elicited by syntactic anomaly *J. Mem. Lang.* 31, 785–806
16 Osterhout, L. and Holcomb, P.J. (1993) Event-related potentials and syntactic anomaly: evidence of anomaly detection during the perception of continuous speech *Lang. Cogn. Proc.* 8, 413–438
17 Osterhout, L., Holcomb, P.J. and Swinney, D.A. (1994) Brain potentials elicited by garden-path sentences: evidence of the application of verb information during parsing *J. Exp. Psychol. Learn. Mem. Cogn.* 8, 507–526
18 Osterhout, L. and Mobley, L.A. (1995) Event-related brain potentials elicited by failure to agree *J. Mem. Lang.* 34, 739–773
19 Osterhout, L. and Nicol, J. On the distinctiveness, independence, and time course of the brain responses to syntactic and pragmatic anomalies *Lang. Cogn. Proc.* (in press)
20 Münte, T.F., Heinze, H-J. and Mangun, G.R. (1993) Dissociation of brain activity related to syntactic and semantic aspects of language *J. Cogn. Neurosci.* 5, 335–344
21 Rösler, F. et al. (1993) Event-related brain potentials while encountering semantic and syntactic constraint violations *J. Cogn. Neurosci.* 5, 345–362
22 Neville, H.J., Mills, D.J. and Lawson, D.S. (1992) Fractionating language: different neural subsystems with different sensitive periods *Cortex* 2, 244–258
23 Van Petten, C. and Kutas, M. (1991) Influences of semantic and syntactic context on open- and closed-class words *Mem. Cogn.* 19, 95–112
24 Osterhout, L. On the brain response to syntactic anomalies: manipulations of word position and word class reveal individual differences *Brain Lang.* (in press)
25 Just, M.A. and Carpenter, P. (1992) A capacity theory of comprehension: individual differences in working memory *Psychol. Rev.* 99, 122–149
26 Donchin, E. (1981) Surprise!...Surprise? *Psychophysiology* 18, 493–513
27 Johnson, R., Jr (1993) On the neural generators of the P300 component of the event-related potential *Psychophysiology* 30, 90–97
28 Osterhout, L. et al. (1996) On the language specificity of the brain response to syntactic anomalies: is the syntactic positive shift a member of the P300 family? *J. Cogn. Neurosci.* 8, 507–526
29 Coles, M.G.H., Gratton, G. and Donchin, E. (1988) Detecting early communication: using measures of movement-related potentials to illuminate human information processing *Biol. Psychol.* 26, 69–89
30 Van Turennout, M., Hagoort, P. and Brown, C. (1997) Electrophysiological evidence on the time course of semantic and phonological processes in speech production *J. Exp. Psychol. Learn. Mem. Cogn.* 23, 787–806
31 Swaab, T., Brown, C. and Hagoort, P. (1997) Spoken sentence comprehension in aphasia: event-related potential evidence for a lexical integration deficit *J. Cogn. Neurosci.* 9, 39–66
32 Hagoort, P., Brown, C. and Swaab, T. (1996) Lexical-semantic event-related potential effects in patients with left hemisphere lesions and aphasia, and patients with right hemisphere lesions without aphasia *Brain* 119, 627–649
33 Neville, H.J. (1991) Neurobiology of cognitive and language processing: effects of early experience, in *Brain Maturation and Cognitive Development: Comparative and Cross-cultural Perspectives* (Gibson, K.R. and Peterson, A.C., eds), pp. 355–380, de Gruyter
34 Weber-Fox, C. and Neville, H.J. (1996) Maturational constraints on functional specializations for language processing: ERP and behavioral evidence in bilingual speakers *J. Cogn. Neurosci.* 8, 231–256
35 King, J.W. and Kutas, M. (1995) Who did what and when? Using word- and clause-level ERPs to monitor working memory in reading *J. Cogn. Neurosci.* 7, 376–395
36 Osterhout, L., Bersick, M. and McLaughlin, J. (1997) Brain potentials reflect violations of gender stereotypes *Mem. Cogn.* 25, 273–285
37 Nobre, A.C., Allison, T. and McCarthy, G. (1994) Word recognition in the human inferior temporal lobe *Nature* 372, 260–262

Computational approaches to motor control

Daniel M. Wolpert

This review will focus on four areas of motor control which have recently been enriched both by neural network and control system models: motor planning, motor prediction, state estimation and motor learning. We will review the computational foundations of each of these concepts and present specific models which have been tested by psychophysical experiments. We will cover the topics of optimal control for motor planning, forward models for motor prediction, observer models of state estimation and modular decomposition in motor learning. The aim of this review is to demonstrate how computational approaches, as well as proposing specific models, provide a theoretical framework to formalize the issues in motor control.

D.M. Wolpert is at the Sobell Department of Neurophysiology, Institute of Neurology, Queen Square, London, UK WC1N 3BG.

tel: +44 171 837 3611 x4183
fax: +44 171 813 3107
e-mail: wolpert@hera.ucl.ac.uk

This review will focus on several basic theoretical issues in motor control as well as supporting experimental studies. While many of the concepts discussed are applicable to all areas of motor control, including eye movements, speech production and posture, we will focus on arm movements as an illustrative system. From an engineering perspective the arm can be considered as a system whose inputs are the motor commands emanating from the controller within the central nervous system (see Fig. 1). In order to determine the behaviour of the arm in response to this input an additional set of variables, called state variables, must also be known. For example, in a robotic model of the arm the motor command signals the torques generated around the joints and the state variables could be the joint angles and

 1364-6613/97/$17.00 PII: S1364-6613(97)01070-X

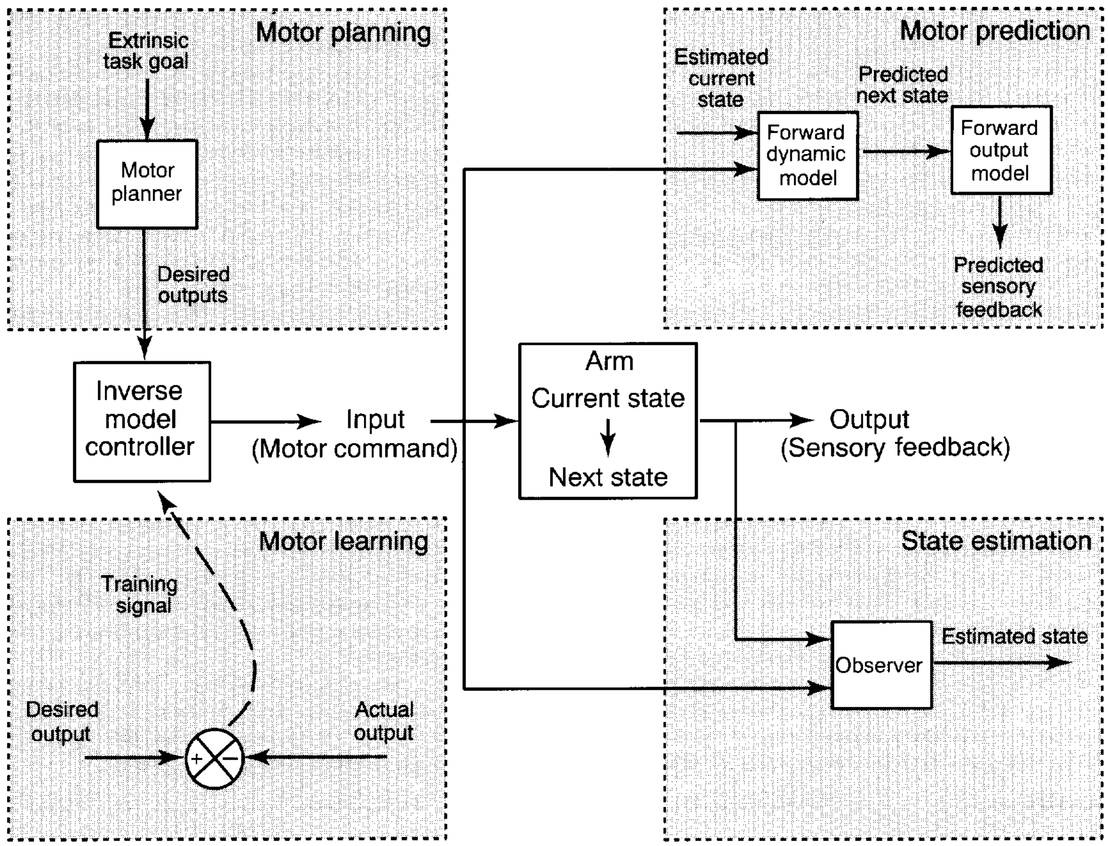

Fig. 1 Motor control. The motor system is shown schematically along with the four themes of motor control that will be reviewed (see text for details). The motor system (centre) has inputs, motor commands, which causes it to change its state and produce an output, sensory feedback. For clarity not all lines are shown.

angular velocities. Taken together, the inputs and the state variables are sufficient to determine the future behaviour of the system.The controller does not usually have direct access to the state of the system but has access to some function of the state, such as sensory feedback, which forms the output of the system.

This review will consider four issues which arise in motor control and each is represented schematically in Fig. 1. The first is motor planning, which I consider to be the computational process by which the desired states or outputs of the system are specified given an extrinsic task goal. Second, I will explore the notion of motor prediction, and the use of forward models which predict the behaviour of the arm given its input. Third, I will consider state estimation, which is a sensorimotor integration process by which the hidden state of the motor system can be estimated by monitoring both its inputs and outputs. I will show how motor prediction can play a fundamental role in such state estimation. Lastly, I will consider motor learning, focusing on modularity of learning.

Motor planning

The computational problem of motor planning arises from a fundamental property of the motor system – the reduction in the degrees of freedom that occurs during the transition from neural commands through muscle activation to movement kinematics[1] (Fig. 2). Even for the simplest of tasks, such as moving the hand to a target location, there are an infinite number of possible paths that the hand could move along and for each of these paths there is an infinite number of velocity profiles (trajectories) that the hand could follow. Having specified the hand path and velocity, each location of the hand along the path could be achieved by multiple combinations of joint angles. Moreover, due to the overlapping actions of muscles and the ability to co-contract, each arm configuration could be achieved by many different muscle activations. Motor planning can therefore be considered as the computational process of selecting a single solution or pattern of behaviour at all levels within this motor hierarchy (Fig. 2) from the many alternatives that may be consistent with the goal of the task.

One computational framework that is natural for such a selection process is optimal control, in which a cost function is chosen in order to evaluate quantitatively the performance of the system under control[2]. The cost function is usually defined as the integral of an instantaneous cost, over a certain time interval, and the aim is to minimize the value of this cost function. Every possible solution, that is every possible movement, has an associated cost and the solution with the lowest cost is selected as the plan. Within this framework the cost function is a mathematical means for specifying the plan. The variables that appear in the cost function, and that are therefore planned, determine the pattern of behaviour that is observed.

While many possible cost functions have been examined[3] there are two main classes of model that have been

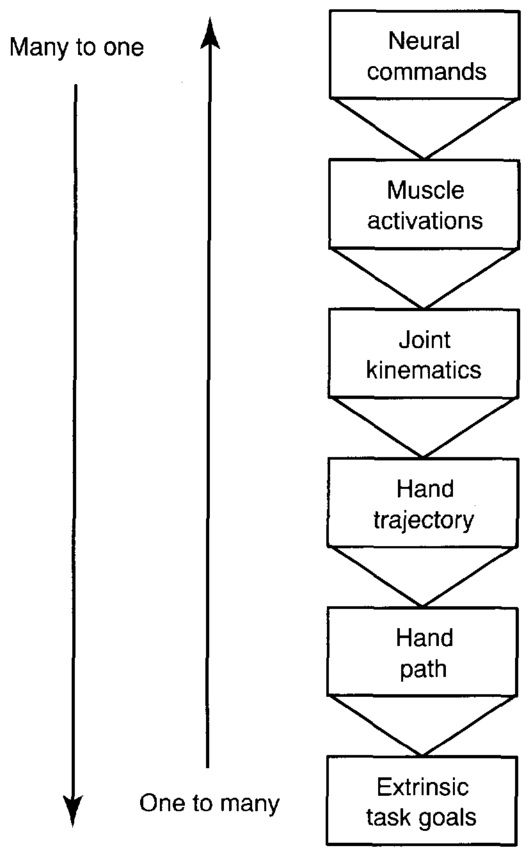

Fig. 2 The compulation problem of motor planning. The levels in the motor hierarchy are shown with the triangles between the levels indicating the reduction in the degrees of freedom between the higher and lower levels. Specifying a pattern of behaviour at any level completely specifies the patterns at the level below (many to one: many patterns at the higher level correspond to one pattern at the lower) but is consistent with many patterns at the level above (one to many). Planning can be considered as the process by which particular patterns, consistent with the extrinsic task goals, are selected at each level.

proposed for point-to-point movements – kinematic and dynamic based models. The cost function in kinematic based models contains only the geometrical and time-based properties of the motion, and the variables of interest are the positions (e.g. joint angles or hand Cartesian coordinates) and their corresponding velocities, accelerations and higher derivatives. Based on the observation that point-to-point movements of the hand are smooth when viewed in a Cartesian framework, it was proposed that the squared first derivative of Cartesian hand acceleration or 'jerk' is minimized over the movement[4,5]. This minimum jerk hypothesis produces a unique solution, given the movement duration and suitable boundary conditions of the initial and final position and velocity, and can be formulated into an on-line feedback rule[6]. The model predicts straight-line Cartesian hand paths with bell-shaped velocity profiles that are consistent with the empirical data for rapid movements made without accuracy constraints[5,7–9].

The cost function in dynamic based models depends on the dynamics of the arm, and the variables of interest include joint torques, forces acting on the hand and the muscle commands. Several models have been proposed in which the cost function depends on dynamic variables such as torque change, muscle tension or motor command[10,11]. One critical difference between the kinematic and dynamic based models is the degree to which planning and execution processes can be separated. The specifications of the movement in kinematic models, such as minimum jerk, are the positions and velocities of the arm as a function of time. Therefore, a separate process is required to achieve these specifications and this model is a hierarchical, serial plan and execute model. In contrast, the solutions to dynamic models, such as minimum torque change, are the motor commands required to achieve the movement and therefore planning and execution are not conceived as separate processes.

Although examination of natural movements has not resolved the debate between kinematic and dynamic based cost functions, the predictions of the two classes of models are different under visual and force field perturbations. Thus, if the visual feedback of the hand path is altered, the cost in kinematic based models, which depend on the perceived movement of the limb, increases. However, provided the visuomotor perturbation is chosen so as to decay to zero at both the start and end of the movement[12], the target can still be reached by using the same series of motor commands such that dynamic based cost, which depends on these motor commands is, not increased. Kinematic based models, therefore predict that under these conditions the actual hand path will change so as to bring the perceived path back to the original path, which has the lowest cost. In contrast, dynamic models predict no such adaptation. However, in the presence of a force field which alters the dynamics of the arm, kinematic based models predict that the normal kinematics of movement will be regained to, once again, minimize the cost. Moreover, under these conditions, dynamic based models predict a new solution (and therefore a new hand path) to the optimization process, that takes account of the new motor commands required to make the movement in the force field. Evidence from both perturbed visual feedback studies[12–14] and from force field studies[15,16] support a kinematic based sequential plan and execute strategy.

However, studies of more complex movements around an obstacle suggest that knowledge of the dynamics of the arm is used in planning. Indeed, subjects tend to select their movement paths so as to ensure that their closest point of approach to an obstacle is on an axis where the arm is most inertially stable[17]. Similarly, external movement constraints may affect kinematics, since there are differences between the hand paths used when the hand is free to move compared to a constrained condition in which it is required to move along the surface of a table. The paths for unconstrained movement are more curved that those made along the tabletop[18]. This suggests that the nature of the interaction with the environment can have a significant effect on movement kinematics. Whether these findings can be incorporated into an optimal control framework is still an open question.

The minimum jerk model has recently been used in an attempt to find a unified framework within which to understand two properties of trajectory formation, local isochrony and the two-thirds power law. Whereas global isochrony applies to the observation that the average velocity of movement increases with the movement distance thereby maintaining a constant movement duration, local

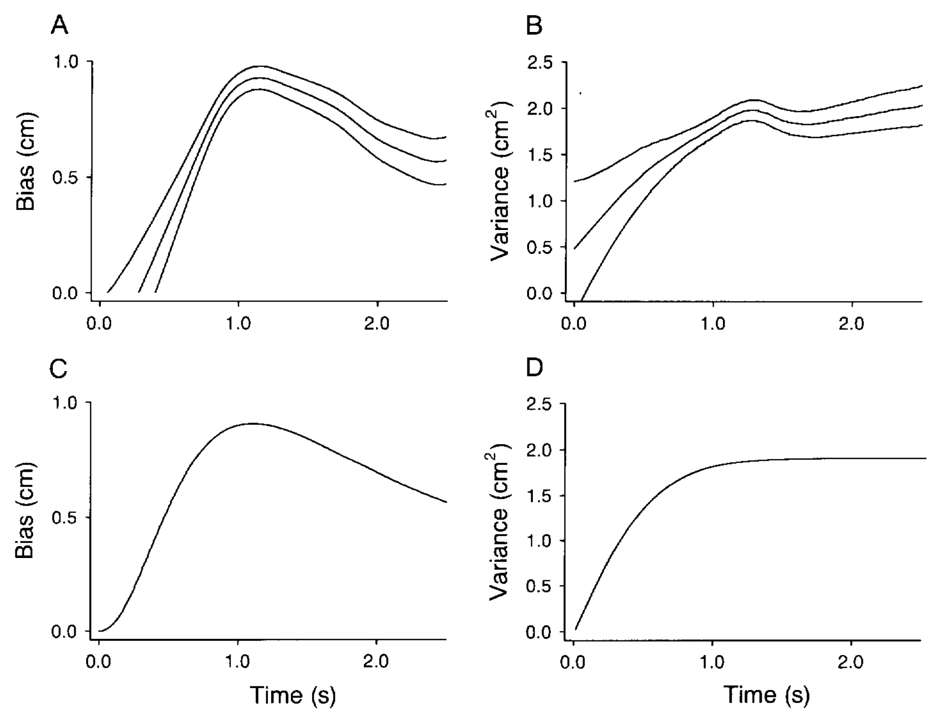

Fig. 3 Sensorimotor integration. The propagation of the bias **(A)** and variance **(B)** of the state estimate against movement duration is shown as a spline fit with outer standard error lines to the data from eight subjects[31]. Also shown are the bias **(C)** and variances **(D)** from a simulation of the Kalman filter (see Box 1 for details). Modified from Ref. 31.

isochrony refers to the subunits of movement. For example, if subjects are required to trace out a figure eight in which the two loops are of unequal size, the time to traverse each loop is approximately equal. By approximating the solution of the minimum jerk when the path is constrained (only the velocity along the path could be varied) local isochrony becomes an emergent property of the optimization of jerk[19]. The two-thirds power law, $A \propto C^{\beta}$ $(\beta \approx \frac{2}{3})$, is based on the observation of the relationship between path curvature (C) and hand angular velocity (A) during drawing or scribbling[20] (for a more general formulation of the law see Ref. 21). It has been shown that the solution of the minimum jerk along a constrained path approximates the solution given by the two-thirds power law[19]. One area of debate is the extent to which the two-thirds power law is a manifestation of a plan rather than a control constraint. Based on a simple model of control it has been shown that the two-thirds power law could be an emergent property of the muscles' visco-elastic properties[22]. However, one feature which has not yet been explained by such emergent property models is the fact that the exponent of the power law, β, changes systematically through development, from a value of 0.77 at age 6 to an adult value of 0.66 ($\frac{2}{3}$) at around 12 years[21].

Motor prediction

The possibility that we predict the consequences of our own action using an internal model of the motor system has emerged as an important theoretical concept in motor control[23]. Such models have become known as forward models as they capture the forward or causal relationship between inputs to the system, such as the arm, and the outputs. Forward dynamic models, for example, predict the next state (for example, the position and velocity) given the current state and the motor command whereas forward output models predict the sensory feedback given this estimated state (Fig. 1). This is in contrast to inverse models which invert the system by providing the motor command which will cause a desired change in state. As inverse models produce the motor command required to achieve some desired result they have a natural use as a controller (see Motor Learning).

Forward models have several possible uses. Forward models are key components in systems that use a copy of the motor command, an efference copy, to anticipate and cancel the sensory effects of a given movement. This reafference process has been extensively investigated in eye movement control (see Ref. 24 for a review). By using such a system for the control of arm movements it is possible to cancel out the effects of self-motion on sensation and thereby distinguish between those sensory events that are due to self-produced motion and those caused by the environment, such as contact with objects. Another use of a forward model is to maintain stability in the presence of feedback delays. In most sensorimotor loops the feedback delays are large, and can result in instability when trying to make rapid movements under feedback control. Two strategies can maintain stability during movement with such delays: intermittency and prediction. Intermittency, in which movement is interspersed with rest, is seen in manual tracking and saccadic eye movement. The intermittency of movement allows time for veridical sensory feedback to be obtained (a strategy often used in adjusting the temperature of a shower where the time delays are large). Such intermittency can arise either from a psychological refractory period after each movement[25] or an error deadzone[26] in which the perceived error must exceed a threshold before a new movement is initiated. Alternatively, in predictive control, a forward model is used to provide internal feedback of the predicted outcome of an action which can be used before sensory feedback is available, thereby preventing instability[27]. Forward models can also be of computational use in motor learning. A forward model which captures the relationship between motor commands and outcome can be used to convert errors in outcome into errors in motor commands, thereby providing a suitable signal for learning – a process known as distal supervised learning[28]. Similarly, by predicting the sensory outcome of an action, without actually performing it, a forward model can be used in mental practice to learn to select between possible actions. Finally, a forward model forms an integral component in systems which integrate sensory and motor information in state estimation.

State estimation

Although the state of a system is not directly available to the controller, it is possible to estimate the state indirectly. Such a

Box 1. The Kalman filter model of sensorimotor integration

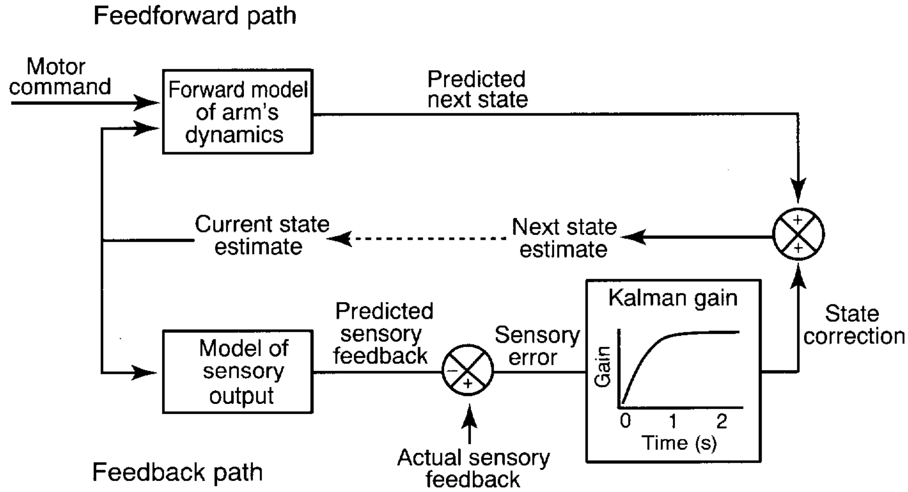

The Kalman filter model of the sensorimotor integration process is based on a formal engineering model from the optimal state estimation field[a]. The Kalman filter is a linear dynamical system that produces an estimate of the location of the hand by using both the motor outflow and sensory feedback in conjunction with a model of the motor system, thereby reducing the overall uncertainty in its estimate[b]. This model assumes that the localization errors arise from two sources of uncertainty, the first from the variability in the response of the arm to the motor command and the second in the sensory feedback given the arm's configuration. The Kalman filter model can be considered as the combination of two processes which together contribute to the state estimate. The first, feedforward process (upper part) uses the efferent outflow along with the current state estimate to predict the next state by simulating the movement dynamics with a forward model. The second, feedback process (lower part) compares the sensory inflow with a prediction of the sensory inflow based on the current state. The sensory error, the difference between actual and predicted sensory feedback, is used to correct the state estimate resulting from the forward model. The relative contributions of the internal simulation and sensory correction processes to the final estimate are modulated by the time varying Kalman gain so as to provide optimal state estimates.

To simulate the experimental data of the sensorimotor integration task (see main text: State Estimation) we modelled the hand as a damped point mass, that was acted on by a force whose temporal profile was chosen to replicate the kinematics of movement[c]. At the start of the movement the subject was given full view of his arm and the initial state estimate of the hand position was set to its veridical value. The Kalman filter with an accurate forward model of the system would be expected to show no bias. To accommodate the observation that the subjects generally tended to overestimate the distance that their arm had moved, we set the gain that couples force to estimated acceleration within the forward model to a value that was larger than its veridical value.

The Kalman filter model demonstrated the two distinct phases of bias propagation observed (Fig. 3C). By overestimating the force acting on the arm the forward model overestimates the distance travelled, an integrative process eventually balanced by the sensory correction. The accuracy of the prediction from the forward model component of the Kalman filter depends on the accuracy of the current state estimate (one of its inputs). Therefore during the early part of the movement, when the current state estimate is accurate, the sensorimotor integration process weights heavily the contribution of the forward model to the final estimate. However, in the later stages of the movement, when the current state estimate is less accurate, the sensory feedback must be relied upon to correct for inaccuracies in the forward model. In the Kalman filter, the relative weighting shifts smoothly from the feedforward process towards the feedback process over the first second of movement and then remains approximately constant, resulting in the asymptote of the variance propagation. The Kalman filter model suggests that the peaking and gradual decline in bias is a consequence of a trade-off between the inaccuracies accumulating in the internal simulation of the arm's dynamics and the feedback of actual sensory information. Figure modified from Ref. c.

References

a Goodwin, G.C. and Sin, K.S. (1984) *Adaptive Filtering Prediction and Control*, Prentice-Hall

b Kalman, R.E. and Bucy, R.S. (1961) New results in linear filtering and prediction *J. Basic Engineering* (ASME) 83D, 95–108

c Wolpert, D.M., Ghahramani, Z. and Jordan, M.I. (1995) An internal model for sensorimotor integration *Science* 269, 1880–1882

state estimator, known as an observer[29], produces its estimate of the current state by monitoring the stream of inputs (motor commands) and outputs (sensory feedback) of the system (Fig. 1). By using both sources of information the observer is able to reduce uncertainty in the state estimate and becomes robust to sensor failure. In addition, as there are delays in sensory feedback, the observer can use the motor command to produce more timely state estimates than would be possible using sensory feedback alone.

Although many studies have examined integration among purely sensory stimuli (for a psychophysical review see Ref. 30) little is known of how sensory and motor information is integrated during movement. When you move your arm in the absence of visual feedback, there are three basic methods the central nervous system can use to obtain an estimate of the current state, the position and velocity, of the hand. The system can make use of sensory inflow (the information available from proprioception), it can make use of motor outflow (the motor commands sent to the arm), or it can combine these two sources of information. While sensory signals can directly cue the location of the hand, motor outflow generally does not. For example, given a sequence of torques applied to the arm (the motor outflow) an internal model of the arm's dynamics is needed to estimate the arm's final configuration. To examine whether an internal model of the arm is used we have studied a sensorimotor integration task in which subjects, after initially viewing their arm in the light, made arm movements in the dark[31]. The subjects' internal estimate of hand location was assessed by asking them to visually localize the position of their hand (which was hidden from view) at the end of the movement. The bias of this location estimate, plotted as a function of movement duration, showed a consistent overestimation of the distance moved (Fig. 3A). This bias showed two distinct phases as a function of movement duration, an initial increase that reached a peak after one second followed by

Box 2. The mixture of experts model of visuomotor learning

Based on the principle of divide-and-conquer, a general computational strategy for designing modular learning systems is to treat the problem as one of combining multiple models, each of which is defined over a local region of the input space. Such a strategy has been introduced in the 'mixture of experts' architecture for supervised learning[a,b]. The architecture involves a set of function approximators known as expert networks or modules (usually neural networks) that are combined by a classifier known as a gating network. These networks are trained simultaneously so as to split the input space into regions where particular experts can specialize. The gating network uses a soft split of the input data, thereby allowing data to be processed by multiple experts; the contribution of each is modulated by the gating module's estimate of the probability that each expert is the appropriate one to use. Each expert is assumed to be responsible for a Gaussian region of input space which leads to the gating unit using a multinomial logit model to partition the input space. As the input varies between two experts' regions the relative contributions of their outputs to the final output will change in a sigmoidal fashion. This model has been proposed as a model both of high-level vision[c] and of the role of the basal ganglia during sensorimotor learning[d]. The mixture of experts approach has been extended to a recursively-defined hierarchical mixture of experts (HME) architecture in which a tree of gating networks combines the expert networks into successively larger groupings that are defined over nested regions of the input space[e]. A maximum likelihood learning algorithm for the HME architecture has been derived[b] based on the Expectation-Maximization (EM) principle from statistics[f].

A modular decomposition model of visuomotor learning is shown in which two different maps can be learned for the same visual target location[g]. This represents the simplest instantiation of the hierarchical mixture of experts, having only one level and two experts. The model maps target and starting locations to motor outputs, m, which could represent, for example, the final hand location or movement vector. Each expert learns a different mapping between target locations and motor outputs appropriate for one of the starting locations, S2 or S6. The contribution of each expert's output, m_{S2} and m_{S6}, to the final motor output, m, is determined by the gating module's output, p. The output p reflects the probability that expert S6 is the correct module to use for a particular starting location. At p values of 1 or 0 the final output is determined solely by the output of expert S6 or expert S2 respectively, whereas at intermediate values of p both experts contribute to the final output. Figure modified from Ref. g.

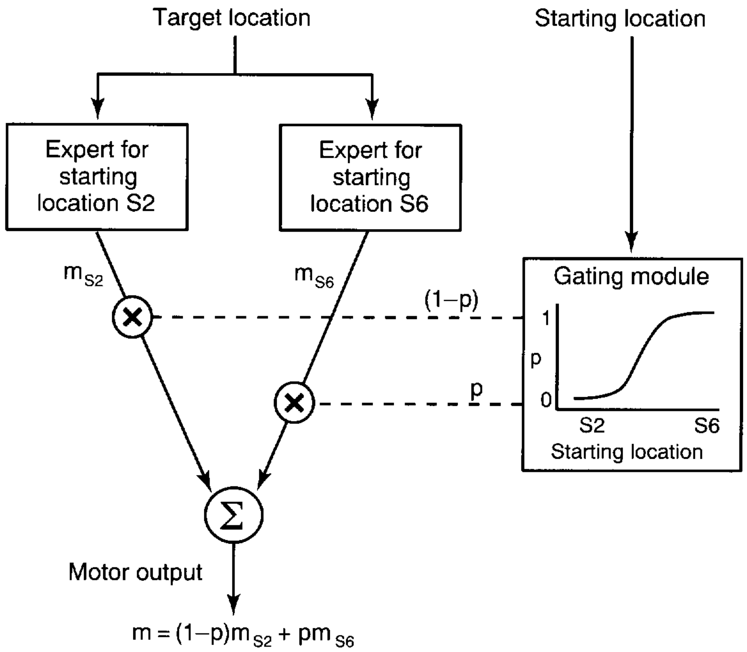

References

a Jacobs, R.A. *et al.* (1991) Adaptive mixture of local experts *Neural Comput.* 3, 79–87
b Jordan, M.I. and Jacobs, R.A. (1994) Hierarchical mixtures of experts and the EM algorithm *Neural Comput.* 6, 181–214
c Jacobs, R.A., Jordan, M.I. and Barto, A.G. (1991) Task decomposition through competition in a modular connectionist architecture: The what and where vision tasks *Cognit. Sci.* 15(2), 219–250
d Graybiel, A.M. *et al.* (1994) The basal ganglia and adaptive motor control *Science* 265, 1826–1831
e Jordan, M.I. and Jacobs, R.A. (1992) Hierarchies of adaptive experts, in *Advances in Neural Information Processing* (Moody, J., Hanson, S. and Lippmann, R., eds), pp. 985–993, Morgan Kaufmann
f Dempster, A.P., Laird, N.M. and Rubin, D.B. (1977) Maximum likelihood from incomplete data via the EM algorithm *J. Roy. Stat. Soc. B* 39, 1–38
g Ghahramani, Z. and Wolpert, D.M. (1997) Modular decomposition in visuomotor learning *Nature* 386, 392–395

a transition to a region of gradual decline. The variance of the estimate also showed an initial increase during the first second of movement after which it plateaued (Fig. 3B).

A model of this sensorimotor integration process which incorporates an internal forward model has been developed to account for these localization errors (see Box 1 for details). Using this internal forward model of the arm, the process integrates two sources of information, efferent outflow and reafferent sensory inflow, in an attempt to provide optimal estimates. This model, unlike simpler models that do not integrate both sensory and motor information, accounts for the empirical data (Fig. 3C,D)[31]. This suggests that a forward model is used to maintain an estimate of the state of the motor system.

Motor learning

Internal models, both forward and inverse, capture information about the properties of the arm. However, these properties are not static but change throughout life both on a short time-scale, due to interactions with objects in the environment, and on a longer time scale, due to growth and injury. Internal models must therefore be able to adapt to changes in the properties of the arm. Several computational approaches to learning an inverse model have been proposed (for a review see Ref. 32).

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 by a robot attached to their hand it has been shown that over time they adapt and are able to move naturally in the presence of the field. This can be interpreted as adaptation of the inverse model or the incorporation of an auxiliary control system to counteract the forces during movement. Several theoretical questions have been addressed using this paradigm. The first explored the representation of the controller and in particular whether it was best represented in joint or

Cartesian space[15]. This was achieved by adapting subjects to the field in one part of the workspace and investigating the generalization of this learning in another part of the workspace. By assessing in which coordinate system the transfer occurred evidence was provided for joint-based control. Another important advance was made in a study designed to answer whether the order in which the states (position and velocities) were visited was important for learning or whether having learned a force field for each state was enough to make learning robust to visiting the states in a novel order[33]. These findings showed that the order was unimportant and argue strongly against rote learning of individual trajectories. Motor learning of such force fields undergoes a period of consolidation after exposure to the field[34]. Indeed, the subjects' ability to perform in a previously experienced field was disrupted if a different field was presented immediately after the initial experience. Consolidation of this motor learning appears to be a gradual process because experience of a second field four hours after the first had no effect on subsequent performance in the first field. This suggests that motor learning undergoes a period of consolidation during which the motor learning or memory is fragile and may be disrupted by different motor learning.

While these approaches have focused on how a single internal model could be used in motor control, recent models have begun to investigate the computational advantages of using a set of internal models. This strategy is based on the principle of 'divide-and-conquer', in which a complex task is decomposed into subtasks, each learnt by a separate module. This strategy has recently been formalized into a computational model of learning known as the 'mixture-of-experts'[35–37]. This model consists of a set of expert modules whose outputs are combined by a single gating module. The system simultaneously learns to partition the task into subtasks (the role of the gating module) and to learn each of these subtasks (the role of the experts). The gating module smoothly combines the output of each of the experts based on the estimated probability that each expert will produce the desired output.

The mixture of experts model has been proposed to account for experimental data on visuomotor learning[38] (see Box 2). The ability of subjects to point accurately at a target (T) from seven different starting locations (S1–S7) was assessed in the absence of visual feedback in two conditions, before and after a novel remapping (Fig. 4). During the exposure phase a virtual reality system was used so that the single visual target location (T) was remapped to two different hand positions depending on the starting location (S2 or S6) of the movement. Subjects repeatedly traced out a visual triangle S2-S6-T-S6-S2-T-S2, thereby alternately pointing to the target from S2 and S6. In Fig. 4A, the dotted lines represent the path taken by the visual feedback of the finger location and the solid line represents the actual path taken by the finger. The single visual location (T) is, therefore, remapped into two distinct finger locations depending on whether the movement starts from S2 or S6. Such a perturbation creates a conflict in the visuomotor map which captures the (normally one-to-one) relation between visually perceived and actual hand locations. One

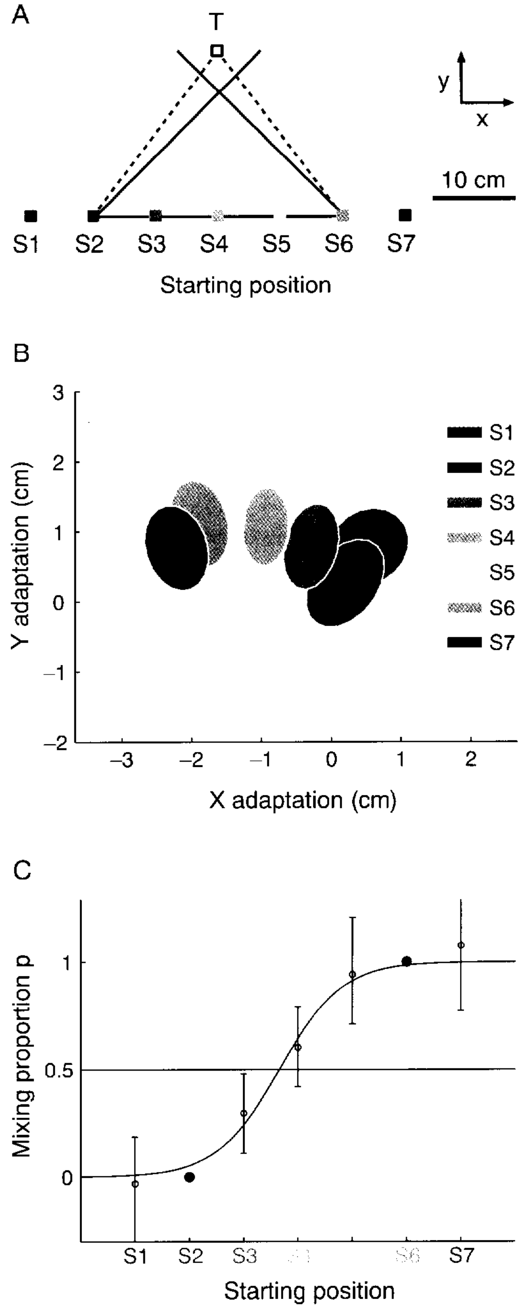

Fig. 4 Modular decomposition and visuomotor learning **(A)** The remapping of the visuomotor map dependent on starting location. Dotted lines represent the path taken by the visual feedback, solid lines represent the actual path of the finger. **(B)** Mean change as a 95% confidence ellipse (eight subjects) in pointing behaviour from each starting location (S1–S7) induced by the visuomotor perturbation. **(C)** Estimated mixing proportion (see Box 2 for details), p, and 95% confidence intervals for the seven starting locations. Modified form Ref. 38.

way to resolve this conflict is to develop two separate visuomotor maps (the expert modules), each appropriate for one of the two starting locations (see Box 2 for details). A separate mechanism (the gating module) then combines the outputs of the two visuomotor maps, based on the starting location of the movement. As in previous studies of the visuomotor system[39], the internal structure of the system can be probed by investigating the generalization properties in response to novel inputs, which in this case are the starting

Outstanding questions

- Is there a unifying principle of planning which can account for the pattern of movement at all levels in the motor hierarchy?
- What is the advantage of making movements which are optimal for a particular cost function such as jerk?
- Which of the possible computational uses of forward models is actually used within the central nervous system?
- Can models of motor learning be developed which, like the human motor system, are robust enough to learn multiple tasks, show powerful generalization to new tasks and an ability to switch between tasks appropriately?

locations on which the system has not been trained. As predicted by the mixture of experts model, subjects were able to learn both conflicting mappings (Fig. 4B), and to smoothly interpolate, in a sigmoidal fashion, from one visuomotor map to the other as the starting location was varied (Fig. 4C). These results therefore suggest that modular decomposition is a feature of visuomotor learning.

Conclusion

The aim of this review was to highlight several important themes in motor control. We have reviewed examples of several models and outlined a computational approach which can form a framework in which experimental studies can be designed and interpreted. The challenge ahead is to develop these computational approaches so that they can be applied, not only to psychophysical experiments, but also studies in neurophysiology, neuropsychology and neuroimaging.

Acknowledgements

I would like to thank Zoubin Ghahramani, Susan Goodbody and Chris Miall for their helpful comments on the manuscript. This work was supported by grants from the Wellcome Trust, the Medical Research Council and the Physiological Society.

References

1 Bernstein, N. (1967) *The Coordination and Regulation of Movements*, Pergamon
2 Bryson, A.E. and Ho, Y.C. (1975) *Applied Optimal Control*, Wiley
3 Nelson, W.L. (1983) Physical principles for economies of skilled movements *Biol. Cybern.* 46, 135–147
4 Hogan, N. (1984) An organizing principle for a class of voluntary movements *J. Neurosci.* 4, 2745–2754
5 Flash, T. and Hogan, N. (1985) The co-ordination of arm movements: An experimentally confirmed mathematical model *J. Neurosci.* 5, 1688–1703
6 Hoff, B. and Arbib, M.I. (1993) Models of trajectory formation and temporal interaction of reach and grasp *J. Motor Behav.* 3, 175–192
7 Kelso, J.A.S., Southard, D.L. and Goodman, D. (1979) On the nature of human interlimb coordination *Science* 203, 1029–1031
8 Morasso, P. (1981) Spatial control of arm movements *Exp. Brain Res.* 42, 223–227
9 Atkeson, C.G. and Hollerbach, J.M. (1985) Kinematic features of unrestrained vertical arm movements *J. Neurosci.* 5, 2318–2330
10 Kawato, M. (1992) Optimization and learning in neural networks for formation and control of coordinated movement, in *Attention and Performance, XIV: Synergies in Experimental Psychology, Artificial Intelligence and Cognitive Neuroscience – A Silver Jubilee* (Meyer, D. and Kornblum, S., eds), pp. 821–849, MIT Press
11 Uno, Y., Kawato, M. and Suzuki, R. (1989) Formation and control of optimal trajectories in human multijoint arm movements: Minimum torque-change model *Biol. Cybern.* 61, 89–101
12 Wolpert, D.M., Ghahramani, Z. and Jordan, M.I. (1995) Are arm trajectories planned in kinematic or dynamic coordinates? An adaptation study *Exp. Brain Res.* 103, 460–470
13 Wolpert, D.M., Ghahramani, Z. and Jordan, M.I. (1994) Perceptual distortion contributes to the curvature of human reaching movements *Exp. Brain Res.* 98, 153–156
14 Flanagan, J.R. and Rao, A. (1995) Trajectory adaptation to a nonlinear visuomotor transformation: Evidence of motion planning in visually perceived space *J. Neurophysiol.* 5, 2174–2178
15 Shadmehr, R. and Mussa-Ivaldi, F. (1994) Adaptive representation of dynamics during learning of a motor task *J. Neurosci.* 14, 3208–3224
16 Lackner, J.R. and DiZio, P. (1994) Rapid adaptation to Coriolis force perturbations of arm trajectory *J. Neurophysiol.* 72, 299–313
17 Sabes, P. and Jordan, M.I. Obstacle avoidance and a perturbation sensitivity model for motor planning *J. Neurosci.* (in press)
18 Desmurget, M. *et al.* (1997) Constrained and unconstrained movements involve different control strategies *J. Neurophysiol.* 77, 1644–1650
19 Viviani, P. and Flash, T. (1995) Minimum-jerk model, two-thirds power law and isochrony: converging approaches to the study of movement planning *J. Exp. Psychol. Hum. Percep. Perform.* 21, 32–53
20 Lacquaniti, F., Terzuolo, C.A. and Viviani, P. (1983) The law relating kinematic and figural aspects of drawing movements *Acta Psychol.* 54, 115–130
21 Viviani, P. and Schneider, R. (1991) A developmental study of the relationship between geometry and kinematics in drawing movements *J. Exp. Psychol. Hum. Percep. Perform.* 17, 198–218
22 Gribble, P.L. and Ostry, D.J. (1996) Origins of the power-law relation between movement velocity and curvature – modeling the effects of muscle mechanics and limb dynamics *J. Neurophysiol.* 76, 2853–2860
23 Miall, R.C. and Wolpert, D.M. (1996) Forward models for physiological motor control *Neural Networks* 9, 1265–1279
24 Jeannerod, M. (1997) *The Cognitive Neuroscience of Action*, Blackwell
25 Smith, M.C. (1967) Theories of the psychological refractory period *Psychol. Bull.* 67, 202–213
26 Wolpert, D.M. *et al.* (1992) Evidence for an error deadzone in compensatory tracking *J. Motor Behav.* 24, 299–308
27 Miall, R.C. *et al.* (1993) Is the cerebellum a Smith predictor? *J. Motor Behav.* 25, 203–216
28 Jordan, M.I. and Rumelhart, D.E. (1992) Forward models: Supervised learning with a distal teacher *Cognit. Sci.* 16, 307–354
29 Goodwin, G.C. and Sin, K.S. (1984) *Adaptive Filtering Prediction and Control*, Prentice-Hall
30 Welch, R.B. and Warren, D.H. (1986) Intersensory interactions, in *Handbook of Perception and Human Performance. Volume I: Sensory Processes and Perception* (Boff, K.R., Kaufman, L. and Thomas, J.P., eds), pp. 25;1–25;36, Wiley
31 Wolpert, D.M., Ghahramani, Z. and Jordan, M.I. (1995). An Internal model for sensorimotor integration *Science* 269, 1880–1882
32 Jordan, M.I. (1995) Computational aspects of motor control and motor learning, in *Handbook of Perception and Action: Motor Skills* (Heuer, H. and Keele, S., eds), Academic Press
33 Conditt, M.A., Gandolfo, F. and Mussa-Ivaldi, F.A. The motor system does not learn dynamics of the arm by rote memorization of past experience *J. Neurophysiol.* (in press)
34 Brashers-Krug, T., Shadmehr, R. and Bizzi, E. (1996) Consolidation in human motor memory *Nature* 382, 252–255
35 Jacobs, R.A. *et al.* (1991) Adaptive mixture of local experts *Neural Comput.* 3, 79–87
36 Jordan, M.I. and Jacobs, R.A. (1994) Hierarchical mixtures of experts and the EM algorithm *Neural Comput.* 6, 181–214
37 Cacciatore, T.W. and Nowlan, S.J. (1994) Mixtures of controllers for jump linear and non-linear plants, in *Advances in Neural Information Processing Systems 6* (Cowan, J.D., Tesauro, G. and Alspector, J., eds), pp. 719–726, Morgan Kaufman
38 Ghahramani, Z. and Wolpert, D.M. (1997) Modular decomposition in visuomotor learning *Nature* 386, 392–395
39 Ghahramani, Z., Wolpert, D.M. and Jordan, M.I. (1996) Generalization to local remappings of the visuomotor coordinate transformation *J. Neurosci.* 16, 7085–7096