

THE CEREBELLUM AS A PREDICTIVE MODEL OF THE MOTOR SYSTEM: A SMITH PREDICTOR HYPOTHESIS

R.C. MIALL and D.M. WOLPERT*

University Laboratory of Physiology
Oxford OX1 3PT, U.K.

*Department of Brain and Cognitive Sciences
M.I.T., Cambridge, MA 02139, U.S.A.

SUMMARY

The performance of motor systems with large feedback delays can be significantly enhanced by the use of internal predictive representations of the motor apparatus. The cerebellum is a likely site for these internal models, and we show that ataxic patients appear to have reduced awareness of the hand position during movement, suggesting that a sensory predictor within the cerebellum is impaired. We recently suggested that the cerebellum holds two types of neural model which together form a 'Smith Predictor'. One is a model of the motor apparatus (limbs and muscles) which provides a rapid prediction of the sensory consequences of each movement. The other model is of the time delays in the feedback control loop (conductance delays, muscle latencies, sensory processing). This delays a copy of the rapid prediction, so that it can be compared with actual sensory feedback; any errors are then used both to correct the movement and to update the internal representations of the motor apparatus. We propose mechanisms by which both parts of the Smith Predictor could be formed within the cerebellum, and present a neural network simulation based on these ideas.

INTRODUCTION

We have recently suggested that the cerebellum acts as a 'Smith Predictor' (Miall *et al.*, 1993a). This is a type of controller originally devised for engineering control systems that suffer long feedback delays -- for example catalytic crackers in steel mills (Smith, 1959). Its principles are however equally well applied to the control of human movement and we propose the scheme for the cerebellar control of visually guided movements. Here the unavoidable delays are due mainly to visual processing, with contributions from visuo-motor integrative processes, axonal conduction delays and muscular latencies. Together these delays may add up to a feedback loop time of 150-200msec, which is long with respect to the duration of many visually guided movements.

The Smith Predictor controller is based on the idea of internal models or neural representations: it holds a predictive model of the motor system, and its output is a prediction of the results of movements (Figure 1). Thus, we propose, the cerebellum receives a copy of a motor command being generated by 'upstream' motor regions (posterior parietal cortex or primary motor cortex), and uses its knowledge of the 'downstream' motor system (joints, muscles etc.) to estimate what the outcome of the movement would be, given the current state of the body. Since this output is a sensory prediction, it is not used to directly control the movement. However, this sensory prediction avoids all the delays within the feedback

system. It can be rapidly compared with the intended outcome of the movement, and the difference between the two forms an error signal used to correct the final stages of the movement (Comparator 2 in Figure 1). So by using a predictive model within an internal negative feedback loop, rapid control can be achieved even in the face of long sensory and motor delays.




Figure 1. A schematic diagram of the Smith Predictor: The Smith Predictor (shaded boxes) lies within a negative feedback loop in which errors (sensed by Comparator 1) are converted by a PD controller (box labelled C) into torques sent to the limb (L). The feedback loop contains delays on the forward and backward paths (D_1 and D_2). The dynamic model of the limb (L') lies on a fast internal feedback loop receiving a moving average of the torques sent to the muscles (MA), and sending its output to Comparator 2. A copy of its output is also passed through a model delay ($D'=D_1+D_2$) before comparison with the delayed feedback (Comparator 3). The dashed lines indicate training signals used to modify the two models.

Two additional properties of this control scheme should be mentioned. First, the prediction is used instead of the actual feedback within the negative feedback loop, so that it cancels the outstanding sensory errors (which remain present in the feedback system until the true sensory feedback returns some 150msec later) and further, unneeded, movements are not generated. It is these inappropriate responses to out-of-date signals that cause instability in feedback systems. Second, the internal model can be used to plan movements independent of their execution -- the brain can ask "what if" questions and rehearse, modify and learn movements on the basis of the internal predictions of their sensory consequences.

Evidence for these ideas is provided by cerebellar patients and monkeys with cerebellar inactivating lesions, whose movements are poorly co-ordinated and ataxic, as might be expected if their internal predictions of their movements were impaired. We have recently examined a group of ataxic patients, shown by MR imaging to have significant damage to the superior cerebellar peduncles, the output tract of the lateral cerebellum (Haggard *et al.*, 1994). These patients performed a visually guided tracking task but showed considerable intermittency in their responses (Figure 2A), indicating that they were especially reliant on visual inputs to perform the task (see Miall *et al.*, 1993b). If the visual target that they were following was briefly blanked from the screen, they could continue to tracking with only slightly increased errors. This implies that they had no difficulty in predicting the target's movement. However, if the joystick-controlled cursor was blanked off briefly, their movements became significantly less accurate (Figure 2A); this was not the case for control subjects (Miall *et al.*, 1993c). This therefore suggests that the cerebellar damage had interrupted an internal prediction of their arm movement, and they were only able to track successfully when provided with visual feedback of their arm position. We have also demonstrated the same effect in a monkey whose dentate nucleus was temporarily inactivated by an infusion of the local anaesthetic lignocaine (Figure 2B,C).

The Smith Predictor includes a second important component that differentiates it from other schemes based on internal predictive models. This second part is an internal model of the very feedback delays that we need to avoid. This time delay model receives the output of the sensory predictor and delays it by an amount equal to the delay suffered by the actual sensory feedback signals (Figure 1). Thus its output is a copy of the prediction delayed to be in synchrony with real feedback. Differences between these two signals (Comparator 3 in Figure 1) indicate a failure of the predictions, and these differences can be treated as sensory errors that must be corrected. Moreover, they can also be used to improve future predictions of the model -- they form an appropriate training signal for the brain (and we suggest specifically the cerebellum) to learn accurate predictive models.

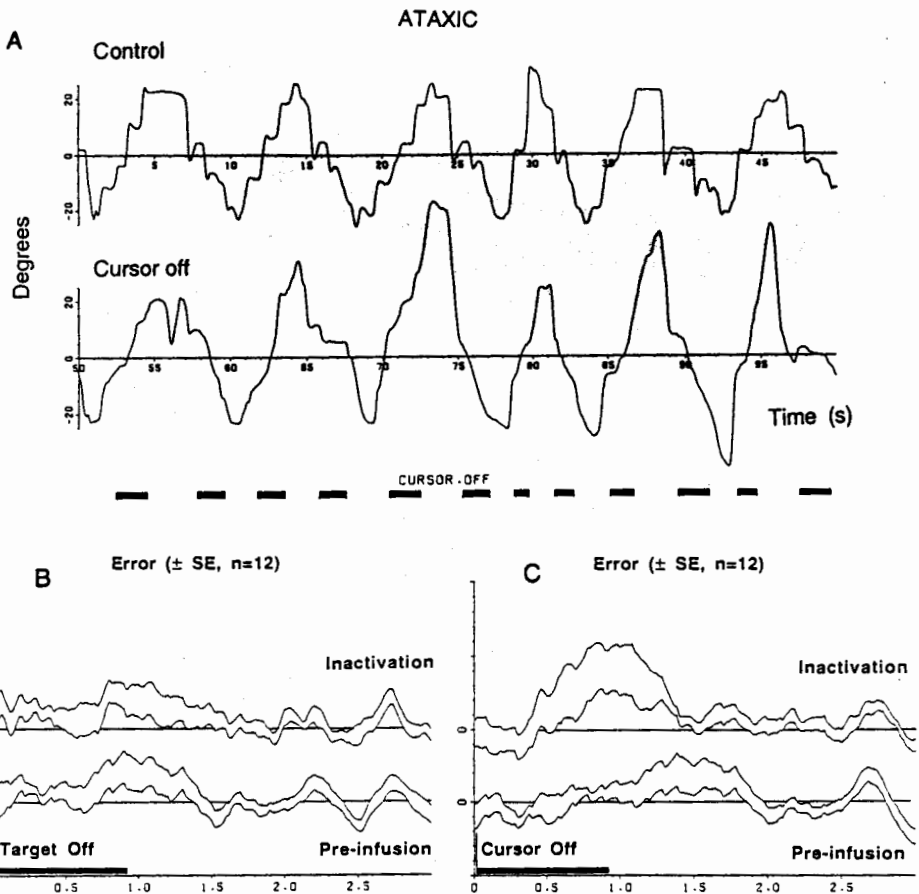


Figure 2. Visually guided tracking during cerebellar dysfunction: **A:** tracking of a ramp waveform by an ataxic patient. The upper record shows his normal tracking movements, which are much less smooth than those of normal subjects. The second record shows his tracking during brief blanking off of the cursor (black bars). Tracking without visual feedback of the cursor was smooth but inaccurate. **B** and **C:** ramp tracking by a monkey during reversible inactivation of the dentate nucleus. **B:** average tracking errors (\pm 1 S.E. of the mean) as the target is blanked off; cerebellar inactivation has little effect on performance. **C:** average tracking errors as the cursor is blanked off; cerebellar inactivation leads to great increase in errors.

The remainder of this chapter will present a neural network model of the Smith Predictor (Wolpert, 1992), and show that both its components -- a predictive model and a time delay model -- could be generated by the cerebellum.

We will consider the control of a planar two jointed arm, comprising two links hinged at a shoulder and elbow joint, and moving in the horizontal plane (Figure 3). Negative feedback control was used as the basic control system driving the arm, and the angular position error and angular velocity error for each joint were scaled by position and velocity gains to produce joint torques (box C in Figure 1). The way in which the desired trajectory of the arm is chosen for a particular movement will not be considered here; it will be assumed that the desired state for each point in time is available to the controller in appropriate co-ordinates. Likewise the problems of visual analysis of the target and of co-ordinate transformations are not examined in this model. Thus the model operates in a kinematic coordinate framework, and may be considered to be downstream of the motor cortex, within the intermediate cerebellum. We make the assumptions that the controller has access to the target's velocity and position as represented in joint angle co-ordinates, and that the arm starts off stationary but on-target. As there is no redundancy in the specifications of the tracking task, and with only two joints in

the arm, the arm configuration required to match hand position and velocity to the target is unique. Figure 3a indicates the need for predictive control -- in this figure the arm was controlled with proportional-derivative (PD) feedback but without a Smith Predictor. Even though the gains were set to get good control with immediate feedback, the arm was unstable when a feedback delay of only 30msec was introduced. In the following simulations, we sought control of the arm with high gains and with 100msec delays.

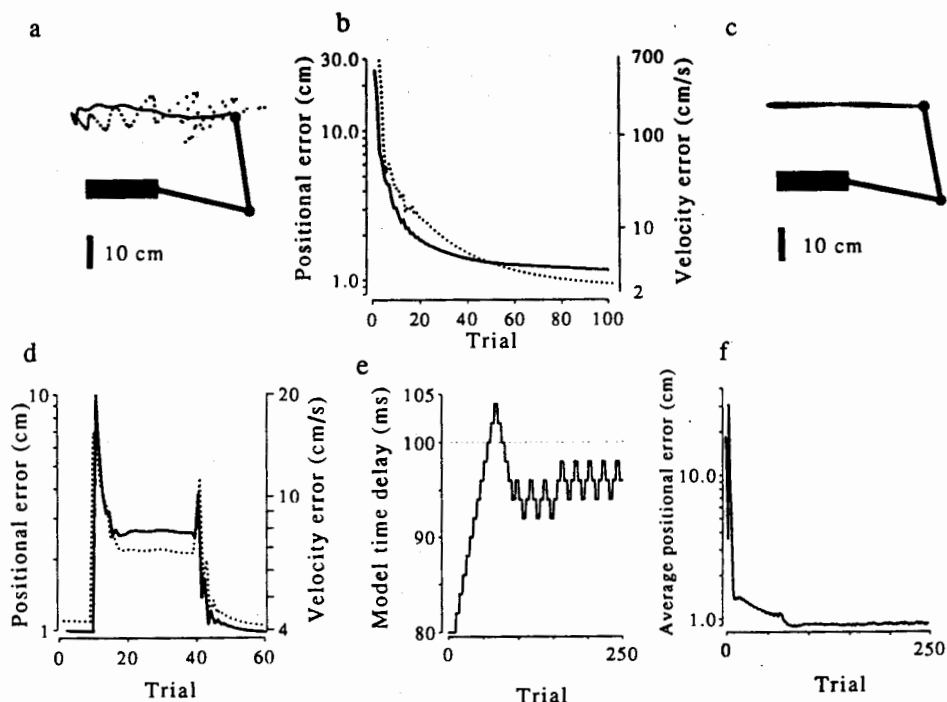


Figure 3. A neural network model of the Smith Predictor: tracking of a sinusoidal target using PD feedback. A target was moved with a 0.5 Hz cosine motion across the arm's workspace. The arm was modelled using rigid body dynamics of two freely pivoting links. The mass was assumed to be uniformly distributed along the links, which were modelled as thin cylinders. The dynamic equations of the arm (Jordan and Rumelhart, 1992) were simulated using a fourth order Runge-Kutta algorithm with a sampling frequency of 500 Hz. **a:** With the feedback delay set to only 30msec, the arm was unstable with open loop gains of 2.5 at the shoulder and elbow. **b:** Learning curves for the simulation when using a predictive neural network learning the arm's dynamics -- these were with a 100msec feedback delay and with open loop gains of 10 (shoulder) and 5 (elbow). The curves represent the position and velocity errors averaged over 10 networks. **c:** Smooth and stable tracking on the 100th trial despite the long delay and high gains. **d:** Learning curves as the model adapts to the addition and then the removal of a 7.5kg weight. **e:** Changes in the model delay time and in the tracking performance (**f**) during simultaneous learning of both the dynamic and time delay models.

Learning the Arm's Dynamics

The Smith Predictor must predict the current joint angles and velocities from the delayed feedback and from a copy of the motor command. Our neural network model therefore has six inputs: feedback of joint angles and velocities for each joint delayed by 100msec, and the outgoing joint torques for each joint. It had 4 outputs which were required to predict the two joint positions and velocities. A standard three layered feedforward neural network was used with the back-propagation algorithm to modify the synaptic weights after each iteration (Rumelhart and McClelland, 1986). Although this is not a physiological training rule, it serves to demonstrate that a solution can be found, and we can later concentrate on how more physiologically accurate networks and learning rules could achieve the same result.

The network's 6 input and 4 output units were linear; it had 50 logistic hidden units with asymptotes of ± 1.0 . Initially the weights were chosen randomly distributed between ± 0.5 .

The outputs of the network were then delayed by the model's delay, and compared to the actual state signals. The mismatch between the four delayed outputs (the estimated positions and velocities of each joint) and the corresponding delayed feedback signals were used as the error terms for the backpropagation algorithm (Comparator 3 in Figure 1). Hence, the output of the dynamic model will only be correct if it anticipates actual feedback by an amount equal to the time delay (Barto, 1990). For this simulation we simply delayed the signals by the known feedback delay without worrying about the delay mechanism (see section on "Fitting the model"). However, we will show how the value of this delay can be learned at the same time as the dynamic model.

To implement backpropagation the network's weights and activations are needed at the same time as the error term; since the error values were known only after a delay, the activities were stored and made available to the backpropagation algorithm. While storing the activation states of the neurones is not physiological impossible (Houk *et al.*, 1990), the exact mechanism that might be used is unknown, and we treat these simulations as a first attempt. The weights did not need to be stored, because the learning rate was very low and the weights changed little over the delay period. Finally, for perfect prediction of the current state of the arm, using delayed feedback, it would be necessary to retain in temporal order the previous 100msec of torques sent to the arm. However, the torques generated during human movement are unlikely to change dramatically over this time-scale because of the filtering properties of motor neurones and muscles, so we simply used a running average of the torques, using a 100msec moving average (MA in Figure 1).

The position and velocity gains of the PD controller were set to 10 for the shoulder and 5 for the elbow. The outer feedback loop had unity gain -- this was just within the stability margin with a 100msec feedback delay, so using gains of 10 and 5 put the controller well outside the stability region. To avoid gross instability while the network was naive (and therefore inaccurate) all the gains were initially set to a value of 0.5, and their values doubled at each trial up to their final values. This 'gain schedule' therefore lasts a maximum of 5 trials. Each trial consisted of two tracking periods of 2 seconds duration. During the first period backpropagation was used at each iteration, whereas during the second period the learning rate was set to zero, effectively turning off backpropagation and allowing us to assess the networks performance in the absence on intra-trial learning.

Figure 3b shows the learning curve for the network: this is the average from 10 runs with different random initialization of the network's weights. As usual with backpropagation learning there was an initial rapid decline in errors over the first few trials followed by a more gradual decline. Figure 3c shows the performance of the network on the 100th trial. The controller was producing near optimal trajectories with average positional errors of under 1 cm. All ten networks produced similar performance after 100 trials. This is the result of accurate predictions -- the network was operating with gains 4 times greater than that in Figure 3a, and with feedback delays of 100msec rather than 30msec. Figure 3d shows the model's ability to adapt to changes in the arm dynamics. After initial training, a mass of 7.5kg was added to the forearm -- the weight was added on trial 10 and removed on trial 40. The learning curves show that the weight significantly increased tracking errors. However, the network soon adapted, and errors reduced. The performance could not return to the original value because the PD feedback gains were unchanged. Thus the controller was driving a heavier mass, and could not be expected to do as well. However, when the weight was removed, the performance again briefly deteriorated as the network relearned the original unweighted dynamics. This indicates that it was indeed adapting to the arm's dynamic behaviour, rather than undergoing some non-specific changes in the face of the added weight.

Learning the Time Delay

Like the arm's dynamics, physiological feedback time delays are not fixed throughout life; they change due to growth and axonal enlargement, and are modality and stimulus dependant. It is therefore reasonable to postulate that the model delay within the Smith Predictor must be adaptable and there are really only two pieces of information which could be used. First, it would be possible to estimate the time between sending out a motor command and perceiving the response via the delayed feedback pathways. The second possibility is to continually change the model delay and try to optimise some performance criteria. It turns out that for these sorts of tracking tasks, the performance is a smooth unimodal function of the model delay relative to the actual delay. So it is possible to perform gradient descent on this error curve to find its minimum, and hence the best value of the model delay.

The model's delay was therefore changed by 2msec every 5 trials. The direction of the change was dependant on the improvement in mean performance over the last 5 trials. If the last change led to an increase in performance, the next change in the model delay was in the same direction; if performance had deteriorated, the next change was in the opposite direction. In this way a crude gradient descent scheme was implemented.

Each change in the model delay had two effects, as can be appreciated from Figure 1. The first is that the dynamic model will learn to predict ahead by an amount equal to the model delay, whether or not this is the same as the actual delay. The second is that the comparator will no longer compare the predictions of sensory feedback in correct synchrony with the actual feedback, and inappropriate adjustments to the dynamic model will therefore be made on the strength of this false error signal.

So it is a potentially difficult problem to learn both the dynamic and time delay models at the same time -- the two models interact and mutually interfere. Our solution, based on evidence from motor psychophysics (Miall *et al.*, 1990), was to have two very different learning rates for these two systems so that one system can fully adapt while the other changes little. To test this, synchronous learning of both the dynamics and the delay was simulated. The feedback delay was 100msec, but the model delay was set initially to 80msec, and the dynamic model neural network was started as usual with randomised weights. Figure 3e shows the time course of the changes in the model delay and Figure 3f shows the concomitant positional errors during tracking. The curves show how the tracking performance rapidly improved as the dynamic network and the model delay became more accurate. The model delay then overshoot, leading to a slight decrement in performance, before settling around a mean value of 96msec. Because of our crude gradient descent algorithm, it could not stay at any one delay value indefinitely, but alternated up and down one delay increment.

Fitting the Model into Cerebellar Physiology

Our hypothesis requires that there should be separate dynamic and time delay neural models. The models should receive as inputs an efferent copy of the motor command being sent to the limb, and also proprioceptive information about the current state of the body. The latter is needed for an accurate internal representation of the limb, as the arm's mechanical properties depend on its position and motion. Hence the internal dynamic model must be updated by proprioceptors. The models should lie on feedback loops, so that their output feeds back onto the input as indicated in Figure 1. Finally, there must be mechanisms to allow the models to be adapted to predict accurately the behaviour of the controlled object, i.e. neural learning mechanisms.

We believe that the cerebellum is an obvious candidate site for these neural models (Miall *et al.*, 1993a). Hence, sufficient information should be available to allow the cerebellum not only to generate but also modify each neural model. The dynamic responses of a limb change greatly during growth, while delays can change either because of increased axonal lengths, or following changes in the sensory responses (Deno *et al.*, 1989; Wolpert *et al.*, 1993). So to lay down useful internal representations of the motor system requires that the controller actively explores the environment to assess the outcome of its actions (Barto, 1989). The responses received back from the environment tell the controller not only about its effects on the external world but also about the controlled object, i.e. the arm. Thus, there are two processes to be learned: an early estimate of the outcome of actions and an estimate of the delay before actual feedback will be received; these must be learned on the basis of delayed signals from the periphery.

Climbing fibre input from the inferior olive may provide a training input to the cerebellar cortex (see Ito, 1984; or Strata, 1989). The climbing fibres should therefore signal the need for adaptation, in other words signal back the fact of a mismatch in predicted and actual feedback. A mismatch could mean that the prediction was inaccurate, or the effector's behaviour had changed, but in either case the models would require adaptation. Gilbert and Thach (1977) showed that the average rate of climbing fibre activity increased as monkeys adapted their movements to a novel load; Gellman, Gibson and Houk (1985) and Andersson and Armstrong (1987) have shown that the most potent stimulus for climbing fibres is an unexpected sensory event, and that similar "reafferent" sensory stimuli resulting from the animal's own actions do not excite the climbing fibres. There is now strong evidence that coincidence of climbing fibre and parallel fibre inputs to Purkinje cells results in long term depression of the parallel fibre:Purkinje cell synapse (Ito, 1989; Crepel and Jaillard, 1991). This may allow the cerebellar cortex to learn or modify a neural representation of the limb

dynamics. These data therefore suggest that the inferior olive signals differences between the prediction of sensory re-afference and the actual reafference. In other words, we suggest that the delayed prediction from the cerebellum is fed to the inferior olive for comparison with sensory feedback signals. There are indeed inhibitory pathways direct from the cerebellar nuclei to the inferior olive and also indirectly via the red nucleus (Weiss *et al.*, 1990). Thus the inferior olivary signals should modify the dynamic model, to ensure that the inner loop of the Smith Predictor accurately mimics actual performance. However, the climbing fibre inputs probably cannot signal the quantitative size of the mismatch. Instead, they may signal the event of a mismatch, and perhaps also the direction of the error (Houk, 1990). Barto (1990) reviews techniques to train neural networks with this sort of reinforcement signal.

Thus the major role that we would attribute to the inferior olive is the comparison between expected and actual sensory signals (Comparator 3 in Figure 1). This is vital to provide a teaching signal for the cerebellum.

The second model of the Smith Predictor requires an output which is delayed to match the returning visual feedback, perhaps 150 - 250msec later. The size of the feedback time delay could be estimated by measuring the delay between issuing a motor command and assessing its result. This would be most easy to do if the motor command were discrete (Miall *et al.*, 1993b), as the refferent signal would then change abruptly. The parallel fibres within the cerebellar cortex could act as a "tapped delay line" (Braitenberg, 1961), thus allowing the encoding of a time interval; Desmond and Moore (1988) propose that a chains of pontine nuclear cells do the same thing. However, we prefer the idea of using a predictive neural network trained to predict backwards in time: back-prediction of a signal is equivalent to delaying it. Thus, we suggest that both models within the Smith Predictor consist of predictive neural networks; the first model makes a forward prediction of the outcome of the movement. The second model makes a backward prediction, based on the output of the first model, and results in a delayed copy of the controller's actions.

The fact that the Smith Predictor contains two separate models requires training signals for both. We propose that the inferior olive provides one, allowing the dynamic model to be learned. The second training signal could be provided by the diffuse noradrenergic and serotonergic inputs from the locus coeruleus and raphe nucleus. These could provide a "performance measure", reporting to the cerebellum on the overall success of the behaviour (Gilbert, 1975). We imagine that this measure could be something like a running average of positional errors in a tracking task, or the retinal slip accumulated over a few minutes in a VOR task. In support of this view, Van Neerven, Pompeiano, Collewinj and Van der Steen (1990) have shown that beta-noradrenaline can interfere with VOR adaptation in the rabbit, while D'Ascanio, Manzoni and Pompeiano (1991) has shown that noradrenaline-blockers reduce the gain of vestibulo-spinal reflexes. Thus the time delay model of the Smith Predictor might be trained with reinforcement learning on the basis of non-specific performance criteria (Barto, 1990), while the more specific signals provided by climbing fibres train the dynamic model.

These two learning schemes could have very different learning rates. We have shown that humans (Miall *et al.*, 1990) and monkeys (unpublished data) are very much faster to adapt to changes in the gain or load of a tracking manipulandum than they are to a change in its feedback delay. This would suggest that the dynamic model is rapidly modified within the cerebellum, whereas the temporal delay model is much slower to adapt. Deno *et al.*, (1989) have shown that oculomotor adaptation to feedback delays does occur over several days; and long experience with delayed feedback in tracking paradigms certainly improves performance (unpublished data). Thus adaptation to time delays does take place in primates, as would be expected from an adaptive Smith Predictor, but is slow. Hence, we propose that the dynamic model would adapt rapidly, driven by inferior olivary input, whereas the time delay model would adapt more slowly, driven either by the same inferior olive signal or by the noradrenergic or serotonergic inputs. The difference in learning rates may be functionally unimportant, as in every-day experience it is the dynamic behaviour of the motor system that changes rapidly, for example when carrying heavy objects, rather than the feedback delays.

Finally, the model we have presented here operates in kinematic coordinates, and we propose this as a model of the intermediate cerebellum. There may actually be two independent Smith Predictors within the cerebellum, and we suggest the second one would operate in egocentric coordinates and be situated in the lateral cerebellum (Miall *et al.*, 1993a). The lateral cerebellum forms a link between visual association areas, especially the posterior parietal cortex, and the motor and premotor cortices. This cerebro-cerebellar pathway may well be the major route by which visual information reaches the cortical motor areas for the

guidance of the limbs (Stein and Glickstein, 1992). If this route was to contain a Smith Predictor, it would serve to transform a movement command (an instruction to reach a desired goal) specified by the posterior parietal cortex in visual, egocentric coordinates into a motor control signal and transfer it to the motor cortex. This then provides a common computational role for both lateral and intermediate cerebellum, and can explain many of the symptoms of cerebellar damage which are so obvious during visually guided movement.

ACKNOWLEDGEMENTS. We would like to thank the Wellcome Trust, the Medical Research Council, and the McDonnell Pew Centre for Cognitive Neuroscience for their support.

REFERENCES

- Andersson, G. and Armstrong, D.M. (1987). Complex spikes in Purkinje cells in the lateral vermis (b zone) of the cat cerebellum during locomotion. *Journal of Physiology (London)*, **385**, 107-134.
- Barto, A.G. (1989). From chemotaxis to cooperativity: Abstract exercises in neuronal learning strategies. In R. Durbin, R.C. Miall, and G. Mitchison (Eds.), *The computing neuron* (pp. 73-98). Wokingham: Addison-Wesley.
- Barto, A.G. (1990). Connectionist Learning for Control: An Overview. In T. Miller, R.S. Sutton, and P.J. Werbos (Eds.), *Neural networks for control* (pp. 5-58). Cambridge, Mass: MIT Press.
- Braitenberg, V. (1961). Functional interpretation of cerebellar histology. *Nature*, **190**, 539-540.
- Crepel, F. and Jaillard, D. (1991). Pairing of pre- and postsynaptic activities in cerebellar Purkinje cells induces long-term changes in synaptic efficacy in vitro. *Journal of Physiology (London)*, **432**, 123-141.
- D'Ascanio, P., Manzoni, D., and Pompeiano, O. (1991). Changes in gain of vestibulospinal reflexes after local injection of Beta-adrenergic substances in the cerebellar vermis of decerebrate cats. *Acta Otolaryngologica*, **111**, 247-250.
- Deno, D.C., Keller, E.L., and Crandall, W.F. (1989). Dynamical neural network organization of the visual pursuit system. *IEEE Transactions on Biomedical Engineering*, **36**, 85-92.
- Desmond, J.E. and Moore, J.W. (1988). Adaptive timing in neural networks: The conditioned response. *Biological Cybernetics*, **58**, 405-415.
- Gellman, R.S., Gibson, A., and Houk, J.C. (1985). Inferior olivary neurons in the awake cat: detection of contact and passive body displacement. *Journal of Neurophysiology*, **54**, 40-60.
- Gilbert, P.F.C. (1975). How the cerebellum could memorise movements. *Nature*, **254**, 688-689.
- Gilbert, P.F.C. and Thach, W.T. (1977). Purkinje cell activity during motor learning. *Brain Research*, **128**, 309-328.
- Haggard, P.N., Miall R.C., Wade, D.T., Anslow, P., Renowden, S., Fowler, S. and Stein, J.F. (1994) Damage to cerebellocortical pathways following closed head injury: an MRI and behavioural study. *Neurosurgery and Psychiatry* (in Press).
- Houk, J.C. (1990). Role of cerebellum in classical conditioning. *Society of Neuroscience Abstracts*, **16**, 205.8 (Abstract)
- Houk, J.C., Singh, S.P., Fisher, C. and Barto A.G. (1990) An adaptive sensorimotor network inspired by the anatomy and physiology of the cerebellum. In T. Miller, R.S. Sutton, and P.J. Werbos (Eds.), *Neural networks for control* (pp. 301-348). Cambridge, Mass: MIT Press.
- Ito, M. (1984). *The cerebellum and neural control*. New York: Raven Press.
- Ito, M. (1989). Long-term depression. *Annual Review of Neuroscience*, **12**, 85-102.
- Jordan, M.I. and Rumelhart, D.E. (1992). Forward models: Supervised learning with a distal teacher. *Cognitive Science*, **16**, 307-354.
- Miall, R.C., Kerr, G.K., Wolpert, D.M., and Forsyth, D. (1990). Adaptation to task dynamics and visual feedback in human visually guided movements. *Neuroscience Letters Supplement*, **38**, S51.
- Miall, R.C., Weir, D.J., Wolpert, D.M. and Stein, J.F. (1993a). Is the cerebellum a Smith Predictor? *Journal of Motor Behaviour*, **25**, 203-216.
- Miall, R.C., Weir, D.J., and Stein, J.F. (1993b). Intermittency in human manual tracking tasks. *Journal of Motor Behavior*, **25**, 53-63.
- Miall, R.C., Haggard, P.N., and Stein, J.F. (1993c) Visuo-motor pursuit movements in ataxic patients with cerebellar damage. *Journal of Physiology (London)*, **473**, 24P.
- Rumelhart, D. and McClelland, J. (1986) *Parallel distributed processing*. MIT Press, Cambridge Mass.
- Smith, O.J.M. (1959). A controller to overcome dead time. *ISA Journal*, **6**, 28-33.
- Stein, J.F. and Glickstein, M. (1992) The role of the cerebellum in the visual guidance of movement. *Physiological Reviews*, **72**, 967-1017.
- Strata, P. (1989). *The olivocerebellar system in motor control*. Berlin: Springer-Verlag.
- Van Neerven, J., Pompeiano, O., Collewyn, H., and Van der Steen, J. (1990). Injections of beta-noradrenergic substances in the flocculus of rabbits affect adaptation of the VOR gain. *Experimental Brain Research*, **79**, 249-260.

- Weiss, C., Houk, J.C., and Gibson, A.R. (1990). Inhibition of sensory responses of cat inferior olive neurons produced by stimulation of the red nucleus. *Journal of Neurophysiology*, **64**, 1170-1185.
- Wolpert, D.M., Miall, R.C., Cumming, B.C. and Boniface, S.J. (1993). Retinal adaptation of visual processing time delays. *Vision Research*, **33**, 1421-1430.
- Wolpert, D.M. (1992) *Overcoming time delays in visuo-motor control*. D.Phil. Thesis, Lincoln College, Oxford.