

Computational motor control

Daniel M. Wolpert* & Zoubin Ghahramani⁺

*Sobell Department of Motor Neuroscience, Institute of Neurology, Queen Square
University College London, London WC1N 3BG, UK

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

Abstract

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.

Introduction

From a computational perspective the sensorimotor system allows us to take actions to achieve goals in an uncertain and varying world. We will consider a very general framework to phrase the computational problems of motor control and show how the main concepts of sensorimotor control arise from this framework. Consider a person who interacts with the environment by producing actions. The actions or motor outputs will cause muscle activations, and based on the physics of the musculoskeletal system and the outside world will lead to a new state of both the person and environment. By *state* we refer to the set of time-varying parameters which taken together with the fixed parameters of the system, the equations of motion of the body and world, and the motor output allow a prediction of the consequences of the action. For example to predict how a pendulum responds to a torque acting on it you would need to know the pendulum's angle and angular velocity, which together would form its state. However, fixed parameters such as the length and mass of the pendulum would not form part of the state. 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 central nervous system (CNS) does not have direct access to the state but receives as its input sensory feedback. The sensory inputs provide information about the state of world, such as the location of objects, as well as information about the state of our own body, such as the position and velocity of the hand. In addition to these sensory inputs, the central nervous system can monitor its own activity. For example, a copy of the motor output can be used to provide information about the ongoing movement. This signal is known as an *effference copy* to reflect that it is copy of the signal flowing out of the

central nervous system to the muscles. We can also consider some sensory inputs as providing reward, for example the taste of chocolate or warmth on a cold day, or punishment, such as hunger or pain. While some rewards or punishments are directly specified by the environments others may be indirectly or internally generated.

Within this framework we can consider the goal of motor control as selecting actions to maximize future rewards. For example, an infant may generate actions and receive reward if the actions bring food into its mouth, but punishment (negative rewards) if it bites its own fingers. Therefore, it has to choose actions to maximize food intake while minimizing the chance of biting itself. Often in computational motor control we specify a discount factor so that we regard an action that will lead to a reward tomorrow of less value than another action which will lead to the same reward immediately. Conversely, we only choose an action that will achieve a reward at some distant time if the reward greatly exceeds the immediate reward we would get for all other actions.

We will show how all the main themes of computational motor control, such as planning, control and learning arise from considering how *optimality* can be used to plan movements, *motor commands* are generated, *states* and *contexts* are estimated and predicted, and *internal models* are represented and learned. Recent progress in motor control has come both from more sophisticated theories and from the advent of virtual reality technologies and novel robotic interfaces. Using these technologies it has been possible, for the first time, to create sophisticated computer-controlled environments. Having such control over the physics of the world that subjects interact with has allowed detailed tests of computational models of planning, control and learning (e.g. Shadmehr and Mussa-Ivaldi 1994 ; Wolpert et al. 1995 ; Gomi and Kawato 1996; Ghahramani and Wolpert 1997; Cohn et al. 2000).

Optimal Control

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 has to 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 kinematics (Bernstein 1967). In fact, almost any task can in principle be achieved in infinitely many different ways. Given all these possibilities, it is surprising that almost every study of the way the motor system solves a given task shows highly stereotyped movement patterns, both between repetitions of a task and between individuals on the same task location (e.g. Morasso 1981; Flash and Hogan 1985). The concept that some movements will lead to reward and other to punishment links naturally to the field of optimal motor control. Specifically, a cost (that can be thought of as punishment-reward) is specified as some function of the movement, and the movement with the lowest cost is executed. In the same way that being able to rank different routes from home to work allows us to select a particular route 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, therefore, an elegant framework for dealing with just such a selection

problem and can, therefore, translate from high-level tasks into detailed motor programs (Bryson and Ho 1975). While optimal control can be motivated from the point of view of reducing redundancy, one should always take into account the ultimate evolutionary role of behavior. From an evolutionary point of view the purpose of action is to maximize the chances of passing on genetic material. Clearly some forms of action are more likely to lead to passing on genetic material, and the brain may have learned to indirectly represent this through cost functions ranking actions. 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.

Flash & Hogan (1985) and Uno and colleagues (1989) proposed optimal control models of movement based on maximizing smoothness of the hand trajectory and of the torque commands respectively. Although these models have been successful at reproducing a range of empirical data it is unclear why smoothness is important, and how it is measured by the CNS over the movement. Moreover, these models are limited to a single motor system such as the arm. Harris & Wolpert (1998) have proposed an alternative cost which provides a unifying model for goal-directed eye and arm movements. This model assumes that there is noise in the motor command and that the amount of noise scales with the magnitude of the motor command. In the presence of such signal-dependent 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 the way aspects of the distribution are penalized, and it is this which 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 1 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 is simple for the CNS to measure. Recently, Todorov and Jordan (2002) have shown that optimal feedback control in the presence of signal-dependent noise may form a general strategy for movement production. This model suggests that parameters that are relevant to achieving the task are controlled at the expense of increase in variance in task-irrelevant parameters. For example, in a tracking movement with the hand the variability of the shoulder, elbow and wrist joints may each be high, but by controlling correlations between them, the hand variability is kept low. Moreover, the optimal feedback control model shows that control can be achieved without the need for the CNS to specify a desired trajectory, such as a time series of desired hand positions or velocities.

Figure 1 near hear

State Estimation and Prediction

For the CNS to implement any form of control, it needs to know the current state of the body. However, the CNS faces two problems. First, considerable delays exist in the transduction and transport of sensory signal to the CNS. Second, the CNS must estimate the state of the system from the sensory signals which may be contaminated by noise and may only provide partial information as to the state. For example, consider a tennis ball we have just hit. If we simply used the retinal location of the ball 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 predictive model. The relationship between our motor commands and the consequences are governed by the physics of musculoskeletal system and outside environment. Therefore, to make such a prediction requires a model of this transformation. Such a system is termed an *internal forward model* as it models the causal or forward relationship between actions and their consequences. The term *internal* is used to emphasize that this model is internal to the CNS. 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. Second, components of the ball's state such as its spin cannot be observed easily. However, the spin can be estimated using sensory information integrated over time. The ball's spin will influence its path. By observing the position of the ball over time, an estimate of its spin can be obtained. The estimate from sensory feedback can be improved by incorporating information based on the forward model's predictions (even in a system with no delays).

This combination, using sensory feedback and forward models to estimate the current state is known as an *observer* (Goodwin and Sin 1984). The major objectives of the observer are to compensate for the delays in the sensorimotor system and to reduce the uncertainty in the state estimate which arises due to noise inherent in both the sensory and motor signals. For a linear system, the Kalman filter is the optimal observer in that it produces estimates of the state with the least squared error (Figure 2). Such a model has been supported by empirical studies examining estimation of hand position (Wolpert et al. 1995), posture (Kuo 1995) and head orientation (Merfeld et al. 1999).

Figure 2 near here

Using the observer framework it is a simple computational step from estimating the current state, to predicting future states and sensory feedback. Such prediction have many potential benefits (Wolpert and Flanagan 2001). 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 shown to rely on prediction (for a review see Johansson and Cole 1992). Modeling the performance of subjects who were asked to balance a pole on their fingertip has also provided evidence for predictive models. Examining a variety of control schemes, Mehta and Schaal (2002) concluded, through a process of elimination, that a forward predictive model was likely to be employed.

A sensory prediction can be derived from the state prediction and used to cancel the sensory effects of movement, that is reafference. 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, and relies on a cerebellum-like structure (e.g. Bell et al. 1997). In primates, neurophysiological studies by Duhamel and colleagues (1992) have shown 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 or force, is felt less intensely when it is self-applied. It has been shown that the reduction of the felt intensity of self-applied tactile stimuli critically depends upon on the precise spatio-temporal alignment between the predicted and actual sensory consequences of the movement (Blakemore et al. 1999).

Motor Command Generation

In general the CNS can employ two distinct strategies to generate actions. One strategy is to represent the muscle activations or forces required to compensate for the dynamics of the body or an externally imposed perturbation. This compensation can be achieved by a system which can map desired behavior into the motor commands required to achieve the behavior. Such a system is termed an *inverse model* as it inverts the relationship of the motor system which converts motor command to inputs. When a perfect inverse model is cascaded with the motor system it should produce an identity mapping in that the actual outcome should match the desire outcome. Therefore, to learn model-based compensations for the dynamics of objects we interact with, our CNS needs to learn *internal models* of these objects. An alternative to this model-based compensation is to use co-contraction of the muscles to increased the stiffness of the arm, thereby reducing the displacement caused by an external or inter-segmental forces (Fel'dman 1966 ; Bizzi et al. 1984; Hogan 1984).

Both forms of compensation to perturbations are seen experimentally when subjects are exposed to novel force-fields (Figure 3). By force field we mean a force usually generated by a robotic interface that is related to the state of the hand such as its position. When reaching in a predictable force-field the CNS tends to employ a low-stiffness strategy and learns to represent the compensatory forces. Early in learning the stiffness of the arm reduces systematically as these compensatory responses are learned (Shadmehr and Mussa-Ivaldi 1994 ; Nezafat et al. 2001; Wang et al. 2001). When manipulating an external object with internal degrees of freedom, like a mass-spring system, people also employ low-stiffness control (Dingwell et al. 2002). However, in several situations it is not possible to reliably predict the forces the hand will experience, and therefore model-based compensation is difficult. For example, when drilling into a wall with a power drill, the aim is to maintain the drill bit perpendicular to the wall while applying an orthogonal force. This situation is inherently unstable in that any deviations from orthogonality lead

to forces which destabilize the posture (Rancourt and Hogan 2001). In this situation the stiffness of the hand can be increased in all directions thereby stabilizing the system. Burdet et al (2001) have used an analogous task in which the instability was present in only one direction (shown schematically in Figure 3 right). Subjects reached in a force field in which any deviation of the hand from the straight line between starting point and target was exacerbated by a force perpendicular to the line. They showed that subjects tailored the stiffness of the hand to match the requirements of the task, stiffening the hand only in the perpendicular direction. This is the first demonstration that stiffness can be controlled independently in different directions. Therefore, it seems that the CNS employs both high- and low-stiffness control strategies with the high-stiffness control reducing the effect of any perturbations that a compensation mechanism can not represent.

Figure 3 near here

Bayesian 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 for such an estimation problem is the Bayesian approach which can be used to estimate probabilities for each possible context. The probability of each context can be factored into two terms, the likelihood and the prior. The likelihood of a particular context is the probability of receiving the current sensory feedback given the hypothesized 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 parallel (Wolpert and Kawato 1998; Haruno et al. 2001). Each is tuned to one context and estimates the relative likelihood of its context. This array of models therefore acts as a set of hypothesis testers. The prior contains information about the structured way contexts change over time and how likely a context is prior to 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 a schematic example of picking up what appears to be a full milk carton, which is in reality empty. This shows how the predictive models correct on-line for erroneous priors which initially weighted output of the controller for a full milk carton more than that for an empty. 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 (Figure 2 & 4), they both can be considered on-line ways of doing Bayesian inference in an uncertain environment.

Figure 4 near here

The 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 context (Eskandar and Assad 1999), as well as representing the likelihood of the sensory feedback given the context (Kim and Shadlen 1999). An elegant example of context estimation has been provided by Cohn and colleagues (Cohn et al. 2000). When subjects make a reaching movement while rotating their torso, they compensate for the velocity dependent Coriolis forces arising from the rotation, which 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 Coriolis force. This leads to misreaching which over subsequent movements reduces as the sensory consequences of the expected Coriolis force are not experienced.

Motor 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 properties of the sensorimotor system. 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 which lead to such predictive learning in the cerebellum-like structure of electric fish has recently been partially elucidated (Bell et al. 1997).

Acquiring an inverse internal model through motor learning is generally a difficult task. This is because the appropriate training signal, the error in the output of the inverse model, that is the motor command error, is not directly available. When we fail to sink a putt no-one tells us how our muscle activations should change to achieve the task. 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. An original proposal was to use direct inverse modeling (Widrow and Stearns 1985; Miller 1987; Kuperstein 1988; Atkeson and Reinkensmeyer 1988) in which an inverse model could be learned during a motor babbling stage. This controller would simply observe motor commands and sensory outcomes during babbling and try to learn how outcomes (as inputs to the inverse model) map to the motor commands that caused this outcome.. For linear systems such a process can be shown to usually converge to correct parameter estimates (Goodwin and Sin 1984). However, there are several problems with such a system. First it is not *goal directed*; that is, it is not sensitive to particular output goals (Jordan and Rumelhart 1992). The learning process samples randomly during babbling

and there is no guarantee that it will sample appropriately for a give task. Second, the controller is trained “off-line”, that is the input to the controller for the purposes of training is the actual output, not the desired output. For the controller to actually participate in the control process, it must receive the desired plant output as its input. The direct inverse modeling approach therefore requires a switching process; the desired plant output must be switched in for the purposes of control and the actual plant output must be switched in for the purposes of training. Finally, for nonlinear systems a difficulty arises that is related to the general “degrees-of-freedom problem” in motor control (Bernstein 1967). The problem is due to a particular form of redundancy in nonlinear systems (Jordan 1992). In such systems, the “optimal” parameter estimates in fact may yield an incorrect controller. Because of the redundancy in the motor system there may be many motor commands that lead to the same outcome and during direct inverse learning the system may see the same outcome many times caused by different motor commands. Most learning systems when trying to learn to map a single outcome into the multiple motor commands which lead to this outcome will finally map this outcome to the average of all these motor commands. However, for nonlinear systems it is rarely the case that the average of all these motor command will lead to the same outcome, and therefore direct inverse modeling fails for such nonlinear systems.

Two learning mechanism have been proposed to overcome these limitations. Kawato and colleagues (1987; 1992) have proposed an ingenious solution to this problem, feedback-error-learning (Figure 5). They suggest that a hard-wired, but not perfect, feedback controller exists which computes a motor command based on the discrepancy between desired and estimated state. The motor command is the sum of the feedback controller motor command and the output of an adaptive inverse model. They reasoned that if the feedback controller ended 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 would be performing perfectly. Based on this they regarded the output of the feedback controller as the error signal, and used it to train the inverse model, an approach which is highly successful. Therefore, feedback error learning makes use of a feedback controller to guide the learning of the feedforward controller. The feedforward controller is trained “on-line”, that is it is used as a controller while it is being trained. is goal directed. Neurophysiological evidence (Shidara et al. 1993) supports this learning mechanism within the cerebellum for the simple reflex eye movement, called the ocular following response. The suggestion is that the cerebellum constructs an inverse model of the eye's dynamics.

Figure 5 near here

Another solution is to use distal supervised learning (Jordan and Rumelhart 1992). In distal supervised learning, the controller is learned indirectly, through the intermediary of a forward model of the motor apparatus. The forward model must itself be learned from observations of the inputs and outputs of the system. The distal supervised learning approach is therefore composed of two interacting processes, one process in which the forward model is learned, and another process in which the forward model is used in the

training of the controller. The controller and the forward model are joined together and are treated as a single *composite learning system*. If the controller is to be an inverse model, then the composite learning system should be an identity transformation (i.e., a transformation whose output is the same as its input). This suggests that the controller can be trained indirectly by training the composite learning system to be an identity transformation. During this training process, the parameters in the forward model are held fixed. Thus the composite learning system is trained to be an identity transformation by a constrained learning process in which some of the parameters inside the system are held fixed. By allowing only the controller parameters to be altered, this process trains the controller indirectly.

Distal supervised learning and other models (Haruno et al. 2001) have suggested that we use a forward model to train a controller. In an experiment designed to simultaneously assess both forward and inverse model learning subjects were required to move an object along a straight line, while the load on the object was varied during the trial (Flanagan et al. 2003). Over repeated trials, the subjects learned to compensate for the load so that they could produce a straight trajectory. The hand trajectory was used to measure how quickly subjects learned to control the movement, whereas prediction was measured by looking at changes in grip force. In early trials, grip force was changed reflexively as the hand path (and therefore the load force) was perturbed, but subjects quickly learned to alter their grip force predictively. By contrast, it took many trials for them to learn to control the load. This suggests that we learn to predict the consequences of our actions before we learn to control them.

Unifying principles

Computational approaches have started to provide unifying principles for motor control. Several common themes have already emerged in this review. 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. It is our belief 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

This work was supported by grants from the Wellcome Trust, Gatsby Charitable Foundation and the Human Frontiers Science Programme.

Figure 1. A schematic of the Task Optimization in the Presence of Signal-dependent noise (TOPS) model of Harris & Wolpert. Shown are 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, due to noise on the motor commands, they have different final distributions. 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 which optimizes the statistics is selected.


Figure 2. 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. Using a copy of the motor command, that is efference copy, and a model of the dynamics the current state distribution is predicted from this previous state. In general, the uncertainty is increased. 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. Although there are delays in sensory feedback which must be compensated, they have been omitted from the diagram for clarity.


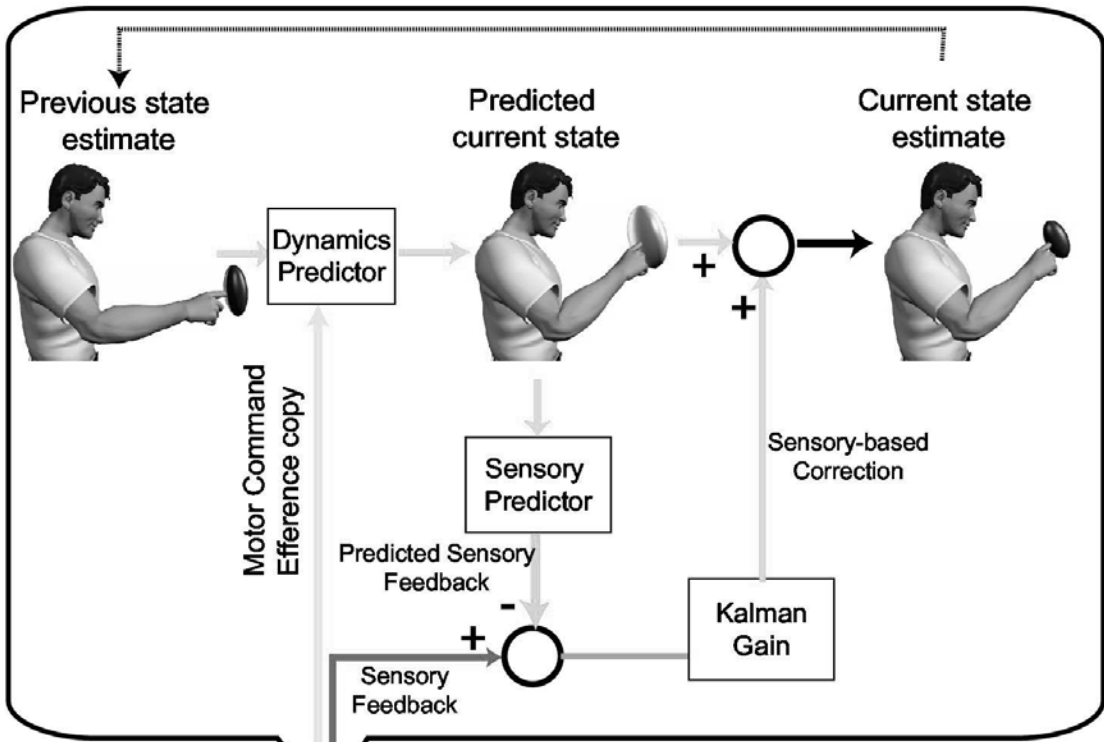
Figure 3. Schematic of two strategies for control when learning to move in a force field. Subjects reach between the two circular targets under a force field generated by a robot (not shown), that depends on the position of the hand. The force experienced for different positions are shown by the arrows. a) Under a stable and predictable force field acting to the left the subjects will learn to produce a straight line movement. If the field is unexpectedly turned off for a movement subjects will show an after-effect (black trajectory) reflecting the compensation they are producing in their motor command to counteract the field. b) The field is unstable as any deviation for a straight hand path will generate force acting in the same direction. Subjects learn to move in straight line but show no after-effects on removal of the field. The task is achieved by increasing the stiffness of the arm, but only in the direction of maximum instability. The stiffness ellipse represents restoring force to a step displacement of the hand in different directions (dotted prior to learning and solid after).

Figure 4. A schematic of Bayesian 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

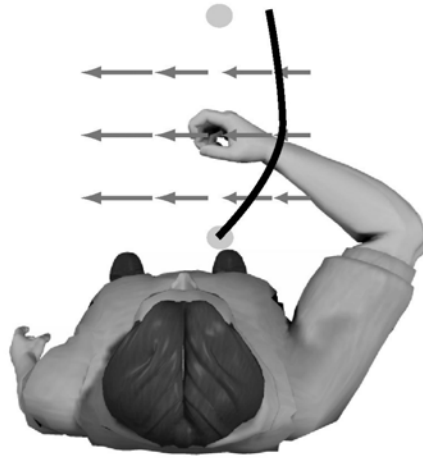
large amount of movement compared to the full carton context. These predictions are compared with actual feedback. As the carton is, in fact, empty the sensory feedback matches the predictions of the empty carton context. This leads to a high likelihood for the empty carton and a low likelihood of the full carton. The likelihoods are combined with the priors using Bayes rule to generate the final (posterior) probability of each context.

Figure 5. A schematic of feedback-error learning. The aim is to learn an inverse model which 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 which is added to the feedforward motor command generated by the inverse model. If the feedback motor command goes to zero then the state error will, in general, 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.




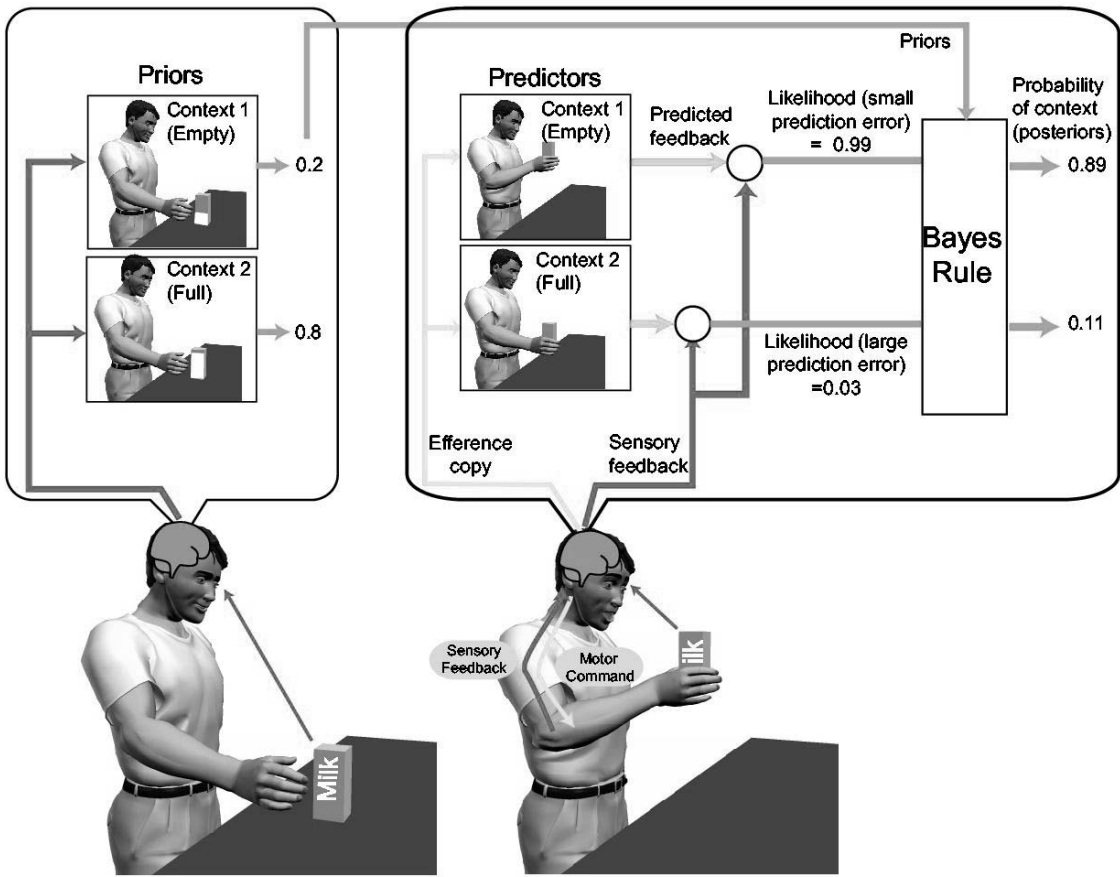



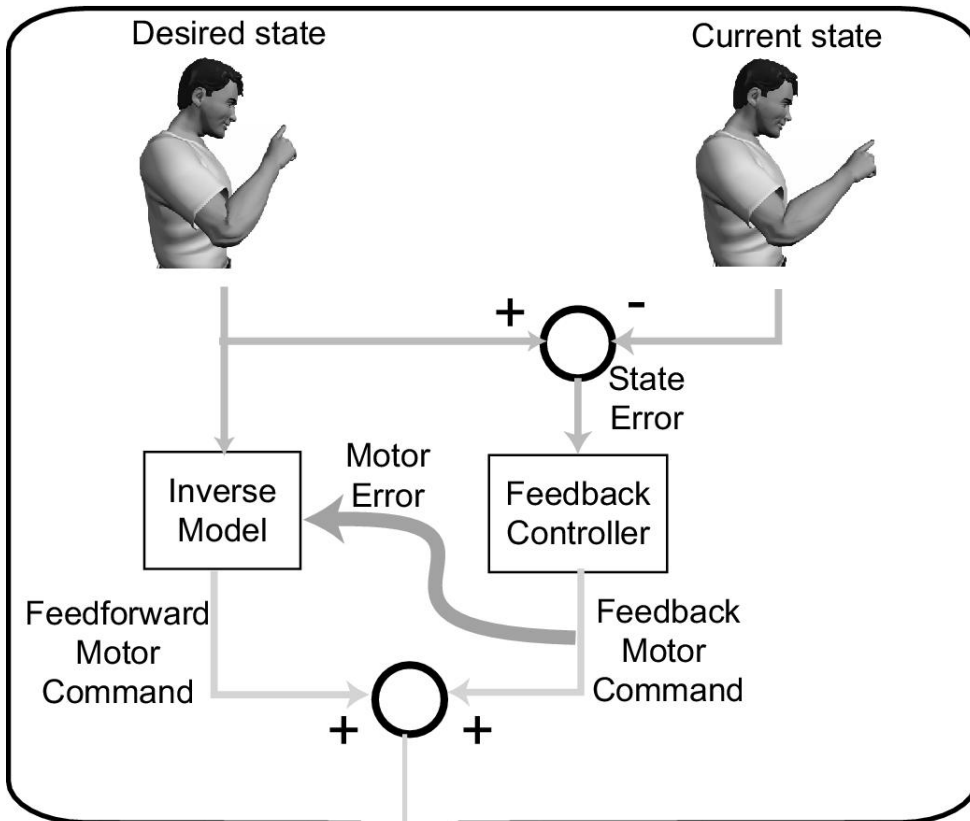
Model-based compensation



Task matched stiffness







References

- Atkeson CG, Reinkensmeyer DJ (1988) Using associative content-addressable memories to control robots. . In: IEEE Conference on Decision and Control
- Bell CC, Han VZ, Sugawara Y, Grant K (1997) Synaptic plasticity in a cerebellum-like structure depends on temporal order. *Nature* 387 278-281
- Bernstein N (1967) *The Coordination and Regulation of Movements*. Pergamon, London
- Bizzi E, Accornerro N, Chapple B, Hogan N (1984) Posture control and trajectory formation during arm movement *J. Neurosci.* 4: 2738-2744
- Blakemore SJ, Frith CD, Wolpert DM (1999) Spatio-temporal prediction modulates the perception of self-produced stimuli. *J Cogn Neurosci* 11: 551-559.
- Bryson AE, Ho YC (1975) *Applied Optimal Control*. Wiley, New York
- Burdet E, Osu R, Franklin DW, Milner TE, Kawato M (2001) The central nervous system stabilizes unstable dynamics by learning optimal impedance. *Nature* 414: 446-449
- Cohn JV, DiZio P, Lackner JR (2000) Reaching during virtual rotation: context specific compensations for expected coriolis forces. *J Neurophysiol* 83: 3230-3240.
- Dingwell JB, Mah CD, Mussa-Ivaldi FA (2002) Manipulating objects with internal degrees of freedom: evidence for model-based control. *J Neurophysiol* 88: 222-235
- Duhamel JR, Colby CL, Goldberg ME (1992) The updating of the representation of visual space in parietal cortex by intended eye movements. *Science* 255: 90-92
- Eskandar EN, Assad JA (1999) Dissociation of visual, motor and predictive signals in parietal cortex during visual guidance. *Nature Neurosci.* 2: 88-93
- Fel'dman AG (1966) 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
- Flanagan JR, Vetter P, Johansson RS, Wolpert DM (2003) Prediction precedes control in motor learning. *Curr Biol* 13: 146-150
- Flash T, Hogan N (1985) The co-ordination of arm movements: An experimentally confirmed mathematical model *J. Neurosci.* 5: 1688-1703
- Ghahramani Z, Wolpert DM (1997) Modular decomposition in visuomotor learning. *Nature* 386: 392-395
- Gomi H, Kawato M (1996) Equilibrium-point control hypothesis examined by measured arm stiffness during multijoint movement. *Science* 272: 117-120
- Goodwin GC, Sin KS (1984) *Adaptive Filtering Prediction and Control*. Prentice-Hall, Englewood Cliffs, NJ
- Harris CM, Wolpert DM (1998) Signal-dependent noise determines motor planning. *Nature* 394: 780-784
- Haruno M, Wolpert DM, Kawato M (2001) Mosaic model for sensorimotor learning and control. *Neural Comput* 13: 2201-2220.
- Hogan N (1984) An organizing principle for a class of voluntary movements *J. Neurosci.* 4: 2745-2754
- Johansson RS, Cole KJ (1992) Sensory-motor coordination during grasping and manipulative actions. *Curr. Opin. Neurobiol* 2: 815-823

- Jordan MI (1992) Constrained supervised learning. *J. of Mathematical Psychology* 36: 396-425
- Jordan MI, Rumelhart DE (1992) Forward models: Supervised learning with a distal teacher *Cognitive Science* 16: 307-354
- Kawato M, Furawaka K, Suzuki R (1987) A hierarchical neural network model for the control and learning of voluntary movements *Biological Cybernetics* 56: 1-17
- Kawato M, Gomi H (1992) The cerebellum and VOR/OKR learning models. *Trends in Neurosciences* 15: 445-453
- Kim J, Shadlen MN (1999) Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nature Neurosci.* 2: 176-185
- Kuo AD (1995) An Optimal-Control Model for Analyzing Human Postural Balance. *Ieee Transactions on Biomedical Engineering* 42: 87-101
- Kuperstein M (1988) Neural model of adaptive hand-eye coordination for single postures *Science* 239: 1308-1311
- Mehta B, Schaal S (2002) Forward models in visuomotor control. *J Neurophysiol* 88: 942-953
- Merfeld DM, Zupan L, Peterka RJ (1999) Humans use internal model to estimate gravity and linear acceleration. *Nature* 398: 615-618
- Miller WT (1987) Sensor-based control of robotic manipulators using a general learning algorithm. *IEEE J. of Robotics and Automation* 3: 157-165
- Morasso P (1981) Spatial control of arm movements *Exp. Brain Res.* 42: 223-227
- Nezafat R, Shadmehr R, Holcomb HH (2001) Long-term adaptation to dynamics of reaching movements: a PET study. *Exp Brain Res* 140: 66-76
- Rancourt D, Hogan N (2001) Stability in force-production tasks. *J Mot Behav* 33: 193-204
- Shadmehr R, Mussa-Ivaldi F (1994) Adaptive representation of dynamics during learning of a motor task. *J. Neurosci.* 14:5: 3208-3224
- Shidara M, Kawano K, Gomi H, Kawato M (1993) Inverse-dynamics encoding of eye movement by Purkinje cells in the cerebellum. *Nature* 365: 50-52
- Todorov E, Jordan MI (2002) Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5: 1226-1235
- Uno Y, Kawato M, Suzuki R (1989) Formation and control of optimal trajectory in human multijoint arm movement. Minimum torque-change model. *Biol Cybern* 61: 89-101
- Wang T, Dordevic GS, Shadmehr R (2001) Learning the dynamics of reaching movements results in the modification of arm impedance and long-latency perturbation responses. *Biol Cybern* 85: 437-448
- Widrow B, Stearns SD (1985) Adaptive signal processing. Prentice-Hall, Englewood Cliffs, NJ
- Wolpert DM, Flanagan JR (2001) Motor prediction. *Curr Biol* 11: R729-732.
- Wolpert DM, Ghahramani Z, Jordan MI (1995) An internal model for sensorimotor integration. *Science* 269: 1880-1882
- Wolpert DM, Kawato M (1998) Multiple paired forward and inverse models for motor control. *Neural Networks* 11: 1317-1329