control even though the initial strategy was incorrect. This example has two modules representing two contexts. However, the modular architecture can, in principle, scale to thousands of modules, that is contexts. Although separate architectures have been proposed for state and context estimation (Figs. 3 and 4), they both can be considered as on-line ways of doing Bayesian inference in an uncertain environment.

This interpretation of the processes necessary for context estimation is consistent with recent neurophysiological studies in primates, showing that the CNS both models the expected sensory feedback for a particular context³⁹, and represents the likelihood of the sensory feedback given the context⁴⁰. In an elegant example of context estimation⁶, when subjects make a reaching movement while rotating their torso, they compensate for the velocity-dependent Coriolis forces that arise from the

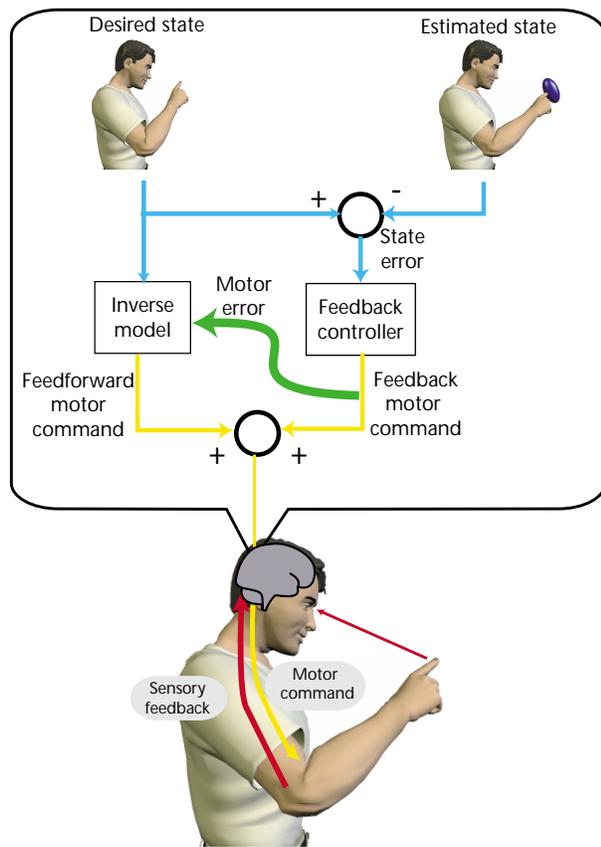


Fig. 5. A schematic of feedback-error learning. The aim is to learn an inverse model that can generate motor commands given a series of desired states. A hard-wired and low-gain feedback controller is used to correct for errors between desired and estimated states. This generates a feedback motor command that is added to the feedforward motor command generated by the inverse model. If the feedback motor command goes to zero, then the state error, in general, will also be zero. Therefore the feedback motor command is a measure of the error of the inverse model and is used as the error signal to train it.

rotation and act on the arm. When subjects experience illusory self-rotation induced by a large moving visual image, they make movements as though they expect, based on the visual priors, the context of the Coriolis force. This leads to misreaching, which is reduced over subsequent movements as the sensory consequences of the expected Coriolis force are not experienced.

Internal models: learning

Internal models, both forward and inverse, capture information about the properties of the sensorimotor system. These properties are not static but change throughout life, both on a short time-scale, due to interactions with the environment, and on a longer time-scale, due to growth. Internal models must therefore be adaptable to changes in the sensorimotor system's properties. The environment readily provides an appropriate training signal to learn predictors of sensory feedback. The difference between the predicted and actual sensory feedback can be used as an error signal to update a predictive model. The neural mechanisms that lead to such predictive learning in the cerebellum-like structure of electric fish are partially understood³³.

Acquiring an inverse internal model through motor learning

is generally a difficult task. This is because the appropriate training signal, the motor command error, is not directly available. When we fail to bowl a googly, no one tells us how our muscle activations should change to achieve this cricketing feat. Instead we receive error signals in sensory coordinates, and these sensory errors need to be converted into motor errors before they can be used to train an inverse model. The feedback-error learning model^{15,41} provides an ingenious solution to this problem (Fig. 5). A hard-wired, but not perfect, feedback controller computes a motor command based on the discrepancy between desired and estimated states. The motor command is the sum of the feedback controller motor command and the output of an adaptive inverse model. The rationale behind this model is that if the feedback controller ends up producing no motor command, then there must be no discrepancy between desired and estimated state, that is, no error in performance, and the inverse model is performing perfectly. Thus the output of the feedback controller can be regarded as the error signal, and used to train the inverse model, an approach that is highly successful. Neurophysiological evidence⁴² supports this learning mechanism within the cerebellum for the simple reflex eye movement called the ocular following response, suggesting that the cerebellum constructs an inverse model of the eye's dynamics.

Recent work on dynamic learning has focused on the representation of the inverse model. If subjects make point-to-point movements in a force field generated either by a robot attached to their hand² or by a rotating room⁴³, over time they adapt and are able to move naturally in the presence of the field. Several theoretical questions have been addressed using this protocol. The learning of dynamics generalizes in joint-based coordinates², learning depends on the states experienced but not on their temporal order⁴⁴, state-dependent fields are learned more efficiently than temporally changing fields⁴⁵, and both forward and inverse models are simultaneously adapted during learning⁴⁶.

Focus has begun to shift away from examining learning of a single context to consider how we are able to learn a variety of contexts. One architecture that can learn to act in multiple contexts is the modular selection and identification for control (MOSAIC) model, which contains multiple predictor-controller pairs³⁸. As described above, the predictors provide the probability of each context, and these probabilities are used to weight the outputs of a set of corresponding controllers tuned to each context. This system can simultaneously learn the multiple predictors and controllers necessary and how to select the controller appropriate for a given context.

Our understanding of the mechanisms of motor learning has gained from examining how learning one context can interfere with learning others. When subjects try to learn two different dynamics^{47,48} or visuomotor rearrangements⁴⁹, interference occurs when they are presented in quick succession, but not when they are separated by several hours. This suggests that motor learning undergoes a period of consolidation, during which time the motor memory is fragile to being disrupted. Interestingly, visuomotor and dynamic perturbations do not interfere and can be learned as fast together as they can individually⁴⁹.

Many situations that we encounter are derived from combinations of previously experienced contexts, such as novel conjoints of manipulated objects and environments. By modulating the contribution of the outputs of the individual controllers to the final motor command, an enormous repertoire of

behaviors can be generated. Therefore, multiple internal models can be regarded conceptually as motor primitives, the building blocks used to construct intricate motor behaviors with an enormous vocabulary. After learning two different contexts, the CNS can appropriately mix the outputs both within the visuomotor domain⁴ and across the visuomotor and dynamic domains⁵⁰.

Unifying principles

Computational approaches have started to provide unifying principles for motor control. Several common themes have already emerged. First, internal models are fundamental for understanding a range of processes such as state estimation, prediction, context estimation, control and learning. Second, optimality underlies many theories of movement planning, control and estimation and can account for a wide range of experimental findings. Third, the motor system has to cope with uncertainty about the world and noise in its sensory inputs and motor commands, and the Bayesian approach provides a powerful framework for optimal estimation in the face of such uncertainty. We believe that these and other unifying principles will be found to underlie the control of motor systems as diverse as the eye, arm, speech, posture, balance and locomotion.

ACKNOWLEDGEMENTS

We thank Pierre Baraduc, Robert van Beers, James Ingram, Kelvin Jones and Philipp Vetter for comments on the manuscript. This work was supported by grants from the Wellcome Trust, the Gatsby Charitable Foundation and the Human Frontiers Science Organization.

RECEIVED 16 JUNE; ACCEPTED 29 SEPTEMBER 2000

- Bellman, R. *Dynamic Programming* (Princeton Univ. Press, Princeton, New Jersey, 1957).
- Shadmehr, R. & Mussa-Ivaldi, F. Adaptive representation of dynamics during learning of a motor task. *J. Neurosci.* 14, 3208–3224 (1994).
- Wolpert, D. M., Ghahramani, Z. & Jordan, M. I. An internal model for sensorimotor integration. *Science* 269, 1880–1882 (1995).
- Ghahramani, Z. & Wolpert, D. M. Modular decomposition in visuomotor learning. *Nature* 386, 392–395 (1997).
- Gomi, H. & Kawato, M. Equilibrium-point control hypothesis examined by measured arm stiffness during multijoint movement. *Science* 272, 117–120 (1996).
- Cohn, J. V., DiZio, P. & Lackner, J. R. Reaching during virtual rotation: context specific compensations for expected coriolis forces. *J. Neurophysiol.* 83, 3230–3240 (2000).
- Flash, T. & Hogan, N. The co-ordination of arm movements: An experimentally confirmed mathematical model. *J. Neurosci.* 5, 1688–1703 (1985).
- Uno, Y., Kawato, M. & Suzuki, R. Formation and control of optimal trajectories in human multijoint arm movements: Minimum torque-change model. *Biol. Cybern.* 61, 89–101 (1989).
- Harris, C. M. & Wolpert, D. M. Signal-dependent noise determines motor planning. *Nature* 394, 780–784 (1998).
- Kitazawa, S., Kimura, T. & Yin, P. Cerebellar complex spikes encode both destinations and errors in arm movements. *Nature* 392, 494–497 (1998).
- Feldman, A. G. Functional tuning of the nervous system with control of movement or maintenance of a steady posture. III. Mechanographic analysis of execution by arm of the simplest motor tasks. *Biophysics* 11, 766–775 (1966).
- Bizzi, E., Accornerro, N., Chapple, B. & Hogan, N. Posture control and trajectory formation during arm movement. *J. Neurosci.* 4, 2738–2744 (1984).
- Hogan, N. An organizing principle for a class of voluntary movements. *J. Neurosci.* 4, 2745–2754 (1984).
- Flash, T. The control of hand equilibrium trajectories in multi-joint arm movements. *Biol. Cybern.* 57, 257–274 (1987).
- Kawato, M., Furukawa, K. & Suzuki, R. A hierarchical neural network model for the control and learning of voluntary movements. *Biol. Cybern.* 56, 1–17 (1987).
- Giszter, S. F., Mussa-Ivaldi, F. A. & Bizzi, E. Convergent force fields organized in the frog's spinal cord. *J. Neurosci.* 13, 467–491 (1993).
- Tresch, M. C., Saltiel, P. & Bizzi, E. The construction of movement by the spinal cord. *Nat. Neurosci.* 2, 162–167 (1999).
- Mussa-Ivaldi, F. A. Modular features of motor control and learning. *Curr. Opin. Neurobiol.* 9, 713–717 (1999).
- Mussa-Ivaldi, F. A. Do neurons in the motor cortex encode movement direction? An alternative hypothesis. *Neurosci. Lett.* 91, 106–111 (1988).
- Sanger, T. Theoretical considerations for the analysis of population coding in motor cortex. *Neural Comput.* 6, 29–37 (1994).
- Georgopoulos, A. P. Current issues in directional motor control. *Trends Neurosci.* 18, 506–510 (1995).
- Scott, S. & Kalaska, J. F. Motor cortical activity is altered by changes in arm posture for identical hand trajectories. *J. Neurophysiol.* 73, 2563–2567 (1995).
- Kakei, S., Hoffman, D. S. & Strick, P. L. Muscle and movement representations in the primary motor cortex. *Science* 285, 2136–2139 (1999).
- Todorov, E. Direct cortical control of muscle activation in voluntary arm movements: a model. *Nat. Neurosci.* 3, 391–398 (2000).
- Goodwin, G. C. & Sin, K. S. *Adaptive Filtering Prediction and Control* (Prentice-Hall, Englewood Cliffs, New Jersey, 1984).
- van Beers, R. J., Sittig, A. C. & van der Gon, J. J. D. Integration of proprioceptive and visual position information: An experimentally supported model. *J. Neurophysiol.* 81, 1355–1364 (1999).
- Kuo, A. D. An optimal-control model for analyzing human postural balance. *IEEE Trans. Biomed. Eng.* 42, 87–101 (1995).
- Merfeld, D. M., Zupan, L. & Peterka, R. J. Humans use internal model to estimate gravity and linear acceleration. *Nature* 398, 615–618 (1999).
- Wolpert, D. M., Goodbody, S. J. & Husain, M. Maintaining internal representations: the role of the superior parietal lobe. *Nat. Neurosci.* 1, 529–533 (1998).
- Miall, R. C. & Wolpert, D. M. Forward models for physiological motor control. *Neural Networks* 9, 1265–1279 (1996).
- Johansson, R. S. & Cole, K. J. Sensory-motor coordination during grasping and manipulative actions. *Curr. Opin. Neurobiol.* 2, 815–823 (1992).
- Sirigu, A. *et al.* The mental representation of hand movements after parietal cortex damage. *Science* 273, 1564–1568 (1996).
- Bell, C. C., Han, V. Z., Sugawara, Y. & Grant, K. Synaptic plasticity in a cerebellum-like structure depends on temporal order. *Nature* 387, 278–281 (1997).
- Duhamel, J. R., Colby, C. L. & Goldberg, M. E. The updating of the representation of visual space in parietal cortex by intended eye movements. *Science* 255, 90–92 (1992).
- Blakemore, S. J., Frith, C. D. & Wolpert, D. M. Perceptual modulation of self-produced stimuli: The role of spatio-temporal prediction. *J. Cogn. Neurosci.* 11, 551–559 (1999).
- Frith, C. D. *The Cognitive Neuropsychology of Schizophrenia* (Lawrence Erlbaum, Hove, UK, 1992).
- Sirigu, A., Daprati, E., Pradatdiehl, P., Franck, N. & Jeannerod, M. Perception of self-generated movement following left parietal lesion. *Brain* 122, 1867–1874 (1999).
- Wolpert, D. M. & Kawato, M. Multiple paired forward and inverse models for motor control. *Neural Networks* 11, 1317–1329 (1998).
- Eskandar, E. N. & Assad, J. A. Dissociation of visual, motor and predictive signals in parietal cortex during visual guidance. *Nat. Neurosci.* 2, 88–93 (1999).
- Kim, J. & Shadlen, M. N. Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat. Neurosci.* 2, 176–185 (1999).
- Kawato, M. & Gomi, H. The cerebellum and VOR/OKR learning models. *Trends Neurosci.* 15, 445–453 (1992).
- Shidara, M., Kawano, K., Gomi, H. & Kawato, M. Inverse-dynamics encoding of eye movement by Purkinje cells in the cerebellum. *Nature* 365, 50–52 (1993).
- Lackner, J. R. & DiZio, P. Rapid adaptation to Coriolis force perturbations of arm trajectory. *J. Neurophysiol.* 72, 299–313 (1994).
- Condit, M. A., Gandolfo, F. & Mussa-Ivaldi, F. A. The motor system does not learn dynamics of the arm by rote memorization of past experience. *J. Neurophysiol.* 78, 554–560 (1997).
- Condit, M. A. & Mussa-Ivaldi, F. A. Central representation of time during motor learning. *Proc. Natl. Acad. Sci. USA* 96, 11625–11630 (1999).
- Bhushan, N. & Shadmehr, R. Computational nature of human adaptive control during learning of reaching movements in force fields. *Biol. Cybern.* 81, 39–60 (1999).
- Brashers-Krug, T., Shadmehr, R. & Bizzi, E. Consolidation in human motor memory. *Nature* 382, 252–255 (1996).
- Gandolfo, F., Mussa-Ivaldi, F. A. & Bizzi, E. Motor learning by field approximation. *Proc. Natl. Acad. Sci. USA* 93, 3843–3846 (1996).
- Krakauer, J. W., Ghilardi, M. F. & Ghez, C. Independent learning of internal models for kinematic and dynamic control of reaching. *Nat. Neurosci.* 2, 1026–1031 (1999).
- Flanagan, J. R. *et al.* Composition and decomposition of internal models in motor learning under altered kinematic and dynamic environments. *J. Neurosci.* 19, B1–B5 (1999).