

2013 Special Issue

Adaptive filters and internal models: Multilevel description of cerebellar function

John Porrill^a, Paul Dean^{a,*}, Sean R. Anderson^b^a Department of Psychology, Sheffield University, Western Bank, Sheffield, S10 2TP, UK^b Department of Automatic Control and Systems Engineering, Sheffield University, Sheffield, S1 3JD, UK

ARTICLE INFO

Keywords:

Cerebellum
Internal model
Adaptive control

ABSTRACT

Cerebellar function is increasingly discussed in terms of engineering schemes for motor control and signal processing that involve internal models. To address the relation between the cerebellum and internal models, we adopt the chip metaphor that has been used to represent the combination of a homogeneous cerebellar cortical microcircuit with individual microzones having unique external connections. This metaphor indicates that identifying the function of a particular cerebellar chip requires knowledge of both the general microcircuit algorithm and the chip's individual connections.

Here we use a popular candidate algorithm as embodied in the adaptive filter, which learns to decorrelate its inputs from a reference ('teaching', 'error') signal. This algorithm is computationally powerful enough to be used in a very wide variety of engineering applications. However, the crucial issue is whether the external connectivity required by such applications can be implemented biologically.

We argue that some applications appear to be in principle biologically implausible: these include the Smith predictor and Kalman filter (for state estimation), and the feedback-error-learning scheme for adaptive inverse control. However, even for plausible schemes, such as forward models for noise cancellation and novelty-detection, and the recurrent architecture for adaptive inverse control, there is unlikely to be a simple mapping between microzone function and internal model structure.

This initial analysis suggests that cerebellar involvement in particular behaviours is therefore unlikely to have a neat classification into categories such as 'forward model'. It is more likely that cerebellar microzones learn a task-specific adaptive-filter operation which combines a number of signal-processing roles.

© 2012 Elsevier Ltd. All rights reserved.

1. Introduction: the chip metaphor

Recent reviews indicate that the possible role of the cerebellum in the formation of internal models is a topic of growing interest (Bastian, 2011; Cermnara & Apps, 2011; Ebner, Hewitt, & Popa, 2011; Imamizu, 2010; Medina, 2011; Shmuelof & Krakauer, 2011). Although the term 'internal model' can be used very generally in this context to refer to any neural representation of a dynamic system (Wolpert, Ghahramani, & Jordan, 1995), many of its most important conceptual features can be captured by a simple example (Fig. 1).

Motor commands (as expressed by motoneurons) act on muscles which move a part of the body. The mechanical properties of the muscles and body part (for convenience referred to as the motor plant) ensure that the dynamics of the movement differ from those of the command (Fig. 1(A)). The circuit shown in Fig. 1(B) allows a model of this plant to be learnt, by sending a copy of the motor commands to an adaptive element (highlighted in red

throughout). The output of this element is compared with sensory feedback from the actual movement, and the discrepancy between the two used as a signal to 'train' the adaptive element (a form of supervised learning). As training proceeds the discrepancy decreases, meaning that the dynamics of the adaptive element resemble more closely the dynamics of the motor plant. In this way a model of the motor plant is learnt.

The model shown in Fig. 1 (usually referred to as a 'forward' model, as explained below) has many uses, including the prediction of the sensory effects of movement. Such a prediction can, for example, help distinguish the sensory signals produced by one's own movements from those arising from external events—the classical reafference problem (e.g. Cullen, 2004), and further uses are discussed in Section 3. Hence proposals that the cerebellum takes part in the formation of internal models seek to provide a crucial link between cerebellar function and proven sensorimotor competences (Blakemore, Frith, & Wolpert, 1999, 2001; Imamizu, Kuroda, Miyauchi, Yoshioka, & Kawato, 2003; Kawato, 1995, 1996, 1999, 2008; Miall, Christensen, Cain, & Stanley, 2007; Miall & Reckess, 2002; Miall, Weir, Wolpert, & Stein, 1993; Miall & Wolpert, 1996; Wolpert, 1997; Wolpert & Ghahramani, 2000; Wolpert et al., 1995; Wolpert & Kawato, 1998; Wolpert, Miall, & Kawato, 1998).

* Corresponding author.

E-mail address: p.dean@sheffield.ac.uk (P. Dean).

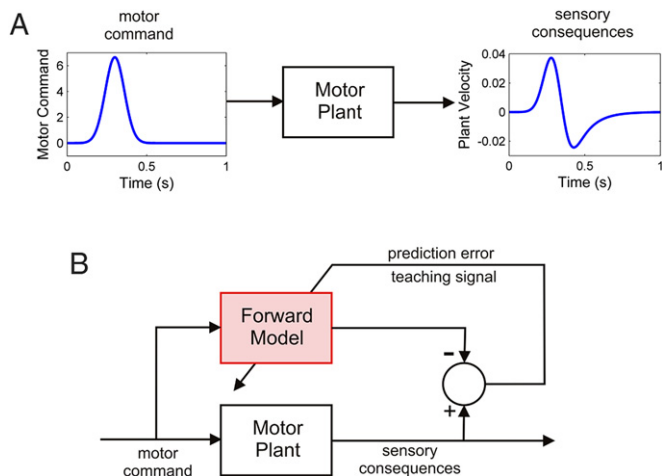


Fig. 1. Dynamic response and forward model of a simple viscoelastic motor plant. A: The motor command from the motoneurons acts on muscles, which move some part of the body. The mechanics of the muscles plus body part (= ‘motor plant’) mean that the temporal trajectory of the movement differs from that of the motor command of the profile of the plant. The example shows the velocity response of a simple viscoelastic plant to a Gaussian motor command. (For convenience, the dynamics of sensory transduction are neglected, so the sensory measurement of the movement introduces no distortions.) B: A circuit for learning a forward model of the motor plant dynamics. The forward model is the adaptive element (highlighted in red; this convention is also applied in later figures). It can be learnt using sensory error (that is, the difference between the predicted and actual sensory consequences of the motor command) as teaching signal. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

To evaluate how far this proposed link is supported by current evidence, the present article first outlines the popular ‘chip’ metaphor for cerebellar organisation, which requires cerebellar functions to be modelled at both microcircuit and external-connectivity levels.

1.1. The ‘chip’ metaphor

The arrangement of neurons and their connections within cerebellar cortex is broadly similar over the entire cortical mantle, whereas each individual region of the cerebellar cortex has a unique pattern of connections with external neural structures. This combination has long been recognised:

“The cerebellar tissues have quite a uniform histological structure. Their role in the actual motor control, however, varies from region to region, depending upon what subcortical structures they are connected with, as pointed out by Herrick (1924)” (Ito, 1970, p. 162);

and has given rise to what might be termed the ‘chip’ metaphor of cerebellar organisation

“Cerebellar microcomplexes are connected to various systems of the brain and so play diverse roles in central nervous system functions. This situation would be similar to that of a computer chip which can be used for a great many purposes.” (Ito, 1997, p. 475).

This metaphor is illustrated in Fig. 2, which shows in schematic form a functional sub-region of cerebellar cortex represented by an identical internal structure and idiosyncratic inputs and outputs. The important implication of the chip metaphor is that the function of any particular cerebellar sub-region depends on *both* the signal-processing capacities of the generic chip, *and* the particular architecture in which it is embedded.

The relevance of the chip metaphor for evaluating internal-model hypotheses of cerebellar function can be illustrated by the




Fig. 2. The cerebellar ‘chip’ metaphor. Each cerebellar microzone has a similar internal organisation, but its own idiosyncratic set of connections, two inputs and one output. The climbing fibre and output connections are unique to a microzone: some of a microzone’s mossy fibre inputs may be shared with other microzones. The climbing fibre teaching signal specifies the learning goals for the chip, hence it is this connectivity which is basic to defining individual microzones. The Purkinje cell output must then be connected to a target region in the deep cerebellar or vestibular nuclei which contributes to achieving this goal, and for which the learning procedure hardwired into the chip is stable and convergent. This provides a strong constraint on the output connectivity. The mossy fibre input connections are the least constrained. They can be regarded as a wide ‘bus’ of possibly relevant sensory and motor signals, from which those signals actually relevant to the task will be chosen by the learning procedure.

‘inverse-model’ circuit shown in Fig. 3. The need for an inverse model of the motor plant arises because of the ‘distorting’ effects of plant dynamics on the motor command, as shown in Fig. 1(A). Motor commands that specify a desired trajectory for a part of the body must therefore be converted into a form that compensates for the characteristics of the plant. This can be achieved by passing the command, not directly to the plant itself, but indirectly through an inverse model of the plant (Fig. 3(A)). As with the ‘forward’ plant model (terminology emphasising the contrast with the inverse plant model) such a model needs to be learnt, and a possible circuit for achieving this is shown in Fig. 3(B). Although, as will be argued later, the circuit shown in Fig. 3(B) is too simple to be biologically realistic, it illustrates an important point about the difference between an adaptive element and the circuit of which it is a component. Comparison of Figs. 1 and 3 shows how the same adaptive element can learn either a forward model, or an inverse model, depending on the details of the external wiring. This is exactly the point captured by the cerebellar ‘chip’ metaphor of Fig. 2.

Evaluating internal-model hypotheses of cerebellar function therefore entails evaluating both the microcircuit model and the way it is wired into any particular system-level architecture. The particular microcircuit model chosen here is the adaptive filter, and this is briefly described in Section 2, and its general computational suitability for internal model formation explained. Particular internal-model architectures are then assessed in two stages. The first asks how they are biologically plausible—how far the signals they require could be in principle provided biologically (Section 3). The second stage considers the problems that arise when the plausible architectures have to be mapped in practice onto complex neural circuits (Section 4). The final section addresses the implications of the internal-model hypothesis for the future understanding of cerebellar functions (Section 5).

2. Microcircuit level: inside the chip

The repeating nature of the cerebellar microcircuit suggests that there is a generic ‘cerebellar algorithm’, and hypotheses about its computational capability (Albus, 1971; Marr, 1969) appeared soon after the microcircuit itself was first described (Eccles, Ito, & Szentágothai, 1967). The Marr–Albus framework was further developed by Fujita (1982), who proposed that the cerebellar circuit acts as an adaptive filter in which the mossy fibre inputs to the cerebellum convey dynamic time-varying signals rather than the static spatial patterns associated with the original Marr–Albus formulations (further details in Section 2.4). Since it appears that

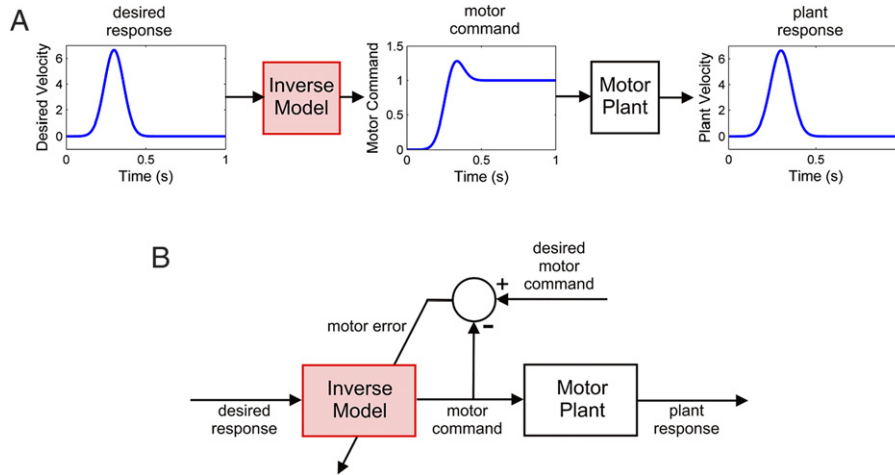


Fig. 3. Inverse model control of a simple viscoelastic motor plant. A: A desired plant response, in the form of a Gaussian signal, is transmitted through an inverse model of the viscoelastic motor plant (also used in Fig. 1). The inverse model transforms the desired response into the pulse-step motor command that is the exact control input required to reproduce the desired response from the plant. B: A circuit for learning an inverse model of the motor plant dynamics. The inverse model is the adaptive element and can be learnt by using motor error (that is, the difference between actual and desired motor command) as teaching signal as shown in the diagram. Since the desired motor command is not known (or if it was, could be used directly to drive the plant) this inverse model architecture is not biologically plausible as it stands (see text).

many of the cerebellar models concerned with behaviour (such as the control of eye or arm movements) resemble the adaptive-filter model (Dean, Porrill, Ekerot, & Jorntell, 2010), this is the microcircuit model considered here. Because the main focus here is on the system-level connectivity, only a brief description is given of the model and its biological plausibility. More detailed analysis is available elsewhere (e.g. Dean & Porrill, 2010, 2011).

2.1. The adaptive filter

A linear in weights adaptive filter (Fig. 4(A)) takes as input a number of time varying signals $u_1(t)$, $u_2(t)$, \dots , $u_N(t)$ (for clarity only a single input is shown in the figure). These inputs are passed through a basis of fixed filters G_i to produce signals $p_i = G[u_1(t), \dots, u_N(t)]$ which are combined with weights (w_1 , w_2 , \dots , w_K) to produce the output

$$z(t) = \sum w_i p_i(t). \quad (1)$$

The term adaptive filter is used because these weights are not pre-calculated, but are changed systematically during the operation of the filter to improve filter performance on a task. The learning rule for filter weights usually takes the form

$$\delta w_i = -\beta \langle e(t) p_i(t) \rangle \quad (2)$$

where β is a positive learning rate parameter, $e(t)$ is called the teaching signal (which is often related to task errors) and the angle brackets denote a time average. This learning rule has been rediscovered in many contexts and given names including the delta rule, the Widrow–Hoff rule (Widrow & Stearns, 1985), the LMS rule, the MIT rule, the covariance rule, etc. We will generally use the term covariance rule (Sejnowski, 1977) since updates to a weight w_i are proportional to the covariance of the signal $p_i(t)$ controlled by the weight and the teaching signal $e(t)$.

This learning rule is guaranteed to converge when the teaching signal is the error in filter output $e_{\text{out}}(t) = z(t) - z_d(t)$, but, as we shall see later, this output error signal is not always available to use as a teaching signal. There is some flexibility in the choice of teaching signal however. Suppose we have access to a signal $e(t) = H * e_{\text{out}}(t)$ related to output error by a linear filter H . Then the learning rule above will still converge as long as the linear filter H is strictly positive real (SPR) (Sastry & Bodson, 1989), that is the phase shift due to H at any frequency must be less than 90° so

that the sign of the correlation in the learning rule is not reversed at any frequency. This SPR requirement means, for example, that if learning is required to be stable up to a maximum frequency of 1 Hz then the learning rule can tolerate delays in the teaching signal of up to 250 ms (in general the allowable delay is related to the maximum frequency by $T = 1/4f$; for further discussion see Porrill and Dean (2007a)).

The filter shown in Fig. 4(A) is often referred to as an analysis–synthesis filter, with the analysis stage corresponding to the transformation of inputs by the bank of fixed filters, and the synthesis stage corresponding to the subsequent recombination of the suitably weighted transformed inputs.

2.2. Similarity to cerebellar microcircuit

Fig. 4(B) shows the basic cerebellar microcircuit in a form that allows comparison with the adaptive filter shown in Fig. 4(A). Mossy fibre inputs carry the signals $u_1(t)$, $u_2(t)$, \dots , $u_N(t)$ to the cerebellum; these are analysed by granule cells whose axons bifurcate to form parallel fibres which carry basis signals $p_1(t)$, $p_2(t)$, \dots , $p_K(t)$. The parallel fibres synapse on Purkinje cells which perform the linear synthesis to produce the signal $z(t)$ carried by the simple spike output. In addition a single climbing fibre winds around the Purkinje cell; spikes on this input cause the PC to fire complex spikes on a 1–1 basis. These complex spikes are interpreted as a trigger for synaptic weight changes with the climbing fibre input acting as a teaching signal. The covariance rule above is broadly consistent with the known behaviour of complex spike dependent LTD and LTP at PF/PC synapse.

An important aspect of the resemblance illustrated in Fig. 4 is that the adaptive filter model offers an explanation of the two distinctive features of the cerebellar microcircuit (e.g. Dean & Porrill, 2011). One is the enormous number of granule cells, estimated at up to 80% of all neurons in the human brain (Herculano-Houzel, 2010), potentially corresponding to the large number of basis functions required by an analysis–synthesis filter, particularly if non-linear bases are needed. The second distinctive microcircuit feature is the unusual behaviour of climbing fibres. These produce complex spikes in Purkinje cells on average at ~ 1 Hz, apparently too low a frequency to have a significant influence on Purkinje cell output (average simple-spike frequency ~ 40 Hz). On the other hand

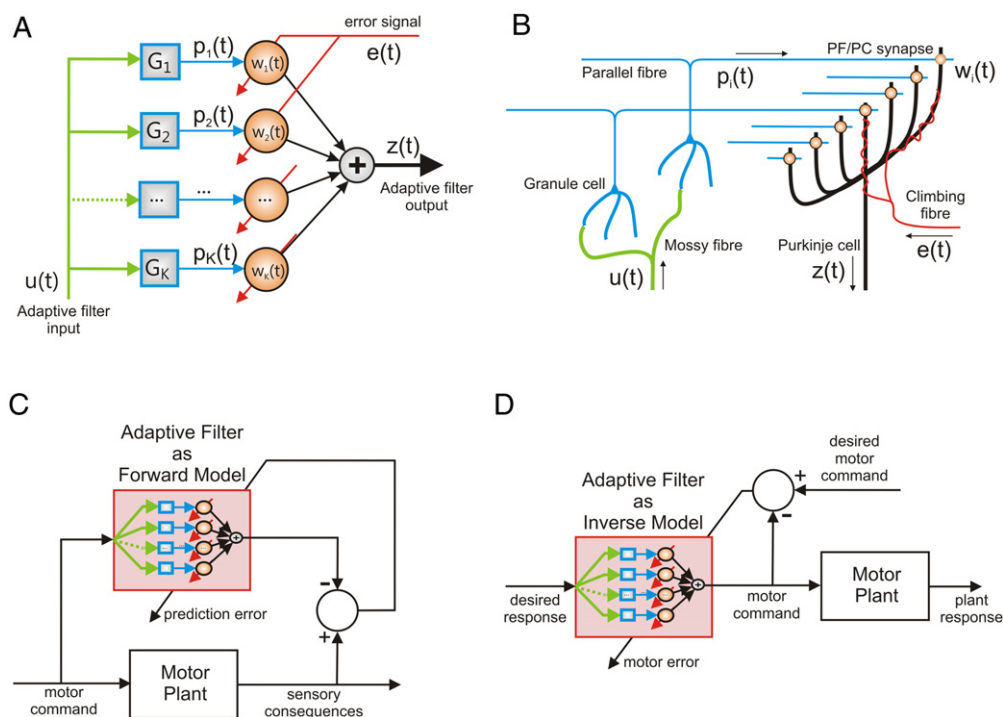


Fig. 4. Adaptive filter model of the cerebellum. A: Systems diagram of the adaptive filter. Input signals (only one is shown for clarity) are analysed by a bank of filters (in engineering applications this could be a bank of delay-lines with a range of delays) to produce a basis of signals; these are recombined with appropriate weights to produce the desired output. The filter weights are adapted over time using the least-mean-squares (LMS) rule. B: Schematic diagram of the cerebellar microcircuit. Mossy fibres transmit the input signal to cerebellar cortex. The granule cell layer transforms the input signal to produce an output transmitted by the parallel fibres. Parallel fibres synapse onto the dendrites of a Purkinje cell. Parallel fibre/Purkinje cell synaptic efficacy is altered by correlational firing of a single climbing fibre and each parallel fibre. C: Forward model approximation of a motor plant using the adaptive filter model of cerebellum. The adaptive filter output is driven by the motor command and the adaptation of filter weights is driven by prediction error. D: Inverse model approximation of a motor plant using the adaptive filter model of cerebellum. The adaptive filter output is driven by the desired response signal and the adaptation of filter weights is driven by the motor error in contrast to (C). This change in connectivity is the only difference between use of the adaptive filter in forward and inverse model roles.

the complex spike produced by a climbing fibre action potential is associated with a large widespread calcium transient throughout the Purkinje-cell dendritic tree, apparently related to plasticity at the estimated 150,000 parallel-fibre synapses on the tree (e.g. Ohtsuki, Piochon, & Hansel, 2009). The peculiar combination of low-frequency firing with very extensive synaptic input that is characteristic of the climbing fibre is exactly what is wanted from a teaching signal which must alter all synaptic weights appropriately without contaminating the filter output.

2.3. Biological plausibility of adaptive-filter model

Although the adaptive filter model of the basic cerebellar microcircuit is popular and appears plausible in certain broad respects, some detailed microcircuit features have been described that appear incompatible with it. For example, some recent studies of granular-layer processing have suggested that mossy fibre signals may be only slightly altered, rather than transformed as required by an adaptive filter. A second example concerns Purkinje cell firing. In certain circumstances Purkinje cells alternate between ‘up’ states with depolarised membrane potential and simple-spike firing, and ‘down’ states with hyperpolarised membrane potential and no simple spikes. In other circumstances, Purkinje cell firing is apparently characterised by complex patterns and pauses, rather than being straightforwardly related to task variables.

Possible explanations for a number of these putative incompatibilities have been discussed previously (Dean & Porrill, 2010; Dean et al., 2010) with the tentative conclusion that for at least some cerebellar microzones the adaptive filter remains a good candidate model (Dean, Jörntell, & Porrill, 2013; Dean & Porrill, 2011).

This conclusion appears consistent with a recent review of cerebellar plasticity (Gao, van Beugen, & De Zeeuw, 2012), which argues that the granular layer increases the diversity of mossy-fibre inputs whereas the Purkinje cell network creates output by selecting appropriately from this diversity, processes that would seem to correspond well to the analysis and synthesis stages of an adaptive filter. Therefore, since the emphasis here is on circuitry external to the cerebellum, the remainder of the review proceeds on the assumption that this conclusion is correct. The central question considered here concerns the biological plausibility of the adaptive filter as a candidate for the adapting element in forward models as shown in Fig. 4(C) and in inverse plant models as in Fig. 4(D) (compare with Figs. 1(B) and 3(B) respectively).

2.4. Computational adequacy of adaptive filter

However, it only makes sense to address questions of biological plausibility if the putative microcircuit algorithm has the required computational power. Can an adaptive filter be used in principle as the learning element in internal model architectures?

A general answer for linear adaptive filters has been provided by Widrow and Stearns (1985), who specifically analysed how they could successfully be used to learn either forward or inverse models, and how such adaptive models could be used for model control, inverse control, interference cancellation, system identification, and signal prediction.

More recent work on the CMAC (Cerebellar Model Articulation Controller) described by Albus (1971) is consistent with this analysis. The CMAC treats the cerebellar microcircuit as a look-up table which stores the desired response to a given set of inputs, and therefore functions as a pattern classifier or feature detector.

However, the learning rule used by the CMAC for adjusting its weights is of the same form as the covariance rule used by adaptive-filter models, so the main difference between the two consists of how inputs are coded. Although Fujita (1982) has claimed that the spatial pattern-classifier used in the CMAC does “not account satisfactorily for processing of time-analog signals conveyed by frequency-modulated nerve impulses” (p. 195), it can be argued that the differences in coding between the CMAC and the adaptive-filter are concerned more with convenience of representation for individual problems rather than a fundamental computational point (Dean et al., 2010).

It turns out that although CMACs are now rarely used for simulating the role of the cerebellum itself (Dean, Mayhew, & Langdon, 1994), they continue to be applied to an extremely wide variety of adaptive control and signal-processing problems (~30 papers per year), including fuzzy control, non-linear control and stock index forecasting (Cheng, 2011; Commuri & Lewis, 1997; Lin, Chen, & Yeung, 2010; Lu & Wu, 2011; Tao & Su, 2011). Of particular relevance here are the CMAC applications to the control of complex robots (Kim & Lewis, 2000; Sabourin, Bruneau, & Buche, 2006). If the above argument concerning CMACs and adaptive filters is accepted, the continuing usefulness of the CMAC in engineering and other contexts can be used as testimony to the computational adequacy of adaptive filters in many control and signal processing architectures, including those required by internal-model hypotheses.

3. Chip connectivity: biological plausibility

The second stage of evaluating internal-model hypotheses concerns the system-level connectivity required by an adaptive filter to implement the variety of internal-model architectures that have been proposed for the cerebellum. To make this evaluation process more tractable, we split it here into two stages. First, which specific internal-model architectures are *in principle* biologically plausible? Second, do these plausible architectures *in practice* map onto neural circuitry in the ways proposed by the models?

The present section deals with the first of these questions: how far are the proposed circuits biologically plausible, in the sense of requiring input signals that at least in principle could be available in the brain. We consider five specific circuits that have been proposed in the literature. (This section is a much fuller exposition of material covered briefly in Supplementary Material of Dean et al., 2010.)

3.1. Forward models

A forward model of a dynamical system describes the relationship between the inputs to the system and its outputs, for example the relationship between the motor commands to a motor plant (e.g. a mechanism consisting of muscles and joints) and the sensory consequences of the movement they produce (Fig. 1). Implementing a forward model of a complex motor plant requires a learning element that can accurately simulate complex input-output transformations, and in biological systems such models can only be acquired by supervised, trial and error, learning. Forward models are thus ideally suited to implementation via an adaptive filter.

A simple architecture in which a forward model can be learned successfully has been shown in Fig. 1. It can be seen that the connectivity requirements are:

1. Mossy fibre input to the forward model consists of an efference copy of the relevant motor commands.
2. The error signal required for learning is error in model output. This signal must be made available on the climbing fibres.

3. In order to calculate the required error signal the model output must target a comparator in which it is subtracted from the actual plant output.

If the teaching signal is output error (i.e. discrepancy between actual and predicted sensory signal) as shown in Fig. 1, then the adaptive filter with covariance learning rule is guaranteed to learn to be an optimal estimate of the plant (in the sense that it minimises the mean square prediction error). Provided sufficient time is available for learning, the accuracy of this estimate is therefore limited only by the completeness of the basis signals generated by the adaptive filter (see Fig. 4(A)).

In engineering systems the architecture shown in Fig. 1 is very often used, since the forward models learnt can be ‘unplugged’ and used elsewhere. It is difficult to see how this option could be implemented by a biological system. However forward models are often required as components of more complex architectures that can be biologically relevant. It is important to realise that successful learning in these architectures does not follow directly from the simple case described here, and that in each case a separate analysis must be made of the availability of the error signal required for learning, based on the specified connectivity. We will consider three further examples of architectures which include a forward model: noise cancellation in Section 3.1.1, the Smith predictor in Section 3.1.2, and the state estimator in Section 3.1.3.

3.1.1. Noise cancellation

An important application of forward models in signal processing is adaptive noise-cancellation (Widrow & Stearns, 1985). The basic structure of the problem is shown in Fig. 5(A). An input signal of interest, $s(t)$, is corrupted by noise $n(t)$, for example unwanted aircraft noise may leak into a headphone audio signal (Dean et al., 2013). Independent information about the noise source is assumed to be available, in the headphone example from a small microphone embedded in the earphones. This noise-source information, or reference signal $r(t)$, is used as the input to an adaptive filter, whose job it is to provide an estimate $\hat{n}(t)$ of the actual noise that is present in the signal. If $r(t)$ and $n(t)$ were identical the task would be very straightforward, because the output of the required filter would be identical to its input. However the noise may have been changed before being added to the signal of interest, a process represented in the diagram by the box labelled ‘noise channel’. The task of the adaptive filter is therefore to learn a forward model of the noise channel.

Previous analyses would suggest that the required teaching signal is filter output error $e(t) = \hat{n}(t) - n(t)$, but this is clearly unavailable since we have no access to the ‘true’ noise signal $n(t)$. An alternative teaching signal can be derived by considering a slightly unusual performance measure, the power of the predicted signal $\hat{s}(t)$, given by

$$E = \frac{1}{2} \langle \hat{s}^2 \rangle. \quad (3)$$

Expanding $\hat{s}(t)$ and assuming statistical independence between both noise and reference noise and the signal $s(t)$ gives

$$E = \frac{1}{2} \langle \hat{s}^2 \rangle = \frac{1}{2} \langle (s + (\hat{n} - n))^2 \rangle = \frac{1}{2} \langle s^2 \rangle + \frac{1}{2} \langle (\hat{n} - n)^2 \rangle \quad (4)$$

(since other cross-correlation terms vanish). Hence this cost function is equal to mean square output error plus a constant, so minimising it minimises output error. Its gradient is

$$\frac{\partial E}{\partial w_i} = \left\langle \hat{s} \frac{\partial \hat{s}}{\partial w_i} \right\rangle = \langle \hat{s} p_i \rangle \quad (5)$$

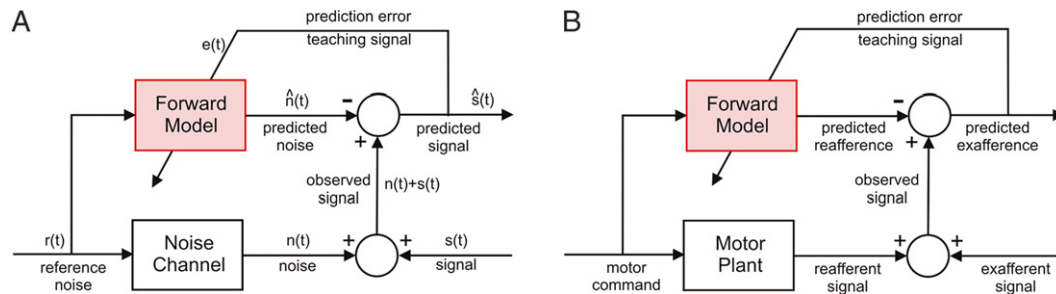


Fig. 5. Adaptive cancellation of reafferent signals. **A:** Adaptive noise cancellation architecture. The problem addressed in noise cancellation is to suppress noise $n(t)$ that additively corrupts a signal of interest $s(t)$. A known reference noise $r(t)$, which is correlated with the unknown disturbance noise, is used to drive a forward model of the noise channel. The predicted noise signal $\hat{n}(t)$ produced by the forward model is used to cancel the noise from the observed signal, resulting in a prediction $\hat{s}(t)$ of the signal of interest, this signal also acts as teaching signal. The architecture is adaptive, so that the forward model can track changing dynamics of the noise channel. **B:** Reafferent signal cancellation. During active sensing reafferent signals are often generated that can interfere with the detection and analysis of exafferent signals. The adaptive noise cancellation architecture in (A) can be used directly to overcome the reafference problem by substituting motor commands for reference noise. Hence, in a biological scenario the animal or human can learn to predict the sensory consequences of their own movements and cancel these reafferent components from observed sensory signals.

so the gradient descent learning rule is a covariance learning rule

$$\delta w_i = -\beta \langle \hat{s}(t) p_i(t) \rangle \quad (6)$$

where the teaching signal is $e(t) = \hat{s}(t)$. This learning rule adjusts the weights until the estimated signal is uncorrelated with all the components $p_i(t)$ of the reference noise $r(t)$. This is an informative example, because the teaching signal we have derived is clearly not a performance error, in fact it is the signal of interest. This means that it also provides an example where successful learning does not reduce the signal carried on the climbing fibre to zero, as would be expected for an error signal.

This architecture shown can be applied directly to the problem of predicting the sensory effects of movement in biological systems (Fig. 5(B)). The signal of interest is now the output of a biological sensor. The task is to separate sensory signals that are produced by events in the outside world (exafference) from the ‘noise’ generated in the sensor by the animal’s own movements (reafference). The nature of this contamination cannot be known directly, but there is information about the movements themselves, provided by the motor commands sent to the relevant muscles. This ‘efference copy’ information is in effect reference information about the noise source, and so could be used as input to an adaptive filter (possibly located in the cerebellum, see below) that learns to mimic the transformation of motor commands into sensory signals. Thus the cerebellum would learn a forward model of ‘Plant + Sensor Dynamics’, including its basic properties of elasticity, viscosity and inertia, together with any post-processing in the sensory apparatus.

Once learning has been achieved, the adaptive filter output is an explicit prediction of the effects of the animal’s own movements on the sensory signal. This prediction is subtracted from the raw sensory input to provide an estimate of the sensory signal generated by objects in the external world. Learning to predict the sensory effects of movement has been suggested as a central cerebellar function (e.g. Miall & Wolpert, 1996). A homely example concerns the difficulty of tickling oneself: the argument is that this difficulty is caused by a prediction of the sensory effects of one’s own movement that is used to diminish the actual sensory effects (Blakemore, Wolpert, & Frith, 1998).

The connectivity requirements for a hypothetical noise cancellation module are thus:

1. Mossy fibre input to the cerebellum is an efferent copy of the motor commands.
2. The error signal required for learning is the noise-cancelled signal.

3. Cerebellar output is an estimate of the self-produced sensations caused by the animal’s own movement, and must target a comparator to produce the noise-cancelled signal. Efferent copy of this signal must be available on the climbing fibre to provide the teaching signal in Fig. 5.

The signals required by the adaptive element in this architecture appear to be available biologically, and its computational effectiveness has been tested in the context of a rat-like whisking robot, using only observed whisker signals and a copy of motor commands as inputs (Anderson et al., 2010). The next question therefore concerns whether the circuit is in fact implemented biologically. This question is addressed in Section 4, for a particular form of noise cancellation that can be used for the detection of novel stimuli.

3.1.2. Smith predictor

In engineering systems the preferred control option for tracking a desired response in a motor system is usually a feedback controller as shown in Fig. 6(A). The motor plant output is compared with the desired response and the error is fed into a high-gain controller (labelled ‘simple controller’ in Fig. 6(A)) to produce the required motor command. To achieve stable and accurate tracking of time varying signals requires careful choice of controller, and there are powerful standard techniques for designing such controllers when a plant model is known.

However if the plant output is subject to delay the use of high-gain controllers can lead to instability and these standard control design techniques break down. In the control engineering literature a scheme known as the Smith predictor (Fig. 6(B)) is often used to overcome the problems associated with designing feedback controllers for systems with delay. The key benefit to the designer is that it allows a simple controller to be designed for the delay-free plant using standard methods. This controller is then used as a component in the Smith predictor (Fig. 6(C)) which also requires a forward model of the plant and a model of the delay. This architecture is guaranteed to give stable control of the delayed plant when the plant and delay models are accurate.

It has been suggested (Miall et al., 1993) that the cerebellum could act as a Smith predictor, compensating for the unavoidable time-delays arising from sensorimotor transmission and central processing of neural signals in biological systems. It is important in assessing this suggestion to make Marr’s distinction between the computational and the algorithmic levels of description. Computationally the Smith predictor implements a control scheme in which the desired response is reproduced with a fixed temporal delay. It is plausible that such a controller could be learned directly by the cerebellar microcircuit in a suitable architecture and

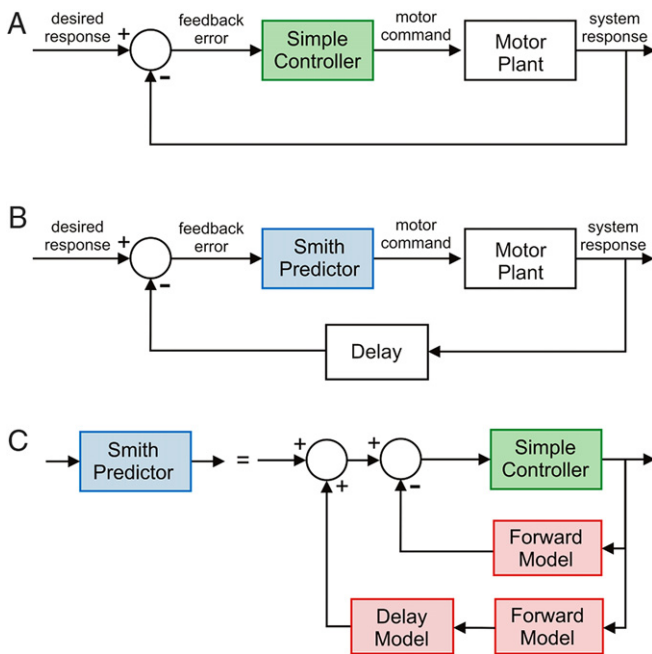


Fig. 6. The Smith predictor. A: Feedback control loop. A controller driven by feedback error causes a motor plant to track a desired response. This control architecture is widely used in engineered systems, and there are standard procedures for designing suitable controllers when the plant is not subject to delay. B: The Smith-predictor control loop. The control loop is now complicated by the presence of a delay, here shown in the sensory feedback pathway. The presence of a delay complicates control but a Smith predictor can be used to simplify the design process, which in effect allows the engineer to use a simple controller designed for a plant without delay. Note that for linear systems, the sum total of delays through a control loop (motor, sensory and plant) can be lumped into a single delay component, which has an equivalent effect on control performance when placed at any location in the loop. C: Smith-predictor controller architecture. The Smith predictor comprises two forward models of the plant, one with delay and one without. Motor commands are primarily generated through control of the undelayed forward model (the inner loop), thereby avoiding control problems associated with delay. However, the sensory feedback loop shown in (B), is required for the compensation of disturbances. Therefore, the outer loop of the Smith predictor (through the delayed forward model) is used to cancel the feedback error for zero disturbances, which permits the transmission of unpredictable disturbances only to the inner control loop (assuming an exact forward model of the plant).

this possibility should be investigated. However the term Smith predictor implies a particular algorithm based on the architecture shown in Fig. 6(C) (or one of its variants) with specific modules implementing the forward plant and delay models found in that diagram. We are not aware that any biologically plausible scheme has been proposed that would allow these individual modules to be learned stably. For example, the teaching signal necessary for learning the plant forward model in Fig. 6(C) is the difference between the model input and the *undelayed* plant output, and the teaching signal necessary for learning the delay model is the difference between the (true) delayed and undelayed plant outputs. Neither signal appears in any natural way in the Smith predictor architecture, and substitution of similar signals containing the delayed components will necessarily lead to learning instability. Some of the stability problems in learning with delayed signals might be countered by a synaptic eligibility trace whose delay matched the plant delay accurately (Porrill & Dean, 2007a) but since this delay is an intrinsic property of the parallel fibre/Purkinje synapse, it could only cope with a limited range of delays (e.g. up to 200 ms Wang, Denk, & Häusser, 2000).

In the light of this analysis we highlight a number of challenges facing the Smith predictor hypothesis for which computational models have yet to be proposed: (i) How is an *undelayed* forward model learnt from signals that are necessarily delayed by

effluent/afferent processing? (ii) How is the Smith predictor architecture specifically subdivided in terms of forward model components and delays? (iii) In this architecture, what is the role and connectivity of the cerebellum? (iv) How are sensory error signals distributed in order to drive learning in separate Smith predictor components (forward model and delay model)? Answering these questions could well provide key insights into the handling of efferent/afferent delays in motor control.

In addition the configuration shown in Fig. 6(C) requires separate models of the controlled plant and of the delay (Miall et al., 1993) suggested a slightly more plausible organisation but one in which the plant and delay are still represented separately). This modular configuration is an unlikely outcome of learning with plausible teaching signals, since it requires a division of the sensory error into individual components due to errors in plant and delay models. Although it seems unlikely that regions of the cerebellum correspond in any direct way to the modular components implied by the Smith predictor architecture, it would be interesting to investigate whether individual components of this architecture – for example the loop via a forward model, which is very similar to the internal feedback loop proposed by Robinson (e.g. Robinson, 1981) for the control of saccades – have plausible cerebellar implementations.

3.1.3. State estimator

The input–output characteristics of a system such as a motor plant often depend on the past history of the control inputs as well as their current values. The state vector of such a system is defined to be a set of variables which completely describes the system, including these ‘historical influences’, such that the control inputs, together with the initial state, completely determine the future behaviour of the system. Physical output from the model can be regarded as a measurement of some combination of these internal state variables. Having access to the internal state of a system is very useful when designing controllers and has important theoretical consequences. For example it can be shown that the optimal controller for many stochastic systems can be implemented as a feedback controller with access to the internal state (Todorov, 2005).

If an accurate theoretical model of the system is unavailable the only way to determine its internal state is by inference from measurements on the system. When the system is linear an optimal (in the sense of least squares) algorithm for state estimation is known, which can be implemented using a Kalman filter in the architecture shown in Fig. 7.

Implementation of this state-estimator requires a great deal of knowledge about the observed system. Firstly we need to know the state update model, which describes how the internal state of the plant changes over time both autonomously, and under the influence of control inputs and system noise. We also need the measurement model, which describes how measurable quantities are related to the system’s internal state, and also includes a model of measurement noise. The state estimator uses these two models to estimate the system’s current state; updating the current estimate of the state vector over time using the state update model and improving this estimate when new measurements arrive. This process uses a Kalman filter algorithm, which determines the optimal values for the gain matrix which determines the influence of the current measurements on the state-estimate. Clearly this influence should be small when the measurements are inaccurate relative to the expected error in the internally predicted state, but large in the opposite case. In essence the Kalman gains determine the relative weighting to be given to the estimates from internal models compared to external measurements.

Kording, Tenenbaum, and Shadmehr (2007) have investigated the use of state estimation in biological motor control systems

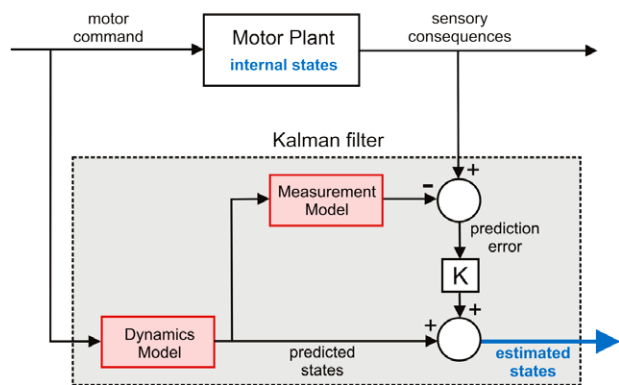


Fig. 7. The Kalman filter. Architecture of a Kalman filter state-estimator. The Kalman filter combines predictive information with sensory observations to produce an optimal estimate of the plant's internal state. The predictive element of the Kalman filter is similar to a typical forward model but the forward model is split into two components: (i) a dynamic state model that predicts future states from current values – the dynamics model, and (ii) a static model that maps states to sensory signals – the measurement model (each highlighted in red, since they could be learned adaptively). The motor commands that drive the plant are also provided as input to the dynamics model to predict the next state. When measurements become available the predicted state is additively combined with the sensory prediction error weighted by the Kalman gain matrix (K), to produce an improved state estimate. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

in the context of saccadic gain adaptation. Saccades become inaccurate if the gain of the saccadic plant varies over time. Such gain changes can arise from a variety of internal causes such as muscle fatigue, or developmental changes in muscle strength. These individual components of the muscle state can have different characteristic time constants and sizes. For example fatigue may cause large, relatively rapid variations while developmental changes cause slower, smaller variations. Although observing saccadic performance only gives information about the combined effect of these components the Kalman filter architecture outlined in Fig. 7 is capable of estimating the contributions of the individual components and making an optimal gain estimate. Comparing the behaviour of this 'ideal observer' with actual performance has shown that in many cases the gains used by real observers are the optimal gains characteristic of state estimation methods.

Studies of this kind can indicate that the system as a whole can approximate state-estimation, but do not indicate whether the site of state-estimation is, as has been proposed, the cerebellum (Miall et al., 2007; Miall & Wolpert, 1996). As far as we are aware no detailed mechanisms or architectures for the biological implementation of state estimation methods have been proposed, so we are unable to give a detailed analysis of the required connectivity or learning characteristics. In particular, no scheme has been suggested that could implement the detailed matrix equations required to estimate the optimal Kalman gain matrices required in the general case.

The argument so far has focused on state-estimation in the context of hidden or unobservable states of a motor plant, states which may in engineering applications have no physical embodiment whatsoever but which have direct relevance to motor control. However, state-estimation as a generic process for combining measurements and model-based predictions can be applied to a very wide range of tasks, including prey localisation (by precerebellar structures in electric fish Paulin, 2005) or estimating the position and velocity of an arm (Miall & Wolpert, 1996). Since the state update model in Fig. 7 is essentially a forward plant model which has to be adaptive to track the true plant, it could in principle be implemented as a (cerebellar) adaptive filter. Perhaps an approximation to the required Kalman gains could also be implemented with an adaptive filter (though Paulin, 2005 used

a particle-filter scheme for estimating prey location that is very unlike the Kalman filter algorithm). As with the Smith predictor, it seems unlikely that in a biological controller these processes could be partitioned as neatly as in the theoretical state estimator. However, future work may show how the overall functionality of these architectures may still be achievable in a simplified form that makes possible their implementation by an adaptive filter.

3.2. Inverse models

We have seen that a feedforward controller can be implemented using the architecture shown in Fig. 3(A). As implemented in its simplest form in Fig. 3(B) this requires the cerebellar adaptive filter C to learn the inverse $C = P^{-1}$ of the motor plant P , hence the name adaptive inverse control.

In the architecture shown in Fig. 3(B) the teaching signal required for supervised learning is the error in adaptive filter output, sometimes called proximal error. Although this error in filter output causes the errors in system output (distal or sensory error) that will be measured by sensors, the two are not usually related in a simple way. The fact that it is only sensory (distal) error that is directly available to a learning system rather than learning element output (proximal) error is called the distal error problem: it is a generic problem in neural net learning systems (Jordan, 1996; Jordan & Wolpert, 2000).

In motor systems, if an adaptive-filter output contributes to a motor command, its output error is called motor error. In this case the distal error problem is also called the motor error problem, and is directly relevant to the issue of biological plausibility of cerebellar internal models. In Fig. 3(B) the teaching signal required is motor error—the difference between actual and desired motor commands to the motor plant. As we have noted this error signal is not directly observable since it passes through the motor plant before producing observable sensory consequences. Nor can the desired motor commands be known in advance in biological systems—hence the circuit shown in Fig. 3(B) as it stands is not biologically plausible.

Two basic solutions to the motor error problem for cerebellar learning have been proposed, motor error learning (Section 3.2.1) and recurrent architecture (Section 3.2.2). In both cases we must take account of the cerebellum acting in conjunction with other control pathways since cerebellar lesions do not completely abolish function. This arrangement appears to be a general feature of cerebellar control, as seen for example in the vestibulo-ocular reflex (Section 4.2.2) where the supplementary controller can be localised in the brainstem (for this reason we will term this non-cerebellar component B).

3.2.1. Motor error learning

The first solution tackles the motor-error problem head on: the required motor error is calculated from sensory errors by hypothetical neural structures called reference structures (Fig. 8(A)). This approach has been most strongly advocated by Kawato and co-workers and implemented in the feedback-error-learning algorithm (Gomi & Kawato, 1992, 1993; Kawato & Gomi, 1992, 1993), where the estimated motor error can also be used in online feedback mode to improve movement accuracy (not shown in figure). In this architecture, taking into account the contribution from B , the cerebellum learns a partial inverse model, $C = P^{-1} - B$.

Because motor error is related to sensory error by the plant transfer function P , sensory error is not a satisfactory training signal unless the plant itself satisfies the Strictly Positive Real condition (Section 3.1). This constraint severely limits the applicability of the circuit shown in Fig. 8(A) if sensory error were used directly as a teaching signal, especially in multi-joint configurations. Feedback-error-learning attempts to solve this problem by

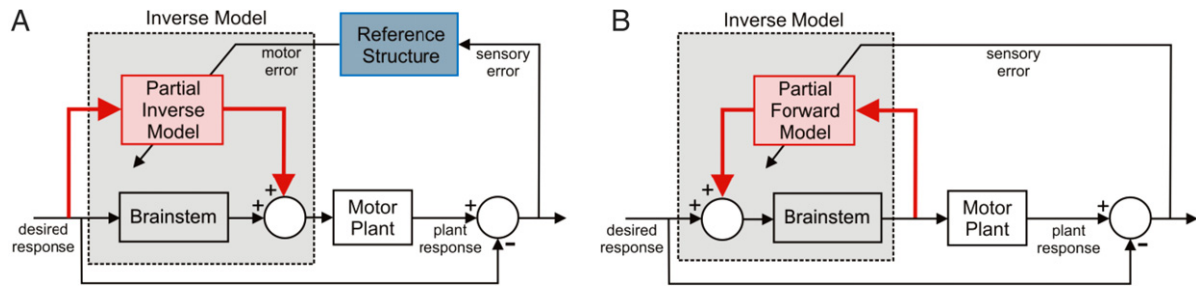


Fig. 8. Alternative architectures for adaptive inverse plant compensation. A: Motor-error architecture. The motor command is produced by a combination of a fixed element (called B in the text and representing e.g. the brainstem in the case of the VOR) and an adaptive element C representing the cerebellum. Together the fixed and adaptive components combine to form an inverse model $P^{-1} = B + C$ of the plant. In order to learn the partial inverse model $C = P^{-1} - B$, the required teaching signal is motor error. However, the motor error is unavailable, and therefore must be estimated from the sensory error by filtering the sensory error through a reference structure that approximates an inverse model of the plant. Since an (approximate) inverse model of the plant must be known in order to learn an inverse model of the plant this reasoning is circular for all but very simple plants. B: Recurrent architecture. This differs from (A) in the direction of the cerebellar arcs (highlighted in red). The cerebellum receives copies of motor command as input and its output is added to the desired response signal. Unlike the forward architecture in (A), the sensory error is used directly, thus avoiding the circularity mentioned above. Although the brainstem and cerebellum combined constitute an inverse model, the cerebellum itself learns a partial forward model $C = B^{-1} - P$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

'back-propagating' sensory error through an approximate inverse of the motor plant (as shown by the blue box labelled 'Reference Structure' in Fig. 8(A)) to obtain an estimate of motor error. Hence the connectivity requirements for this architecture are:

1. Mossy fibre input to the cerebellum consists of a copy of the desired motion.
2. The error signal required for learning is motor error. This signal must be made available on the climbing fibre by passing the sensory error through a suitable reference structure.
3. Cerebellar output must target a brain region involved in transmitting the motor command, and makes an open-loop contribution to this command.

Fig. 8(A) illustrates the central problem with this scheme, which is the requirement for an inverse model (in the Reference Structure) to learn an inverse model. Feedback-error-learning seeks to escape this apparent circularity by using a relatively simple approximation to the true inverse model, such as a single gain term to transform sensory into motor error.

However, although such approximations have been shown to work well in particular theoretical and robotics contexts (e.g. Dean, Mayhew, Thacker, & Langdon, 1991; Kawato, Furukawa, & Suzuki, 1987; Kawato & Gomi, 1993; Miyamura & Kimura, 2002; Nakanishi & Schaal, 2004) they run into problems for more complex plants. Thus, a feedback-error-learning model of the cerebellar control of reaching movements requires that inferior olivary cells "detect 'torque-like' errors in performance" (Schweighofer, Spolstra, Arbib, & Kawato, 1998, p. 99), and a recent application of feedback-error-learning to control of 7-dof robot arm (Tolu, Vanegas, Luque, Garrido, & Ros, 2012) uses a teaching signal explicitly related to individual errors in joint angle. The inescapable difficulty is the *general* constraint that the reference structure required for stable learning (i.e. an implicit inverse Jacobian) be of similar complexity to the inverse model. This point has been made in the context of robotics (e.g. Porrill & Dean, 2007b), with the claim that feedback-error-learning "is rarely used in the literature" (Schenck, 2011, p. 8). For biological plants it has been argued that with multiple-input multiple-output redundant systems such as the vestibulo-ocular reflex in three dimensions, the connectivity needed by the reference structure becomes infeasibly complex—how exactly should a vertical retinal-slip signal be channelled to the six oculorotatory muscles (e.g. Fig. 3 of Porrill, Dean, & Stone, 2004)? As Gomi and Kawato (1992, p. 112) point out: "The most interesting and challenging theoretical problem (raised by FEL) is setting an appropriate inverse reference model in the feedback controller at the spinal and brainstem levels".

A second drawback of the architecture shown in Fig. 8(A) is that it requires the climbing-fibre signal to cerebellar cortex to be an estimate of *motor* error. "Our view that the climbing fibres carry control error information, the difference between the instructions and the motor act, is common to most cerebellar motor-learning models; however ours is unique in that this error information is represented in motor-command coordinates" (Gomi & Kawato, 1992, p. 112). However, experimental investigations of climbing-fibre signals have emphasised their sensory nature (see e.g. Miall & Wolpert, 1996), related to touch, pain, muscle sense, or in the case of the vestibulo-ocular reflex retinal slip (references in Porrill & Dean, 2007b). The fundamental importance of this sensory signalling is reflected in the zonation scheme for cerebellar cortex, which derives from the organisation of inputs to cortex from the inferior olive that are segregated according to the sensory signals they convey (e.g. Voogd, 2011).

3.2.2. Recurrent architecture

The second proposal for a biologically realistic version of Fig. 3(B) uses a recurrent architecture (Fig. 8(B)) in which the input to the adaptive filter is a copy of the motor commands sent to the plant (Dean, Porrill, & Stone, 2002; Jordan, 1996; Porrill et al., 2004). This circuit is based on the actual connectivity of particular cerebellar microzones such as the flocculus. It can be shown that with this connectivity learning using sensory error is stable without the need for reference structures, for example by Lyapunov analysis (Porrill & Dean, 2007b). Note that in this configuration the cerebellum converges to $C = B^{-1} - P$, so that, despite the fact that the *overall* architecture embodies an adaptive inverse controller (Fig. 8(B)), the cerebellum itself implements an incremental *forward* model of the plant. This complexity emphasises that the precise nature of the model learnt by the cerebellar chip cannot be predicted without close analysis of the system connectivity.

The connectivity requirements for the recurrent configuration are:

1. Mossy fibre input to the cerebellum is an efferent copy of the motor commands.
2. The error signal required for learning is sensory error in model output (so no reference structure is required).
3. Cerebellar output must target the brain region producing the desired motion command, forming a closed-loop or recurrent architecture.

This architecture, which is consistent with a range of anatomical and neurophysiological evidence, has previously been shown

in simulations to be capable of solving the 3D VOR plant compensation problem (Porrill et al., 2004), and it has been shown to have significant theoretical advantages for the modular control of non-linear and redundant multi-degree of freedom systems (Porrill & Dean, 2007b). Recurrent architecture was emphasised early on by Eccles (1969) who called it the cerebellar 'dynamic loop', and is now regarded as a generic property of cerebellar connectivity (Middleton & Strick, 2000) of unknown function (Ramnani & Miall, 2001). The computational analysis presented here provides a possible answer. Cerebellar loops allow stable adaptive learning using only observable sensory errors. This allows the cerebellar microcircuit to be treated as a 'cerebellar chip' which can be plugged into a motor system to improve performance, without the need for complex, hard-wired reference structures.

3.3. Conclusions

This section has considered five architectures, three relating to forward models and two to inverse models. We argue that three of these architectures – at least in their present forms – appear to require signals that are not available biologically. The two architectures that do seem to be biologically plausible in principle, namely noise cancellation (forward model) and the recurrent architecture (inverse model) are considered further in the next section, which asks whether they are in fact implemented in neural circuitry. Here we address the distinction between the adaptive filter architectures suitable in engineering applications and those suitable for biological systems. The following appear to be important differences between biology and engineering.

1. The learning rule in biological systems is fixed by the characteristics of synaptic plasticity in the microcircuit. It cannot be adapted *ad hoc* to suit particular tasks, in the way engineers might choose specific methods from an extensive toolkit, for example supplementing an adaptive back-stepping design with a particular anti-windup scheme, to meet the needs of a specific task.
2. The cerebellar chip metaphor should not be over-extended. In engineering applications it is possible to learn a particular component, such as an inverse model (this learning will often be off-line) and then plug the learned component into another slot in a control architecture. In contrast current evidence suggests that the cerebellar chip must learn *in situ*, and must continue to function during the learning process. If correct this simple observation rules out many commonly used engineering solutions to control problems.
3. The teaching signals required for adaptive control and signal processing must be biologically plausible. In many cases these signals will indicate performance errors (but not always, cf. noise cancellation) since these indicate the need to change system parameters. These errors will be provided originally by biological sensors, such as vision or touch, and they should as far as possible be used as teaching signals without requiring extensive processing based on task characteristics (in the way that engineers can choose between output error, equation error, filtered error, etc.). Although this criterion precludes the use of the complex reference structures required to recover motor error it does not exclude the possibility that extensive pre-filtering, together with other operations such as signal gating is implemented in pre-olivary structures.

4. Chip connectivity: biological implementation

In this section we investigate to what extent the two architectures that seem biologically plausible in principle, namely noise cancellation (forward model) and the recurrent architecture (inverse model), are implemented by neural circuitry in practice.

4.1. Forward models in noise cancellation

As mentioned previously, a number of suggestions have been made about the cerebellum's role in learning an internal model that can distinguish self-produced sensory signals from those driven by external stimuli, as in the example of tickling. However, perhaps surprisingly, these suggestions are at present very rarely accompanied by descriptions of the detailed neural circuitry that would be required for such a role. Assessing how far noise-cancellation is in fact carried out by a cerebellar model is therefore a difficult task. Here we consider two possible examples. The first is a hypothesis concerning the circuitry underlying a particular variant of noise-cancellation involving the cerebellum, namely detection of novel vibrissal inputs by the whisking rat. The second concerns experimental evidence directly implicating cerebellar-like structures in noise cancellation in electric fish.

4.1.1. Noise cancellation and novelty detection

The architecture for novelty detection shown in Fig. 9(A) is similar to that of Fig. 5(B), but with the addition of a copy of the sensory signal being sent to the adaptive element. The specific application considered here is the detection of unexpected whisker contacts by rats, so the components of Fig. 5(B) have been re-labelled appropriately. Since the adaptive filter now has information about previous sensory inputs, it is able to cancel any component of the signal predictable from its past history. That is, it becomes a novelty detector, responsive in the whisking application only to new contacts, or sharp changes in the characteristics of a contact.

The connectivity requirements for a novelty-detection module are therefore:

1. Input to the adaptive filter is an efferent copy of the relevant motor commands, together with the original sensory signal.
2. Filter output is an estimate of self-noise, and must target a comparator to produce the noise-cancelled signal.
3. The efferent copy of this signal provides the teaching signal for the filter.

The signal from the whiskers is contaminated by 'self-noise' signals that are generated by the animal's own exploratory movements of the whiskers ('whisking'). The nature of this contamination cannot be known directly, but there is information about the whisking movements themselves, provided by the motor commands sent to the muscles that move the whiskers. This 'efference copy' information in effect is the reference information about the noise source, and so can be used as input to an adaptive filter (with proposed location in zone A2 of the cerebellum) that learns to mimic the transformation of motor commands into sensory signals from the whiskers. Thus, the cerebellum learns a forward model of 'Plant + Sensor Dynamics', including its basic properties of elasticity, viscosity and inertia, together with any post-processing in the sensory apparatus (Anderson et al., 2012).

Once learning has been achieved, cerebellar output would be an explicit prediction of the sensory signal including the effects of the animal's own movements on the sensory signal provided by the whiskers. Evidence suggests that this prediction could be subtracted in the superior colliculus from the raw sensory input to provide an estimate of unexpected whisker signal generated by objects in the outside world. The circuit shown in Fig. 9(A) is consistent with the known connectivity of the superior colliculus, inferior olive and cerebellum in rat, and with the established role of the superior colliculus in detecting (and orienting to) novel vibrissal stimuli. Further experiments are needed to establish whether it corresponds to the firing patterns of the relevant neurons in those structures (Anderson et al., 2012).

The novelty detection architecture illustrates an important theoretical point. To a large extent the function of a cerebellar chip is fixed by its climbing fibre input and its output target. In

- Shmuelof, L., & Krakauer, J. W. (2011). Are we ready for a natural history of motor learning? *Neuron*, *72*(3), 469–476. <http://dx.doi.org/10.1016/j.neuron.2011.10.017>.
- Skavenski, A. A., & Robinson, D. A. (1973). Role of abducens neurons in vestibuloocular reflex. *Journal of Neurophysiology*, *36*, 724–738.
- Tao, T., & Su, S. F. (2011). CMAC-based previous step supervisory control schemes for relaxing bound in adaptive fuzzy control. *Applied Soft Computing*, *11*(8), 5715–5723. <http://dx.doi.org/10.1016/j.asoc.2011.03.019>.
- Todorov, E. (2005). Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system. *Neural Computation*, *17*(5), 1084–1108. <http://dx.doi.org/10.1162/0899766053491887>.
- Tolu, S., Vanegas, M., Luque, N. R., Garrido, J. A., & Ros, E. (2012). Bio-inspired adaptive feedback error learning architecture for motor control. *Biological Cybernetics*, *106*(8–9), 507–522. <http://dx.doi.org/10.1007/s00422-012-0515-5>.
- Voogd, J. (2011). Cerebellar zones: a personal history. *Cerebellum*, *10*, 334–350. <http://dx.doi.org/10.1007/s12311-010-0221-6>.
- Wang, S. S.-H., Denk, W., & Häusser, M. (2000). Coincidence detection in single dendritic spines mediated by calcium release. *Nature Neuroscience*, *3*(12), 1266–1273.
- Widrow, B., & Stearns, S. D. (1985). *Adaptive signal processing*. Engelwood Cliffs, NJ: Prentice-Hall Inc.
- Wolpert, D. M. (1997). Computational approaches to motor control. *Trends in Cognitive Sciences*, *1*(6), 209–216.
- Wolpert, D. M., & Ghahramani, Z. (2000). Computational principles of movement neuroscience. *Nature Neuroscience*, *3*(Suppl.), 1212–1217.
- Wolpert, D. M., Ghahramani, Z., & Jordan, M. I. (1995). An internal model for sensorimotor integration. *Science*, *269*, 1880–1882.
- Wolpert, D. M., & Kawato, M. (1998). Multiple paired forward and inverse models for motor control. *Neural Networks*, *11*(7–8), 1317–1329. doi: S0893-6080(98)00066-5. [pii].
- Wolpert, D. M., Miall, R. C., & Kawato, M. (1998). Internal models in the cerebellum. *Trends in Cognitive Science*, *2*(9), 338–347.




Fig. 10. Early proposal for inverse model of oculomotor plant. This diagram, a reproduction of Fig. 4 from Skavenski and Robinson (1973), shows a simplified model of the vestibulo-ocular reflex. The head-velocity signal (R_v) from the semi-circular canals (SCC) arrives at the vestibular nuclei (VN) and is transmitted directly with gain T via the medial longitudinal fasciculus (mlf) to the oculomotor neurons (OMNs). This signal is also integrated to provide an eye position signal that is sent to the OMNs, to compensate for the elasticity of the oculomotor plant that would otherwise pull the eye back to its starting position when the head movement was finished. The combination of direct and integrator pathways in effect constitute an inverse model of the first-order oculomotor plant.

fibres arising from cells in the inferior olive. But the more complex arrangement provides an explicit estimate of sensory interference that is available for distribution to appropriate targets in the rest of the brain. Also, cerebellar output is no longer constrained to act as a teaching signal, so is freed for other purposes such as cancelling interference by moving the sensor in question, as in the vestibulo-ocular reflex (Dean et al., 2013). However, overall the operations of cerebellar-like structures bear a striking resemblance to those of adaptive noise-cancelling architectures, a resemblance that adds plausibility to the idea that the cerebellum is involved in learning forward models.

4.2. Inverse models

The basic idea of the inverse model as shown in Fig. 3(A) is probably most familiar in the context of the oculomotor system, although not under that name. Skavenski and Robinson (1973) proposed that the head-velocity signal from the semi-circular canals needed to be partly integrated to compensate for the elasticity of the oculomotor plant and provide an accurate vestibulo-ocular reflex (Fig. 10). Extensive subsequent work on the oculomotor system has provided perhaps the best example of the difficulties of mapping concepts such as inverse model onto the complicated realities of neural circuitry. We consider here three of these difficulties (see also Lisberger, 2009). The first two are relevant to any implementation of an inverse model, whereas the third is related to the recurrent architecture.

4.2.1. Distributed plant and motor commands

In diagrams such as Fig. 1 the properties of motor commands and the motor plant are usually represented by relatively simple equations. The reality however is much more complex (Fig. 11). This diagram is itself highly simplified, in that the mechanical and geometrical properties of orbital tissue are not shown, properties that are now known to be much more intricate than originally thought (e.g. Dean & Porrill, 2008a; Demer, 2006; Schutte, van den Bedem, van Keulen, van der Helm, & Simonsz, 2006). It does however show an eye-muscle as composed of many motor units, and that the force each produces depends on both the firing rate of its motoneuron, and the muscle unit's own strength (e.g. as measured by the maximum tetanic force it is capable of exerting).

It appears that in some muscles the control of such an array of motor units is simplified by the 'size principle', whereby strength of muscle unit is correlated with the size of the motoneuron that controls it (Henneman & Mendell, 1981). Since motoneuron

size is related to firing threshold, this principle ensures that even if all motor units receive the same motor-command signal (common drive, De Luca & Erim, 1994), muscle units are recruited automatically in order of increasing strength as the motor command increases. There is evidence that in oculomotor muscles also motor units are recruited in order of increasing strength (Barmack, 1977; Dean, Porrill, & Warren, 1999; Gamlin & Miller, 2012), but modelling studies indicate that at least in primates this is unlikely to be achieved by sending a common synaptic drive to all motoneurons (Dean, 1997; Hazel, Sklavos, & Dean, 2002). The central problem is the large oculomotor range in primates (about $\pm 45^\circ$), which in turn means that the position thresholds at which motoneurons fire range from $< 45^\circ$ in the off-direction of the muscle to $\sim 20^\circ$ in the on-direction. The range of motoneuron intrinsic properties required to produce such a range of thresholds is much larger than that observed experimentally, indicating that low and high threshold motoneurons do not receive the same synaptic drive (Dean, 1997; Hazel et al., 2002).

It therefore appears that in the case of the oculomotor plant, the motor command is a distributed quantity which has a complex internal organisation. This means in turn that the output of any inverse model, represented in simple form in Fig. 3, must also be distributed and organised in a complex manner.

4.2.2. Distributed premotor signals

A further complication in the case of the oculomotor plant is that the motor commands for controlling conjugate horizontal eye-position arise from three different sources (Lisberger, Pavelko, & Broussard, 1994; McFarland & Fuchs, 1992; Scudder & Fuchs, 1992). One is a population of neurons in the vestibular nuclei termed 'position-vestibular-pause' cells, or PVPs (Fig. 12). The second is from cells sometimes called 'eye-position' cells or EPs, situated in the nucleus prepositus hypoglossi now identified as the site of the integrator for horizontal eye-movements postulated by Skavenski and Robinson (1973). The third source of position signals is from floccular target neurons or FTNs, which relay information from the flocculus to ocular motoneurons (OMNs).

Modelling how these inputs combine to form observed patterns of OMN firing (Hazel et al., 2002) suggests that the PVPs provide an excitatory drive that is similar for all OMNs, whereas EPs provide an inhibitory signal that is weak for low-threshold OMNs and strong for high-threshold OMNs, a pattern that acts to increase the range of OMN thresholds. Inhibitory FTNs appear to provide a signal similar to that of the EPs, suggesting this is how the adaptive element of the inverse model (the flocculus) influences OMN firing.

Fig. 12 thus indicates that the inverse model required for oculomotor plant compensation is spread over a number of neural structures, and cannot be identified with any single one of them.

4.2.3. Flocculus has multiple, partial roles

The connectivity of the flocculus shown in Fig. 12 corresponds partly to the recurrent architecture shown in Fig. 8(B), but has additional elements, in particular a vestibular input (mainly via floccular projecting neurons or FPNs in the vestibular nuclei). This gives the flocculus additional functionality. As part of an inverse model for oculomotor plant compensation, it needs to adjust all eye-movement signals sent to the OMNs, of whatever origin. These are provided by the recurrent loop from brainstem and OMNs via neurons in the paramedian tract nuclei (PMT) (e.g. Büttner-Ennever, 1988; Dean & Porrill, 2008b). However, the existence of a specific input from the vestibular system allows the flocculus to compensate for changes that are confined to that system (e.g. a change in semi-circular canal properties). This additional role has been modelled by Haith and Vijayakumar (2009) and corresponds to the flocculus functioning as a partial inverse model of a sensory system.

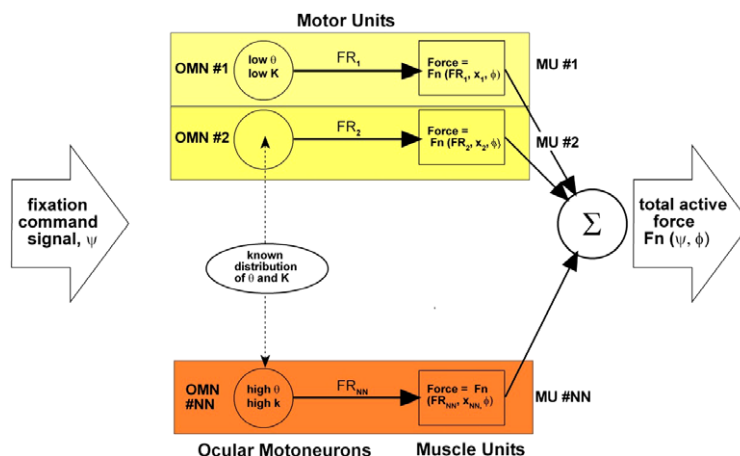


Fig. 11. Schematic diagram of the motor units in an extraocular muscle (EOM). EOMs consist of many motor units ($NN \approx 1000$ in primate) arranged mainly in parallel. In most motor units, a muscle unit is controlled by its ocular motoneuron (OMN). Muscle units vary in 'strength' x_i , and the force that they produce is a function of strength, the firing rate FR of the parent OMN, and muscle length which is determined by eye position (ϕ). During sustained fixation at ϕ , OMNs with a threshold θ that is less than ϕ do not fire at all, while for the others $FR = k_i(\phi - \theta_i)$. These firing rates are determined by the synaptic drive to the motoneurons, shown here as a fixation command signal ψ . Evidence suggests that this drive varies from motoneuron to motoneuron (see text).
Source: From Fig. 1 of Dean (1996).

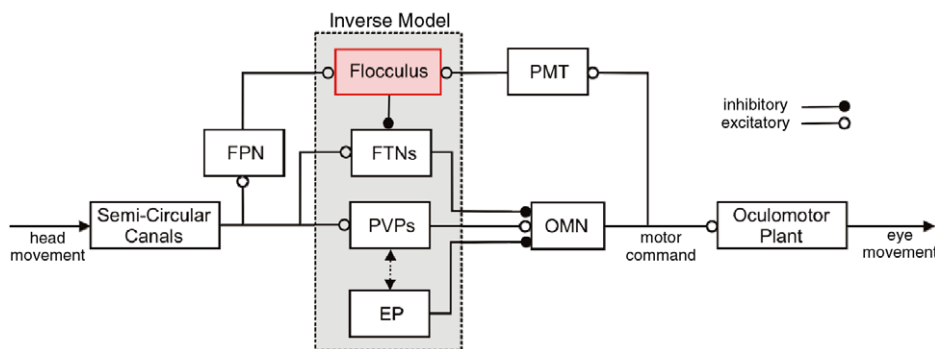


Fig. 12. Sources of premotor drive to ocular motoneurons (OMNs). Conjugate eye-position related firing rates in abducens OMNs are driven from three main sources. 1. Position-vestibular-pause neurons (PVPs) in the medial vestibular nucleus, which are part of the main pathway for the vestibulo-ocular reflex. 2. Floccular target neurons (FTNs), part of a side path for the vestibulo-ocular reflex. These neurons receive an input from the flocculus. 3. Eye-position neurons (EPs) in the nucleus prepositus hypoglossi, constituting the horizontal 'integrator'. The functional connections between EPs and the other two classes of neurons are obscure. Two inputs to the flocculus are shown: one from the semi-circular canal via floccular projecting neurons (FPNs) in the vestibular nuclei, the other from the nuclei of the paramedian tract (PMT) which conveys an efferent copy of the commands sent to the plant. It appears that all three sets of brainstem neurons projecting to the OMNs must be regarded as part of the inverse model, together with the cerebellar flocculus itself (grey box). The flocculus (red) is the adaptive part of the inverse model, as well as the synapses to FTNs conveying vestibular signals. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

A further related function for the flocculus is indicated by evidence relating it to the optokinetic reflex, which is a feedback reflex driven by retinal slip (Carpenter, 1988). Because of delay in the retinal slip signal (~ 100 ms) this reflex can only have substantial gains at low frequencies, but there it complements the vestibulo-ocular reflex which is more effective at higher frequencies. Since the optokinetic reflex is a feedback control system, as illustrated in a simple form in Fig. 6(A), the likely function of the flocculus is to tune the 'simple controller' shown in the diagram to achieve the optimal gain at each frequency given the delays in retinal slip processing, replacing the Smith-predictor architecture shown in Fig. 6(B). This function of the flocculus does not fit easily into categories such as 'forward' or 'inverse' models. A scheme of this kind has been described in Anderson, Porrill, and Dean (2007).

In primates, the contribution of the flocculus to following visual stimuli is enormously expanded, allowing small moving targets to be fixated in 'smooth pursuit' (Carpenter, 1988). Although related to overall gaze stabilisation (inasmuch as it constitutes stabilising gaze on a particular target) it represents a major functional addition, reflected in new classes of Purkinje cell firing patterns found in primates (Lisberger, 2009).

Finally, although the flocculus is represented as the adaptive element in Fig. 12, extensive evidence indicates the existence of plastic synapses in the brainstem itself, in particular vestibular synapses on FTNs. It appears that floccular output itself can act as a teaching signal to guide this plasticity (Boyden, Katoh, & Raymond, 2004; Lisberger, 1998).

5. Conclusions

The analysis presented here clarifies the relationship between the adaptive filter model of the cerebellar microcircuit and descriptions of cerebellar function using the terminology of internal models. It is easy to confuse an adaptive filter with a forward model, even though this is a category mistake. The adaptive filter implemented by a cerebellar microzone need not function as a forward model, but a forward model embodied in the cerebellum necessarily has an adaptive filter as its learning element (we know of no alternative mechanism which will provide the learning processes assumed for forward models). In fact the adaptive filter can be used as a component of many adaptive control schemes (hence the cerebellar chip metaphor). As such it is not to be identified with any particular scheme such as a

'forward model' or a 'state-estimator'; it is both lower-level and more general. In order to learn internal models the cerebellum draws upon the power of the adaptive filter, implemented in an appropriate connectivity.

If the adaptive filter model is correct then the cerebellar chip is indeed a very powerful computational device, as evidenced by the range of applications of the adaptive filter in engineering. We describe this property as 'computational adequacy'. Computational adequacy should be regarded as an essential feature of any proposed model of cerebellar function, given that the cerebellum is known to be involved in such a wide range of sensory, motor and possibly even cognitive, behaviours.

We have analysed the impact of the constraints imposed by biological implementation, such as the need for a suitable teaching signal, on the computational adequacy of the cerebellar chip. At first sight these constraints are quite limiting, since it seems that direct biological implementation of many useful control architectures, such as the Smith-predictor or state-estimation, is implausible. However we have seen that equivalent functionality can often be achieved in a different way, for example by embedding the adaptive filter element in a circuit with a well-chosen macroscopic connectivity, as seen in the recurrent architecture for plant inversion. In fact it seems that engineers may have something to learn from biological design principles. Current engineering design methodologies have access to an enormous range of sophisticated techniques, but control systems must be designed on a case-by-case basis to suit particular problems; this is often a non-trivial task, especially for compliant actuators and structures. In contrast biological systems achieve remarkable levels of performance for such systems using a relatively flat and homogeneous control structure.

Two further conclusions of the analysis should be emphasised. Firstly, we have looked at only a relatively small subset of the kind of tasks with which the cerebellum is known to be involved. This is because the relevant cerebellar connectivity is known in sufficient detail for very few cases (Dean & Porrill, 2011), and hence discussion of the internal model hypothesis is necessarily focused on a small number of examples. However, each example has suggested a different identification of the cerebellar chip with a particular kind of internal model. This indicates that cerebellar function in general is unlikely to be explained under a single, simple, heading such as 'state-estimation' or 'inverse modelling'. Secondly, more detailed analysis of a particular region, the flocculus, has shown that identification of function is likely to be further complicated by the fact that the cerebellum both works alongside other brain areas, and controls distributed systems, so its contribution will necessarily be incremental and distributed. In these situations it is more likely that cerebellar microzones, while implementing a common algorithm, learn task-specific adaptive filters which combine a number of roles.

Acknowledgements

Preparation of this article was supported by a European Union Grant (REALNET, 270434 FP7). SA was supported by a European Union Grant on Biomimetic Technology for Vibrissal Active Touch (BIOTACT) ICT-215910, and an EPSRC grant (EP/1032533/1).

References

- Albus, J. S. (1971). A theory of cerebellar function. *Mathematical Biosciences*, 10, 25–61.
- Anderson, S. R., Pearson, M. J., Pipe, A., Prescott, T., Dean, P., & Porrill, J. (2010). Adaptive cancellation of self-generated sensory signals in a whisking robot. *IEEE Transactions on Robotics*, 26(6), 1065–1076. <http://dx.doi.org/10.1109/tro.2010.2069990>.
- Anderson, S. R., Porrill, J., & Dean, P. (2007). Adaptive filter model of cerebellum improves performance of simulated optokinetic reflex. In *2007 abstract viewer/itinerary planner*. Washington, DC: Society for Neuroscience, (pp. Prog. No. 718.712).
- Anderson, S. R., Porrill, J., Pearson, M. J., Pipe, A., Prescott, T., & Dean, P. (2012). An internal model architecture for novelty detection: implications for cerebellar and collicular roles in sensory processing. *PLoS One*, 7(9), e44560.
- Barmack, N. H. (1977). Recruitment and suprathreshold frequency modulation of single extraocular muscle fibers in the rabbit. *Journal of Neurophysiology*, 40(4), 779–790.
- Bastian, A. J. (2011). Moving, sensing and learning with cerebellar damage. *Current Opinion in Neurobiology*, 21(4), 596–601. <http://dx.doi.org/10.1016/j.conb.2011.06.007>.
- Bell, C., Bodznick, D., Montgomery, J., & Bastian, J. (1997). The generation and subtraction of sensory expectations within cerebellum-like structures. *Brain Behaviour and Evolution*, 50(31), 17–31.
- Bell, C. C., Han, V., & Sawtell, N. B. (2008). Cerebellum-like structures and their implications for cerebellar function. *Annual Review of Neuroscience*, 31, 1–24.
- Blakemore, S. J., Frith, C., & Wolpert, D. (1999). Spatio-temporal prediction modulates the perception of self-produced stimuli. *Journal of Cognitive Neuroscience*, 11(5), 551–559.
- Blakemore, S. J., Frith, C. D., & Wolpert, D. M. (2001). The cerebellum is involved in predicting the sensory consequences of action. *NeuroReport*, 12(9), 1879–1884.
- Blakemore, S. J., Wolpert, D., & Frith, C. (1998). Central cancellation of self-produced tickle sensation. *Nature Neuroscience*, 1(7), 635–640.
- Boyden, E. S., Katoh, A., & Raymond, J. L. (2004). Cerebellum-dependent learning: the role of multiple plasticity mechanisms. *Annual Review of Neuroscience*, 27, 581–609.
- Büttner-Ennever, J. A. (Ed.) (1988). *Neuroanatomy of the oculomotor system* (vol. 2). Amsterdam: Elsevier.
- Carpenter, R. H. S. (1988). *Movements of the eyes* (2nd ed.). London: Pion.
- Cerminara, N. L., & Apps, R. (2011). Behavioural significance of cerebellar modules. *Cerebellum*, 10(3), 484–494. <http://dx.doi.org/10.1007/s12311-010-0209-2>.
- Cheng, K. H. (2011). Adaptive fuzzy CMAC-based nonlinear control with dynamic memory architecture. *Journal of the Franklin Institute—Engineering and Applied Mathematics*, 348(9), 2480–2502. <http://dx.doi.org/10.1016/j.jfranklin.2011.07.011>.
- Commuri, S., & Lewis, F. L. (1997). CMAC neural networks for control of nonlinear dynamical systems: structure, stability and passivity. *Automatica*, 33(4), 635–641. [http://dx.doi.org/10.1016/S0005-1098\(96\)00180-X](http://dx.doi.org/10.1016/S0005-1098(96)00180-X).
- Cullen, K. E. (2004). Sensory signals during active versus passive movement. *Current Opinion in Neurobiology*, 14, 698–706.
- Dean, P. (1996). Motor unit recruitment in a distributed model of extraocular muscle. *Journal of Neurophysiology*, 76(2), 727–742.
- Dean, P. (1997). Simulated recruitment of medial rectus motoneurons by abducens internuclear neurons: synaptic specificity vs. intrinsic motoneuron properties. *Journal of Neurophysiology*, 78, 1531–1549.
- Dean, P., Jörntell, H., & Porrill, J. (2013). Adaptive filter models. In M. Manto, D. Gruoi, J. D. Schmahmann, N. Koibuchi, & F. Rossi (Eds.), *Handbook of the cerebellum and cerebellar disorders* (pp. 1315–1335). Heidelberg: Springer Netherlands.
- Dean, P., Mayhew, J. E. W., & Langdon, P. (1994). Learning and maintaining saccadic accuracy: a model of brainstem–cerebellar interactions. *Journal of Cognitive Neuroscience*, 6(2), 117–138.
- Dean, P., Mayhew, J. E. W., Thacker, N., & Langdon, P. M. (1991). Saccade control in a simulated robot camera-head system: neural net architectures for efficient learning of inverse kinematics. *Biological Cybernetics*, 66, 27–36.
- Dean, P., & Porrill, J. (2008a). Eye orbital mechanics. In U. Windhorst, M. D. Binder, & N. Hirokawa (Eds.), *Encyclopedic reference of neuroscience* (pp. 1533–1539). Heidelberg: Springer Verlag.
- Dean, P., & Porrill, J. (2008b). Oculomotor anatomy and the motor-error problem: the role of the paramedian tract nuclei. *Progress in Brain Research*, 171, 177–186. [http://dx.doi.org/10.1016/S0079-6123\(08\)00624-9](http://dx.doi.org/10.1016/S0079-6123(08)00624-9).
- Dean, P., & Porrill, J. (2010). The cerebellum as an adaptive filter: a general model? *Functional Neurology*, 25(3), 1–8.
- Dean, P., & Porrill, J. (2011). Evaluating the adaptive-filter model of the cerebellum. *Journal of Physiology London*, 589(Pt 14), 3459–3470. <http://dx.doi.org/10.1113/jphysiol.2010.201574>. [pii].
- Dean, P., Porrill, J., Ekerot, C. F., & Jörntell, H. (2010). The cerebellar microcircuit as an adaptive filter: experimental and computational evidence. *Nature Reviews Neuroscience*, 11(1), 30–43. <http://dx.doi.org/10.1038/nrn2756>. [pii].
- Dean, P., Porrill, J., & Stone, J. V. (2002). Decorrelation control by the cerebellum achieves oculomotor plant compensation in simulated vestibulo-ocular reflex. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 269(1503), 1895–1904.
- Dean, P., Porrill, J., & Warren, P. A. (1999). Optimality of static force control by horizontal eye muscles: a test of the minimum norm rule. *Journal of Neurophysiology*, 81, 735–757.
- De Luca, C. J., & Erim, Z. (1994). Common drive of motor units in regulation of muscle force. *Trends in Neurosciences*, 17(7), 299–305.
- Demer, J. L. (2006). Evidence supporting extraocular muscle pulleys: refuting the platygian view of extraocular muscle mechanics. *Journal of Pediatric Ophthalmology and Strabismus*, 43(5), 296–305.

- Ebner, T. J., Hewitt, A. L., & Popa, L. S. (2011). What features of limb movements are encoded in the discharge of cerebellar neurons? *Cerebellum*, 10(4), 683–693. <http://dx.doi.org/10.1007/s12311-010-0243-0>.
- Eccles, J. C. (1969). The dynamic loop hypothesis of movement control. In K. N. Leibovic (Ed.), *Information processing in the central nervous system* (pp. 245–269). New York: Springer.
- Eccles, J. C., Ito, M., & Szentágothai, J. (1967). *The cerebellum as a neuronal machine*. Berlin: Springer-Verlag.
- Fujita, M. (1982). Adaptive filter model of the cerebellum. *Biological Cybernetics*, 45, 195–206.
- Gamlin, P. D., & Miller, J. M. (2012). Extraocular muscle motor units characterized by spike-triggered averaging in alert monkey. *Journal of Neuroscience Methods*, 204(1), 159–167. <http://dx.doi.org/10.1016/j.jneumeth.2011.11.012>.
- Gao, Z. Y., van Beugen, B. J., & De Zeeuw, C. I. (2012). Distributed synergistic plasticity and cerebellar learning. *Nature Reviews Neuroscience*, 13(9), 619–635. <http://dx.doi.org/10.1038/nrn3312>.
- Gomi, H., & Kawato, M. (1992). Adaptive feedback control models of the vestibulo-cerebellum and spinocerebellum. *Biological Cybernetics*, 68(2), 105–114.
- Gomi, H., & Kawato, M. (1993). Neural network control for a closed-loop system using feedback–error–learning. *Neural Networks*, 6, 933–946.
- Haith, A., & Vijayakumar, S. (2009). Implications of different classes of sensorimotor disturbance for cerebellar-based motor learning models. *Biological Cybernetics*, 100, 81–95.
- Hazel, T. R., Sklavos, S. G., & Dean, P. (2002). Estimation of premotor synaptic drives to simulated abducens motoneurons for control of eye position. *Experimental Brain Research*, 146(2), 184–196.
- Henneman, E., & Mendell, L. M. (1981). Functional organization of motoneuron pool and its inputs. In V. B. Brooks (Ed.), *Handbook of physiology, the nervous system, motor control, vol. II, Sect. I, Part 1* (pp. 423–507). Bethesda, Md.: American Physiological Society.
- Herculano-Houzel, S. (2010). Coordinated scaling of cortical and cerebellar numbers of neurons. *Frontiers in Neuroanatomy*, 4. <http://dx.doi.org/10.3389/fnana.2010.00012>. Article 12.
- Imamizu, H. (2010). Prediction of sensorimotor feedback from the efference copy of motor commands: a review of behavioral and functional neuroimaging studies. *Japanese Psychological Research*, 52(2), 107–120. <http://dx.doi.org/10.1111/j.1468-5884.2010.00428.x>.
- Imamizu, H., Kuroda, T., Miyauchi, S., Yoshioka, T., & Kawato, M. (2003). Modular organization of internal models of tools in the human cerebellum. *Proceedings of the National Academy of Sciences of the United States of America*, 100(9), 5461–5466.
- Ito, M. (1970). Neurophysiological aspects of the cerebellar motor control system. *International Journal of Neurology (Montevideo)*, 7, 162–176.
- Ito, M. (1997). Cerebellar microcomplexes. *International Review of Neurobiology*, 41, 475–487.
- Jordan, M. I. (1996). Computational aspects of motor control and motor learning. In H. Heuer, & S. Keele (Eds.), *Handbook of perception and action, volume 2: motor skills* (pp. 71–120). London: Academic Press.
- Jordan, M. I., & Wolpert, D. M. (2000). Computational motor control. In M. S. Gazzaniga (Ed.), *The new cognitive neurosciences* (2nd ed.) (pp. 601–618). Cambridge MA: MIT Press.
- Kawato, M. (1995). Cerebellum and motor control. In M. A. Arbib (Ed.), *The handbook of brain theory and neural networks* (pp. 172–178). Cambridge, MA: The MIT Press.
- Kawato, M. (1996). Learning internal models of the motor apparatus. In J. R. Bloedel, T. J. Ebner, & S. P. Wise (Eds.), *The acquisition of motor behaviors in vertebrates* (pp. 409–430). Cambridge MA: MIT Press.
- Kawato, M. (1999). Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology*, 9, 718–727.
- Kawato, M. (2008). From ‘understanding the brain by creating the brain’ towards manipulative neuroscience. *Philosophical Transactions of the Royal Society, Series B (Biological Sciences)*, 363(1500), 2201–2214.
- Kawato, M., Furukawa, K., & Suzuki, R. (1987). A hierarchical neural-network model for control and learning of voluntary movement. *Biological Cybernetics*, 57, 169–185.
- Kawato, M., & Gomi, H. (1992). The cerebellum and VOR/OKR learning models. *Trends in Neuroscience*, 15(11), 445–453.
- Kawato, M., & Gomi, H. (1993). Feedback–error–learning model of cerebellar motor control. In N. Mano, I. Hamada, & M. R. DeLong (Eds.), *Role of the cerebellum and basal ganglia in voluntary movement* (pp. 51–61). North Holland: Elsevier Science Publishers B.V.
- Kim, Y. H., & Lewis, F. L. (2000). Optimal design of CMAC neural-network controller for robot manipulators. *IEEE Transactions on Systems, Man, and Cybernetics Part C: Applications and Reviews*, 30(1), 22–31. <http://dx.doi.org/10.1109/5326.827451>.
- Kording, K. P., Tenenbaum, J. B., & Shadmehr, R. (2007). The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nature Neuroscience*, 10(6), 779–786.
- Lin, C. M., Chen, L. Y., & Yeung, D. S. (2010). Adaptive filter design using recurrent cerebellar model articulation controller. *IEEE Transactions on Neural Networks*, 21(7), 1149–1157. <http://dx.doi.org/10.1109/tnn.2010.2050700>.
- Lisberger, S. G. (1998). Physiologic basis for motor learning in the vestibulo-ocular reflex. *Otolaryngology—Head and Neck Surgery*, 119(1), 43–48.
- Lisberger, S. G. (2009). Internal models of eye movement in the floccular complex of the monkey cerebellum. *Neuroscience*, 162(3), 763–776. <http://dx.doi.org/10.1016/j.neuroscience.2009.03.059>.
- Lisberger, S. G., Pavelko, T. A., & Broussard, D. M. (1994). Responses during eye movements of brain stem neurons that receive monosynaptic inhibition from the flocculus and ventral paraflocculus in monkeys. *Journal of Neurophysiology*, 72(2), 909–927.
- Lu, C. J., & Wu, J. Y. (2011). An efficient CMAC neural network for stock index forecasting. *Expert Systems with Applications*, 38(12), 15194–15201. <http://dx.doi.org/10.1016/j.eswa.2011.05.082>.
- Marr, D. (1969). A theory of cerebellar cortex. *Journal of Physiology London*, 202, 437–470.
- McFarland, J. L., & Fuchs, A. F. (1992). Discharge patterns in nucleus prepositus hypoglossi and adjacent medial vestibular nucleus during horizontal eye movement in behaving macaques. *Journal of Neurophysiology*, 68(1), 319–332.
- Medina, J. F. (2011). The multiple roles of Purkinje cells in sensori-motor calibration: to predict, teach and command. *Current Opinion in Neurobiology*, 21(4), 616–622. <http://dx.doi.org/10.1016/j.conb.2011.05.025>.
- Miall, R. C., Christensen, L. O., Cain, O., & Stanley, J. (2007). Disruption of state estimation in the human lateral cerebellum. *PLoS Biology*, 5(11), 2733–2744.
- Miall, R. C., & Reckess, G. Z. (2002). The cerebellum and the timing of coordinated eye and hand tracking. *Brain and Cognition*, 48, 212–226.
- Miall, R. C., Weir, D. J., Wolpert, D. M., & Stein, J. F. (1993). Is the cerebellum a Smith predictor? *Journal of Motor Behavior*, 25(3), 203–216.
- Miall, R. C., & Wolpert, D. M. (1996). Forward models for physiological motor control. *Neural Networks*, 9(8), 1265–1279.
- Middleton, F. A., & Strick, P. L. (2000). Basal ganglia and cerebellar loops: motor and cognitive circuits. *Brain Research Reviews*, 31, 236–250.
- Miyamura, A., & Kimura, H. (2002). Stability of feedback error learning scheme. *Systems and Control Letters*, 45, 303–316.
- Montgomery, J. C., Bodznick, D., & Yopak, K. E. (2012). The cerebellum and cerebellum-like structures of cartilaginous fishes. *Brain Behaviour and Evolution*, 80(2), 152–165. <http://dx.doi.org/10.1159/000339868>.
- Munuera, J., Morel, P., Duhamel, J. R., & Deneve, S. (2009). Optimal sensorimotor control in eye movement sequences. *Journal of Neuroscience*, 29(10), 3026–3035. <http://dx.doi.org/10.1523/jneurosci.1169-08.2009>.
- Nakanishi, J., & Schaal, S. (2004). Feedback error learning and nonlinear adaptive control. *Neural Networks*, 17(10), 1453–1465.
- Ohtsuki, G., Piochon, C., & Hansel, C. (2009). Climbing fiber signaling and cerebellar gain control. *Frontiers in Cellular Neuroscience*, 3. <http://dx.doi.org/10.3389/neuro.03.004.2009>. Article 4.
- Paulin, M. G. (2005). Evolution of the cerebellum as a neuronal machine for Bayesian state estimation. *Journal of Neural Engineering*, 2, S219–S234.
- Porrill, J., & Dean, P. (2007a). Cerebellar motor learning: when is cortical plasticity not enough? *PLoS Computational Biology*, 3(10), 1935–1950.
- Porrill, J., & Dean, P. (2007b). Recurrent cerebellar loops simplify adaptive control of redundant and nonlinear motor systems. *Neural Computation*, 19(1), 170–193.
- Porrill, J., Dean, P., & Stone, J. V. (2004). Recurrent cerebellar architecture solves the motor error problem. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 271, 789–796.
- Ramrani, N., & Miall, C. (2001). Expanding cerebellar horizons. *Trends in Cognitive Sciences*, 5(4), 135–136.
- Requarth, T., & Sawtell, N. B. (2011). Neural mechanisms for filtering self-generated sensory signals in cerebellum-like circuits. *Current Opinion in Neurobiology*, 21(4), 602–608. <http://dx.doi.org/10.1016/j.conb.2011.05.031>.
- Roberts, P. D. (1999). Computational consequences of temporally asymmetric learning rules: I. Differential Hebbian learning. *Journal of Computational Neuroscience*, 7(3), 235–246.
- Roberts, P. D., & Bell, C. C. (2000). Computational consequences of temporally asymmetric learning rules: II. Sensory image cancellation. *Journal of Computational Neuroscience*, 9(1), 67–83.
- Robinson, D. A. (1981). The use of control systems analysis in the neurophysiology of eye movements. *Annual Review of Neuroscience*, 4, 463–503.
- Sabourin, C., Bruneau, O., & Buche, G. (2006). Control strategy for the robust dynamic walk of a biped robot. *International Journal of Robotics Research*, 25(9), 843–860. <http://dx.doi.org/10.1177/0278364906069151>.
- Sastry, S., & Bodson, M. (1989). *Adaptive control: stability, convergence, and robustness*. Englewood Cliffs, NJ: Prentice-Hall International.
- Sawtell, N. B., & Williams, A. (2008). Transformations of electrosensory encoding associated with an adaptive filter. *Journal of Neuroscience*, 28(7), 1598–1612.
- Schenck, W. (2011). Kinematic motor learning. *Connection Science*, 23(4), 239–283. <http://dx.doi.org/10.1080/09540091.2011.625077>.
- Schutte, S., van den Bedem, S. P., van Keulen, F., van der Helm, F. C., & Simonsz, H. J. (2006). A finite-element analysis model of orbital biomechanics. *Vision Research*, 46(11), 1724–1731.
- Schweighofer, N., Spoelstra, J., Arbib, M. A., & Kawato, M. (1998). Role of the cerebellum in reaching movements in humans. II. A neural model of the intermediate cerebellum. *European Journal of Neuroscience*, 10(1), 95–105.
- Scudder, C. A., & Fuchs, A. F. (1992). Physiological and behavioral identification of vestibular nucleus neurons mediating the horizontal vestibuloocular reflex in trained monkeys. *Journal of Neurophysiology*, 68, 244–264.
- Sejnowski, T. J. (1977). Storing covariance with nonlinearly interacting neurons. *Journal of Mathematical Biology*, 4, 303–321.

- Shmuelof, L., & Krakauer, J. W. (2011). Are we ready for a natural history of motor learning? *Neuron*, *72*(3), 469–476. <http://dx.doi.org/10.1016/j.neuron.2011.10.017>.
- Skavenski, A. A., & Robinson, D. A. (1973). Role of abducens neurons in vestibuloocular reflex. *Journal of Neurophysiology*, *36*, 724–738.
- Tao, T., & Su, S. F. (2011). CMAC-based previous step supervisory control schemes for relaxing bound in adaptive fuzzy control. *Applied Soft Computing*, *11*(8), 5715–5723. <http://dx.doi.org/10.1016/j.asoc.2011.03.019>.
- Todorov, E. (2005). Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system. *Neural Computation*, *17*(5), 1084–1108. <http://dx.doi.org/10.1162/0899766053491887>.
- Tolu, S., Vanegas, M., Luque, N. R., Garrido, J. A., & Ros, E. (2012). Bio-inspired adaptive feedback error learning architecture for motor control. *Biological Cybernetics*, *106*(8–9), 507–522. <http://dx.doi.org/10.1007/s00422-012-0515-5>.
- Voogd, J. (2011). Cerebellar zones: a personal history. *Cerebellum*, *10*, 334–350. <http://dx.doi.org/10.1007/s12311-010-0221-6>.
- Wang, S. S.-H., Denk, W., & Häusser, M. (2000). Coincidence detection in single dendritic spines mediated by calcium release. *Nature Neuroscience*, *3*(12), 1266–1273.
- Widrow, B., & Stearns, S. D. (1985). *Adaptive signal processing*. Engelwood Cliffs, NJ: Prentice-Hall Inc.
- Wolpert, D. M. (1997). Computational approaches to motor control. *Trends in Cognitive Sciences*, *1*(6), 209–216.
- Wolpert, D. M., & Ghahramani, Z. (2000). Computational principles of movement neuroscience. *Nature Neuroscience*, *3*(Suppl.), 1212–1217.
- Wolpert, D. M., Ghahramani, Z., & Jordan, M. I. (1995). An internal model for sensorimotor integration. *Science*, *269*, 1880–1882.
- Wolpert, D. M., & Kawato, M. (1998). Multiple paired forward and inverse models for motor control. *Neural Networks*, *11*(7–8), 1317–1329. doi: S0893-6080(98)00066-5. [pii].
- Wolpert, D. M., Miall, R. C., & Kawato, M. (1998). Internal models in the cerebellum. *Trends in Cognitive Science*, *2*(9), 338–347.