

Implications of different classes of sensorimotor disturbance for cerebellar-based motor learning models

Adrian Haith · Sethu Vijayakumar

Received: 13 November 2007 / Accepted: 15 September 2008 / Published online: 22 October 2008
© Springer-Verlag 2008

Abstract The exact role of the cerebellum in motor control and learning is not yet fully understood. The structure, connectivity and plasticity *within* cerebellar cortex has been extensively studied, but the patterns of connectivity and interaction with other brain structures, and the computational significance of these patterns, is less well known and a matter of debate. Two contrasting models of the role of the cerebellum in motor adaptation have previously been proposed. Most commonly, the cerebellum is employed in a purely feedforward pathway, with its output contributing directly to the outgoing motor command. The cerebellum must then learn an inverse model of the motor apparatus in order to achieve accurate control. More recently, Porrill et al. (Proc Biol Sci 271(1541):789–796, 2004) and Porrill et al. (PLoS Comput Biol 3:1935–1950, 2007a) and Porrill et al. (Neural Comput 19(1), 170–193, 2007b) have highlighted the potential importance of these recurrent connections by proposing an alternative architecture in which the cerebellum is embedded in a recurrent loop with brainstem control circuitry. In this framework, the feedforward connections are not necessary at all. The cerebellum must learn a forward model of the motor apparatus for accurate motor commands to be generated. We show here how these two models exhibit contrasting yet complimentary learning capabilities. Central to the differences in performance between architectures is that there are two distinct kinds of disturbance to which a motor system may need to adapt (1) changes in the relationship between the motor command and the observed outcome and (2) changes in the relationship between the stimulus and the desired outcome. The computational distinction between these two kinds of transformation is subtle and has there-

fore often been overlooked. However, the implications for learning turn out to be significant: learning with a feedforward architecture is robust following changes in the stimulus-desired outcome mapping but not necessarily the motor command-outcome mapping, while learning with a recurrent architecture is robust under changes in the motor command-outcome mapping but not necessarily the stimulus-desired outcome mapping. We first analyse these differences theoretically and through simulations in the vestibulo-ocular reflex (VOR), then illustrate how these same concepts apply more generally with a model of reaching movements.

Keywords Cerebellum · Motor adaptation · VOR · Kinematics

1 Introduction

Humans and other biological systems demonstrate remarkable abilities to adapt their motor behaviour to novel circumstances and to learn to perform new motor tasks. The need to adapt existing controllers may arise for a variety of reasons including growth, ageing, injury, disease or experimental intervention in the laboratory. Acquisition of accurate motor behaviour from birth during infancy also requires similar or perhaps even identical learning mechanisms to motor learning during adulthood.

In general, a motor control task involves generating appropriate motor commands in response to some stimulus to bring about a desired outcome. There are two fundamental types of change which can alter what the appropriate motor commands are in response to a given stimulus (see Fig. 1). Firstly, the relationship between the motor commands and the resulting outcome can be altered. This typically involves changes in the motor plant dynamics (e.g. through injury,

A. Haith (✉) · S. Vijayakumar
School of Informatics, University of Edinburgh,
Edinburgh EH8 9AB, UK
e-mail: adrian.haith@ed.ac.uk

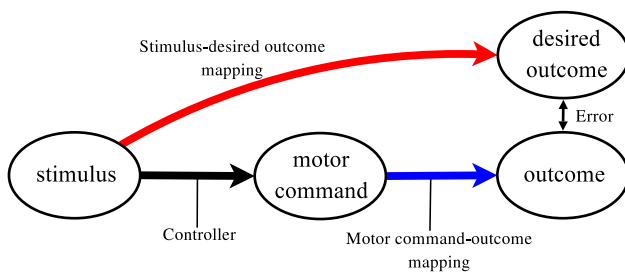


Fig. 1 Illustration of the two kinds of mappings described in Sect. 1

disease, growth or ageing); however some changes which are more kinematic in nature, such as distortions of visual feedback, can also be grouped into this category.

A second kind of change is in the relationship between the initial stimulus and the desired outcome following that stimulus. This kind of change is more subtle than the previous one but examples do occur in the context of most motor behaviours—either naturally or under experimental conditions. After such a change, the original response to the stimulus will no longer be appropriate and a new pattern of responses must be learnt.

In many circumstances, the stimulus and the desired outcome can be considered to be equivalent. For example, in the case of reaching, the stimulus is the location of an object in the visual field and the desired outcome is that the hand be in that same location (note that this is true even when visual feedback is tampered with). Nevertheless, there are numerous examples where the stimulus-desired outcome relationship is not so trivial and subject to change. As we will describe in later sections, many common experimental paradigms in oculomotor adaptation actually fall into the latter category rather than the former. Adaptation to these kinds of changes can also be induced in reaching tasks (Lurito et al. 1991) where they are sometimes referred to as ‘non-standard mappings’ or ‘transformational mappings’ (Shadmehr and Wise 2005).

The possibility of having to adapt in the face of an unknown relationship between the stimulus and the desired outcome has been noted before (Jordan and Rumelhart 1992), but a thorough examination of the extent to which this applies to human motor system in practice has been previously lacking: Indeed, the solutions proposed in (Jordan and Rumelhart 1992) do not aim at biological plausibility since they rely on backpropagation of error signals through a learnt internal model.

We examine the problem of adapting to both kinds of change from a biologically plausible cerebellar learning perspective. In the next section, we describe in detail two existing models of cerebellar-based motor learning in the context of a variety of commonly studied motor behaviours and examine their suitability for adapting to each of the two kinds of change described above.

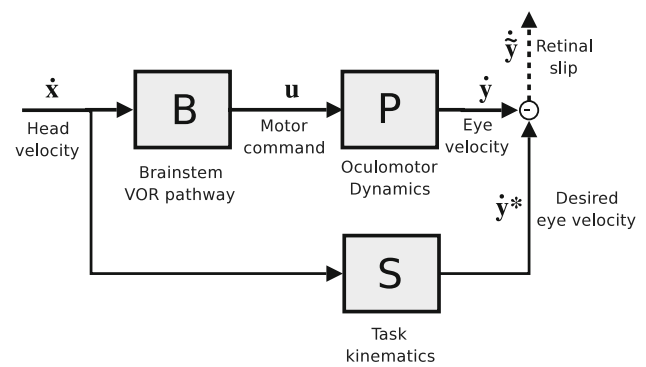


Fig. 2 Schematic of the VOR

Before going into the details of the different architectures, we will first illustrate the discussion with a concrete example in which both kinds of sensorimotor change occur naturally—the vestibulo-ocular reflex (VOR).

1.1 Kinematics and dynamics of the VOR

The vestibulo-ocular reflex in mammals acts to stabilize gaze during head rotations by counter-rotating the eyes. The characteristics of this reflex are not fixed, but can be modified through experience when conditions change. After a suitable amount of training under these new conditions, the VOR becomes recalibrated so that even in the dark, the response of the VOR is altered (Boyden et al. 2004).

For illustrative purposes, and since the oculomotor plant is known to have negligible inertia (Robinson 1964), we consider a first-order dynamics model of the oculomotor plant. Denoting the current eye position by \mathbf{y} , the relationship between the motor command \mathbf{u} and the resulting eye velocity $\dot{\mathbf{y}}$ (see Fig. 2) is determined by the forward dynamics of the oculomotor plant,

$$\dot{\mathbf{y}} = P(\mathbf{y}, \mathbf{u}). \quad (1)$$

The inverse dynamics model is correspondingly defined as

$$\mathbf{u} = P^{-1}(\mathbf{y}, \dot{\mathbf{y}}). \quad (2)$$

The inverse dynamics P^{-1} map the current eye position \mathbf{y} and eye velocity $\dot{\mathbf{y}}$ to a motor command \mathbf{u} which would achieve that eye velocity when acting through the plant.

The desired outcome in this case is that the gaze be stabilized, i.e. that the eye velocity $\dot{\mathbf{y}}$ be equal to some gaze-stabilizing eye velocity $\dot{\mathbf{y}}^*$. Any deviations of eye velocity from this desired value will be perceived as retinal slip—movement of the visual image across the retina. We denote this retinal slip by $\tilde{\dot{\mathbf{y}}}$ and it is given by

$$\tilde{\dot{\mathbf{y}}} = \dot{\mathbf{y}}^* - \dot{\mathbf{y}}. \quad (3)$$

We define \mathbf{u}^* as the motor command which achieves the desired eye velocity $\dot{\mathbf{y}}^*$ when supplied to the plant,

$$\mathbf{u}^* = P^{-1}(\mathbf{y}, \dot{\mathbf{y}}^*). \tag{4}$$

In most VOR models, desired eye velocity is taken as equal and opposite to head velocity, i.e. $\dot{\mathbf{y}}^* = -\dot{\mathbf{x}}$. However, in general, this is not the case. Most VOR gain adaptation experiments work by directly manipulating the relationship between head velocity and desired eye velocity, not by changing the properties of the oculomotor plant. This includes any experiment using prisms or lenses or vestibular mismatch experiments in which an external visual stimulus is moved concurrently with head movements. If the stimulus is moved in phase with head movements, this achieves a change in the gain of the required response.

VOR adaptation therefore cannot be regarded as simply a process of learning the motor command-outcome mapping of the oculomotor plant. Fundamental to VOR adaptation is that the kinematic relationship between the stimulus (vestibular signal, indicating head velocity) and the desired outcome (gaze-stabilizing eye velocity) is also subject to change and must be compensated for, as illustrated in Fig. 2. This kinematic component is effectively what many experimental paradigms manipulate, such as when the visual image is transformed before reaching the eye or in visual/vestibular mismatch conditions. It also encompasses more natural disturbances such as off-axis effects (Coenen and Sejnowski 1996) and inaccuracies or nonlinearities in the relationship between the vestibular signal and the true head velocity $\dot{\mathbf{x}}$. VOR adaptation must, either explicitly or implicitly, reflect learning of this kinematic relationship as well as of the plant dynamics.

Mathematically, we can describe this relationship as a function mapping head velocity $\dot{\mathbf{x}}$ and current eye position \mathbf{y} to a desired eye velocity $\dot{\mathbf{y}}^*$,

$$\dot{\mathbf{y}}^* = S(\mathbf{y}, \dot{\mathbf{x}}) \tag{5}$$

and a corresponding inverse kinematics mapping

$$\dot{\mathbf{x}}' = S^{-1}(\mathbf{y}, \dot{\mathbf{y}}). \tag{6}$$

S maps the current head velocity $\dot{\mathbf{x}}$ to an appropriate gaze-stabilizing eye velocity $\dot{\mathbf{y}}$, while S^{-1} gives the head rotation $\dot{\mathbf{x}}'$ that would have required an eye movement $\dot{\mathbf{y}}$ to stabilize gaze.

Hence, from (2) and (5), the overall mapping which must be learnt is a composite function

$$\mathbf{u}^* = P^{-1}(\mathbf{y}, S(\mathbf{y}, \dot{\mathbf{x}})). \tag{7}$$

This shows exactly how each kind of mapping, stimulus-desired outcome or motor command-outcome, influences the choice of motor command \mathbf{u} .

In later sections, we will give further examples of behaviours where both of these kinds of mappings are subject

to change. First, however, we show the implications that changes in each of these kinds of mappings has for different models of cerebellar-based adaptation.

2 Cerebellar-based learning models

We now analyse in more detail the alternative architectures for cerebellar connectivity with other brain regions and the implications for learning. In this section, we describe a model of cerebellar-based VOR adaptation. The VOR is used here as an illustrative example, particularly because the underlying neuroanatomy has been extensively studied. The arguments we make in this section, however, are quite general to all error-driven cerebellar learning scenarios and central to the theme of this paper.

It is well known that the cerebellum plays a crucial role in VOR adaptation. Although a variety of brainstem circuitry contributes to the VOR, the basic reflex is largely mediated by a fast three-synapse pathway in the brainstem (Boyden et al. 2004). This brainstem pathway is augmented by an adaptive pathway through the cerebellum. It is known that among its many inputs, the region of the cerebellum involved in the VOR receives a variety of parallel-fibre inputs including vestibular-related signals and efferent copies of outgoing oculomotor commands (Hirata and Highstein 2001).

Two distinct kinds of simplified models have previously been proposed to approximate this circuitry. Most models have tended to disregard the motor command efferent copies, modelling the brainstem and cerebellar pathways as having a purely feedforward organization (Gomi and Kawato 1990; Kawato and Gomi 1992; Shibata and Schaal 2001). This kind of feedforward architecture is illustrated in Fig. 3a. More recently, Porrill et al. (2004) have argued that the efferent copy inputs should not be ignored and, furthermore, efferent copy information *alone* is sufficient for successful VOR adaptation. They demonstrate this with a model of cerebellar VOR adaptation which has a *purely* recurrent architecture i.e., having no feedforward vestibular input to the cerebellum. This architecture is illustrated in Fig. 3b.

These two alternative models of VOR adaptation each effectively assume that either the feedforward or the recurrent (efferent copy) inputs to the cerebellum dominate, with the other inputs playing a more minor role in generating the cerebellar output. The true contributions from each of these inputs to the cerebellar output is unknown. Although both architectures are quite capable of accurate control, the learning properties of the feedforward and recurrent architectures turn out to be quite different in adapting to different kinds of disturbance.

Before analysing the different architectures in detail, we briefly discuss a basic model of learning within cerebellar cortex.

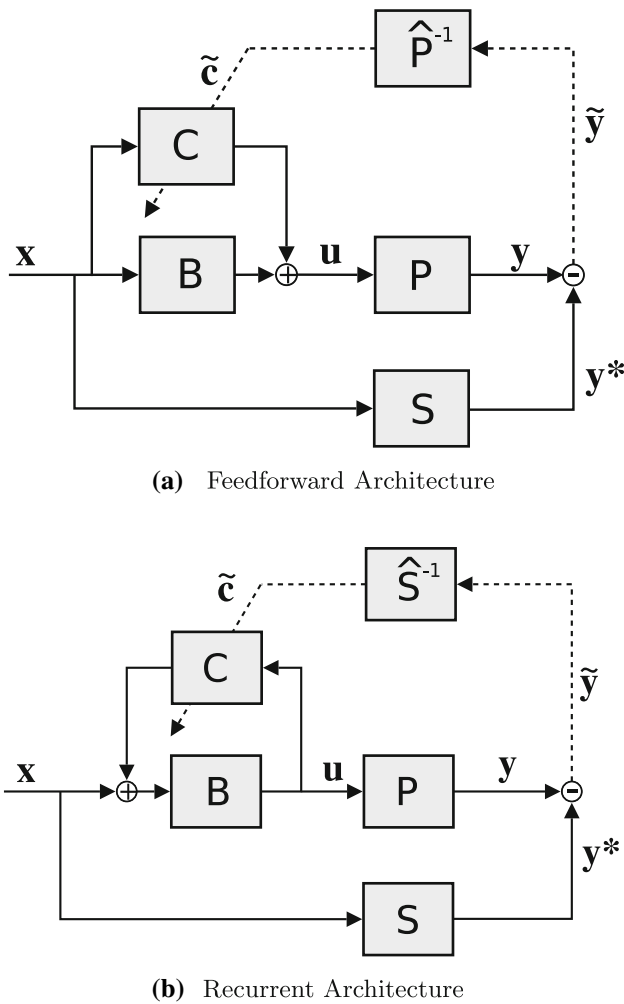


Fig. 3 Schematic of feedforward and recurrent architectures

2.1 Cerebellar plasticity model

For the purpose of illustration we will assume a simple Marr-Albus-Ito type model of cerebellar learning (Ito 2000). The arguments presented here, however, apply to any iterative error-driven learning scheme, including CMACs and more abstract, machine-learning inspired cerebellum models (e.g. Shibata and Schaal 2001).

Let us denote the input to the cerebellum by $\mathbf{z}(t)$ and its output by $\mathbf{c}(t)$. In this model, the cerebellar output $\mathbf{c}(t) = C(\mathbf{z}(t))$ is given by a weighted sum of parallel fibre activities:

$$c_j(t) = \sum_i w_{ij} p_i(t),$$

where, $p_i(t)$ is the activity (i.e. firing rate) of the i th parallel fibre and w_i is the strength of the corresponding parallel fibre-Purkinje cell synapse. This can be written more compactly in vector notation as

$$\mathbf{c}(t) = \mathbf{w}^T \mathbf{p}(t). \tag{8}$$

Learning occurs through adaptation of the synaptic weights \mathbf{w} over time. This plasticity is driven by the climbing fibre signal $\tilde{\mathbf{c}}(t)$. It is widely believed that the climbing fibre signals directly represent errors in motor performance. Some studies have claimed that this corresponds to sensory error (Simpson et al. 1996), while others have suggested that it is a motor-based representation (Wolpert et al. 1998; Yutaka et al. 2005). We assume that the climbing fibre signal corresponds to the error in the cerebellar output $\mathbf{c}(t)$. Whether this equates to a sensory or a motor error will become clear further down. A simple learning rule which approximately captures the known plasticity laws at the synapses is

$$\dot{w}_{ij} = -\beta \tilde{c}_j(t) p_i(t), \tag{9}$$

or in vector notation

$$\dot{\mathbf{w}} = -\beta \tilde{\mathbf{c}}(t) \mathbf{p}(t)^T. \tag{10}$$

Although firing rates clearly cannot be negative, $\mathbf{p}(t)$ and $\tilde{\mathbf{c}}(t)$ can be interpreted as deviations from some baseline firing rate. This learning rule is equivalent to performing gradient descent on the squared error, provided $\tilde{\mathbf{c}}(t)$ really does reflect the error in the cerebellar output.

We do not, however, have explicit knowledge of the error in the cerebellar output. We can only measure the error in performance in terms of retinal slip. This poses a distal learning problem which can be solved by finding an appropriate mapping between the observed outcome error (retinal slip \tilde{y} in the case of the VOR) and the error in the cerebellar output $\tilde{\mathbf{c}}$.

In order to calculate this, we need to know exactly what the desired output of the cerebellum is. This turns out to depend strongly on the architecture (feedforward or recurrent). Previous analyses of cerebellar VOR adaptation have tended to consider only changes in the motor command-outcome mapping (i.e. they have neglected the fact that it is often the stimulus-desired outcome relationship which changes in an experiment). In this case (i.e. if we assume head velocity and desired eye velocity are equal and opposite), the error in the cerebellar output $\tilde{\mathbf{c}}$ under the feedforward architecture is equal to the error in the motor command,

$$\tilde{\mathbf{c}} = \mathbf{u}^* - \mathbf{u}. \tag{11}$$

This motor error can be estimated from the retinal slip by transforming it through an inverse model of the plant dynamics.

Under the recurrent architecture, on the other hand, Porrill et al. (2004) have shown that the error in the cerebellar output is equal to the raw retinal slip signal \tilde{y} and thus the distal learning problem is circumvented. Their analysis of the properties of the recurrent architecture was restricted to adaptation to changes in the oculomotor plant dynamics (although they do acknowledge that other kinds of disturbances are possible).

We will now examine learning within each of these architectures in the case that the head velocity-eye movement relationship S is also subject to change. In both cases, we require an expression for the error in the cerebellar output, $\tilde{\mathbf{c}}$, in terms of the observed output error (retinal slip, $\hat{\mathbf{y}}$, in the case of the VOR).

2.2 Learning in the feedforward architecture

Let us refer to the feedforward architecture shown in Fig. 3a. The inputs to the cerebellum in this model are head velocity $\dot{\mathbf{x}}$ and head position \mathbf{x} ($\dot{\mathbf{x}}$ is omitted from the figure for clarity). We also assume that an optimal cerebellar model C^* exists (this corresponds to a set of optimal weights \mathbf{w}^* for the cerebellum model outlined in the previous section). The error in the cerebellar output is then defined as

$$\tilde{\mathbf{c}}(\mathbf{x}, \dot{\mathbf{x}}) = C^*(\mathbf{x}, \dot{\mathbf{x}}) - C(\mathbf{x}, \dot{\mathbf{x}}), \tag{12}$$

and we wish to express this in terms of the retinal slip $\hat{\mathbf{y}}$.

The motor command is generated by combining the output from the brainstem and cerebellum

$$\mathbf{u} = C(\mathbf{x}, \dot{\mathbf{x}}) + B(\mathbf{x}, \dot{\mathbf{x}}), \tag{13}$$

where $B(\mathbf{x}, \dot{\mathbf{x}})$ describes the brainstem dynamics. Similarly for the optimal cerebellum model

$$\mathbf{u}^* = C^*(\mathbf{x}, \dot{\mathbf{x}}) + B(\mathbf{x}, \dot{\mathbf{x}}). \tag{14}$$

Noting that the optimal motor command \mathbf{u}^* is given by $P^{-1}(\mathbf{y}, S(\mathbf{y}, \dot{\mathbf{x}}))$, we can see that the optimal cerebellum model C^* satisfies

$$C^*(\mathbf{x}, \dot{\mathbf{x}}) = P^{-1}(\mathbf{y}, S(\mathbf{y}, \dot{\mathbf{x}})) - B(\mathbf{x}, \dot{\mathbf{x}}). \tag{15}$$

The cerebellum must therefore learn a composite of a forward kinematics model and an inverse dynamics model, while compensating for the contribution from the brainstem B .

Now, taking the difference between (13) and (14) and comparing it to (12) illustrates that we can express the error in the cerebellar output as:

$$\tilde{\mathbf{c}} = \mathbf{u}^* - \mathbf{u}, \tag{16}$$

that is, $\tilde{\mathbf{c}}$ is equal to the *motor error*. Rewriting the right hand side of (16) in terms of the inverse plant model (2), we have

$$\tilde{\mathbf{c}} = P^{-1}(\mathbf{y}, \hat{\mathbf{y}}^*) - P^{-1}(\mathbf{y}, \hat{\mathbf{y}}). \tag{17}$$

For linear plant dynamics, we can directly simplify and rewrite the expression in terms of the retinal slip $\hat{\mathbf{y}}$,

$$\tilde{\mathbf{c}} = P^{-1}(\mathbf{y}, \hat{\mathbf{y}}). \tag{18}$$

For nonlinear plant dynamics, (17) can be approximated by the first term of the Taylor expansion of P^{-1} about $(\mathbf{y}, \hat{\mathbf{y}})$:

$$\tilde{\mathbf{c}} \approx J_{P^{-1}}(\mathbf{y}, \hat{\mathbf{y}})\tilde{\mathbf{y}}, \tag{19}$$

where $J_{P^{-1}}(\mathbf{y}, \hat{\mathbf{y}})$ is the Jacobian of P^{-1} at the point $(\mathbf{y}, \hat{\mathbf{y}})$.

Equations (18) and (19) show that the error in the cerebellar output can be calculated from the retinal slip via the inverse dynamics of the plant, i.e. the inverse of the mapping from motor commands to observed outcome. We assume that some internal model is available to compute this, which may reside elsewhere in the cerebellum (Wolpert et al. 1998). Alternatively, the error signal might be the direct output of a feedback controller, as in feedback-error learning (Gomi and Kawato 1990). In any case, if the plant dynamics change, an unadapted internal model will still reflect the old dynamics and we can no longer be confident that our estimate of the cerebellar output error is accurate.

The required training signal is, however, independent of the kinematics S , i.e. the relationship between the stimulus (head velocity $\dot{\mathbf{x}}$) and the desired outcome (eye velocity $\dot{\mathbf{y}}$). This is an important but usually overlooked advantage of employing a feedforward cerebellar architecture.

So, in general, we expect learning under the feedforward architecture to be impaired (i.e. converge more slowly) or even made entirely unstable (not converge at all) following a change in the motor command-outcome (dynamics) mapping. However, we expect learning to be unaffected by a change in the kinematics.

2.3 Learning in the recurrent architecture

Next, we derive an expression for the error in the cerebellar output in terms of the measured retinal slip for the recurrent architecture (Fig. 3b). We assume that the inputs to the cerebellum are the head position \mathbf{x} and the afferent motor command \mathbf{u} (the head-position input is omitted in Fig. 3b for clarity).

We begin the derivation by noting that the input to the brainstem model is given by $C(\mathbf{x}, \mathbf{u}) + \dot{\mathbf{x}}$, which is equal to the motor command transformed under the brainstem inverse model, i.e.

$$C(\mathbf{x}, \mathbf{u}) + \dot{\mathbf{x}} = B^{-1}(\mathbf{x}, \mathbf{u}). \tag{20}$$

Again, as in the feedforward case, we assume there exists an optimal cerebellar model C^* which yields exactly the desired motor command (this corresponds to optimal weights \mathbf{w}^* in the cerebellum model outlined previously).

Note that the motor command \mathbf{u} would be optimal for some alternate head velocity $\dot{\mathbf{x}}'$, i.e.

$$C^*(\mathbf{x}, \mathbf{u}) + \dot{\mathbf{x}}' = B^{-1}(\mathbf{x}, \mathbf{u}), \tag{21}$$

with $\dot{\mathbf{x}}' = S^{-1}(\mathbf{y}, P(\mathbf{y}, \mathbf{u}))$, by definition of S^{-1} from (1) to (6). Rearranging this, we have

$$C^*(\mathbf{x}, \mathbf{u}) = B^{-1}(\mathbf{x}, \mathbf{u}) - S^{-1}(\mathbf{y}, P(\mathbf{y}, \mathbf{u})). \tag{22}$$

Under the recurrent architecture then, the cerebellum must learn a composite of a forward dynamics model and an

inverse kinematics model. This is in direct contrast to the feedforward case in (15).

Taking the difference between (20) and (21), we can express the cerebellar output error as:

$$C^*(\mathbf{x}, \mathbf{u}) - C(\mathbf{x}, \mathbf{u}) = \dot{\mathbf{x}}' - \dot{\mathbf{x}} \\ = S^{-1}(\mathbf{y}, \dot{\mathbf{y}}) - S^{-1}(\mathbf{y}, \dot{\mathbf{y}}^*). \quad (23)$$

If we assume that S is linear then we can express this simply in terms of the retinal slip,

$$\tilde{\mathbf{c}}(\mathbf{x}, \mathbf{u}) = S^{-1}(\mathbf{y}, \dot{\mathbf{y}}). \quad (24)$$

If S is nonlinear, a first order Taylor approximation can be used,

$$\tilde{\mathbf{c}} \approx J_{S^{-1}}(\mathbf{y}, \dot{\mathbf{y}})\dot{\mathbf{y}}, \quad (25)$$

where $J_{S^{-1}}(\mathbf{y}, \dot{\mathbf{y}})$ is the Jacobian of S^{-1} at the point $(\mathbf{y}, \dot{\mathbf{y}})$.

Equations (24) and (25) show that the error in the cerebellar output is given by the retinal slip transformed via the inverse kinematics, i.e. the inverse of the mapping from the stimulus to the desired outcome. This can be thought of as an error in the original vestibular signal $\dot{\mathbf{x}}$.

So in general, we expect learning under the recurrent architecture to be impaired under changes in the stimulus-desired outcome mapping, but to be unaffected by changes in the motor command-outcome mapping.

This reveals a duality between the feedforward and recurrent architecture models. The properties of learning in the feedforward architecture are mirrored by those of learning in the recurrent architecture with the roles of the two kinds of transformation transposed.

2.4 Summary

Although we have illustrated the argument with the specific example of the VOR, the arguments presented here are entirely general and can be applied to any other motor behaviour. Furthermore, the only point at which linearity was assumed was in the final step in each derivation combining the terms in Eqs. (17) and (23). For nonlinear P and S , a Taylor expansion gives a simple approximation to the cerebellar output error in terms of the observed output error, provided the error is not too large.

Under the feedforward architecture model, in order for the observed outcome error to act as an appropriate training signal for the cerebellum, it must be transformed into the motor domain via the motor command-outcome mapping. If this mapping changes drastically, for instance, through a change of the plant dynamics, the error signal being used may no longer accurately reflect the error in the cerebellar output and learning may proceed less efficiently or even be disrupted entirely. The transformed error is, however, independent of the head velocity-desired eye velocity mapping. Thus

the feedforward architecture is guaranteed to learn stably following changes in this mapping.

For the recurrent architecture model, in order to train the cerebellum, the observed error must be transformed via the inverse of the stimulus-desired outcome mapping, giving something analogous to an error in the initial stimulus. If the true mapping from stimulus to desired outcome changes substantially, learning may be slowed down or even disrupted entirely. The transformation of the observed error into cerebellar output error is, however, independent of the motor command-outcome relationship and thus, the recurrent architecture is guaranteed to learn stably following a change in that mapping.

3 Simulation of VOR adaptation

In order to test the performance of each of the two alternative VOR models in adapting to a range of changes in both the motor command-outcome mapping and the stimulus-desired outcome mapping, we simulated adaptation of a 2 degree-of-freedom oculomotor plant under a range of transformations of the kinematics and the plant dynamics. The simulated oculomotor plant had simplified dynamics initially given by

$$\dot{\mathbf{y}} = \mathbf{u}, \quad (26)$$

and an initial relationship between head velocity and gaze-stabilizing eye velocity given by

$$\dot{\mathbf{y}}^* = -\dot{\mathbf{x}}. \quad (27)$$

To simulate a change in the relationship between motor command and observed outcome, we changed the dynamics from the ordinary resistive viscosity field described by (26) to a viscous curl field (Fig. 4a) in which there is an angle ϕ between the eye velocity and the force, i.e.

$$\mathbf{u} = P_1 \dot{\mathbf{y}}, \quad (28)$$

where

$$P_1 = \begin{pmatrix} \cos \phi & \sin \phi \\ -\sin \phi & \cos \phi \end{pmatrix}. \quad (29)$$

Due to the first-order dynamics assumed here, this had the effect of rotating the angle of actuation for a given motor command.

To change the relationship between head velocity and gaze-stabilizing eye velocity, we employed a rotation of the visual field (Fig. 4b) by angle ψ . Following this transformation, the desired eye velocity is rotated by angle ψ relative to the head velocity, i.e.

$$\dot{\mathbf{y}}^* = S_1 \dot{\mathbf{x}}^*, \quad (30)$$

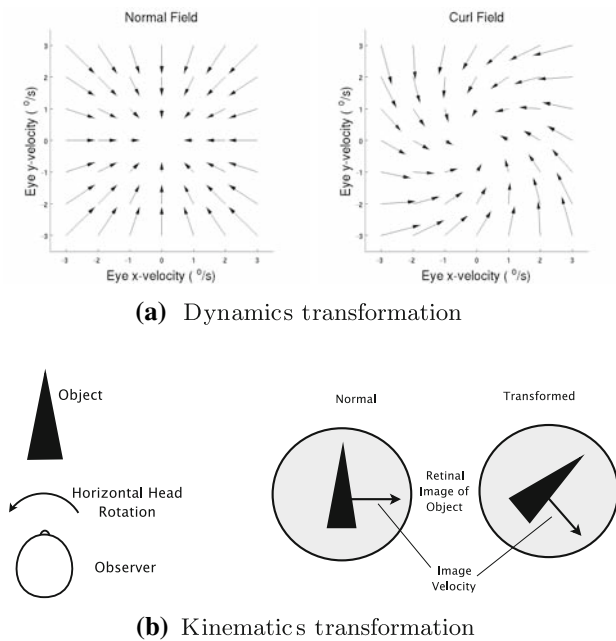


Fig. 4 Example of dynamic and kinematic transformations


where

$$S_1 = \begin{pmatrix} \cos \psi & \sin \psi \\ -\sin \psi & \cos \psi \end{pmatrix}. \tag{31}$$

In all experiments, the head position repeatedly traced out a figure-of-eight:

$$\mathbf{x}(t) = \begin{bmatrix} \sin(0.1t) \\ \sin(0.2t) \end{bmatrix}^T. \tag{32}$$

All experiments were run 10 times, with different initial positions around the figure-of-eight on each trial. Full implementation details are given in Appendix 7.



3.1 Experiments

First, we tested the performance of the feedforward architecture in adapting to the visuomotor rotation. Analyzing the normalised mean-squared (nMSE) velocity error / retinal slip, we found no significant difference in the learning trace when adaptation to different magnitudes of rotation ($\psi = [15, 135]$) were compared. Figure 5a plots the average nMSE over time for $\psi = 45^\circ$ which is representative of all values of ψ . The error bars represent one standard deviation above and below the mean.

We then tested the performance of the feedforward architecture in adapting to novel dynamics. Figure 5a plots the evolution of the nMSE over time for different values of ϕ . For $\phi = 15^\circ$, performance is the same as under the visuomotor rotation. As ϕ increases, however, the rate of improvement drops. Error bars are plotted for the $\phi = 60^\circ$ case to show that this difference is significant (error bars on other plots are omitted for clarity). At $\phi = 90^\circ$, the VOR no longer converges and updates of the cerebellar weights no longer improve performance. This is equivalent to always moving perpendicularly to the direction of steepest slope. For $\phi > 90^\circ$, changes in the cerebellar weights led to deteriorating performance and unstable adaptation.

For the recurrent architecture, we first tested the performance under the change in dynamics. For $\phi < 60^\circ$, we found no significant difference in performance between different values of ϕ . For larger values of ϕ , however, the recurrent loop tended to become unstable after a period of initial improvement. Figure 5b shows the nMSE over time for $\phi = 45^\circ$ which was representative of all trials for $\phi < 60^\circ$.

Finally, we tested the performance of the recurrent architecture in adapting to the visuomotor rotation. Results from these trials are also plotted in Fig. 5b. Again, for clarity, error bars are only plotted for representative transformations. For




Fig. 5 Timecourse of adaptation for VOR model using a feedforward architecture and b recurrent architecture. Both figures display average normalized mean squared eye velocity error (retinal slip), averaged over

ten trials. Different traces show response to different conditions—either a change in dynamics (viscous curl field) or a change in kinematics (visual field rotation) of differing magnitudes

$\psi = 15^\circ$, performance is similar to that under the change in the dynamics. For $\psi = 45^\circ$, however, the adaptation is significantly slower. For $\psi = 50^\circ$ and greater, the recurrent loop tended to become unstable resulting in an exponential increase of the error over a very short timescale. The plots have therefore been curtailed at this point. The initial rate of improvement in performance is nevertheless reflective of the quality of the estimate of the cerebellar output error.

For the linear models considered in this section, the problem of instability in the brainstem-cerebellum loop under the recurrent architecture can be attributed to the eigenvalues λ of the matrix BC (refer Fig. 3b) having magnitude $|\lambda| > 1$. It may, however, be possible to avoid entering into unstable regions of the parameter (cerebellar weight) space by also adapting B , using the output of C as a training signal. This ‘learning transfer’ from C to B would steer the loop away from regions of instability by ensuring that C^* (which would now depend on B) would tend asymptotically to 0. Learning transfer of this kind is supported by physiological evidence (Boyden et al. 2004) and Porrill and Dean (2007a) have suggested it may be used as a mechanism to enhance VOR response at high frequencies.

In summary, as predicted by the theory, performance of the feedforward architecture was impaired following changes in the oculomotor plant dynamics, but was not affected by changes in the kinematics. Performance of the recurrent architecture, on the other hand, was affected by changes in the kinematics but not by changes in the dynamics.

4 Reaching and catching

Arm movements are considerably more complex than eye movements, having more degrees of freedom and nonlinear dynamics. Nevertheless, many of the control principles observed in oculomotor control can be easily generalized to this setting, as we shall describe in this section.

While the exact role of the cerebellum in reaching movements is not entirely understood, cerebellar involvement has been demonstrated in compensating for interaction torques between limb segments (Bastian et al. 1996), adapting to altered dynamics (Smith and Shadmehr 2005), adapting to altered visual feedback (Baizer et al. 1999) and in learning to manipulate an on-screen cursor (Imamizu et al. 2000). The role of motor command generation in reaching is largely attributable to the primary motor cortex (Todorov 2000; Shadmehr and Wise 2005). It therefore seems reasonable to extend the cerebellar-based adaptation frameworks for the VOR from Sects. 2 to model the cerebellum’s role in reaching adaptation.

Indeed, Schweighofer et al. (1998) have suggested that the cerebellum generates feedforward motor commands which refine those generated by the cortex, effectively assuming

a feedforward cerebellar architecture. By contrast, Porrill and Dean (2007b) have proposed a model of arm control which employs a recurrent cerebellar architecture to learn an inverse kinematics model of the arm. Each of the architectures considered in the previous sections might therefore be considered plausible abstract models of the cerebellum’s role in the adaptive control of reaching.

The most common experimental paradigms for eliciting adaptation of reaching movements include application of an external force field to the hand via a robotic manipulandum (e.g., Shadmehr and Mussa-Ivaldi 1994), or tampering with visual feedback—usually in the form of a rotation of the visual field about the initial position of the hand (Krakauer et al. 2000). Both of these manipulations amount to changing the relationship between the motor command and the eventual observed outcome.

A few studies have, however, explored adaptation to changes in the relationship between stimulus and desired outcome. Magescas and Prablanc (2006) trained human subjects on a target-shift paradigm similar to that commonly used in saccade adaptation. A visual target was extinguished as subjects began the reach and reappeared in a new position slightly shifted from the where it had been initially. The final location of the target (i.e. the desired outcome) depends only on the initial stimulus and not on the motor command. Subjects were able to adapt to the task and exhibited robust aftereffects in subsequent reaches. Similar target-jump conditions have also been performed elsewhere (Lurito et al. 1991; Diedrichsen et al. 2005).

This kind of disturbance corresponds to a change in the stimulus-desired outcome mapping. The required motor command in each of these cases is very similar, however the nature of learning under the proposed cerebellar learning models is quite different and depends on the cerebellar architecture assumed.

4.1 Simulations

4.1.1 Reaching model

Reaching movements appear to be planned as a visually estimated difference vector \mathbf{dx} between current hand location and target location (Krakauer et al. 2000; Shadmehr and Wise 2005). A simplified model is illustrated in Fig. 6. Based on this difference vector, a suitable change in joint angles \mathbf{du} is selected by the fixed controller B . We maintain our label B for this controller from the previous model of the VOR where it denotes ‘brainstem’, although here it denotes primary motor cortex. In general, the ‘motor command’ \mathbf{du} will depend also on the current set of joint angles \mathbf{u} ; for simplicity, in our experiments, we consider a unique starting joint position \mathbf{u} . The final observed hand position is then given by the kinematics P which represents a mapping from

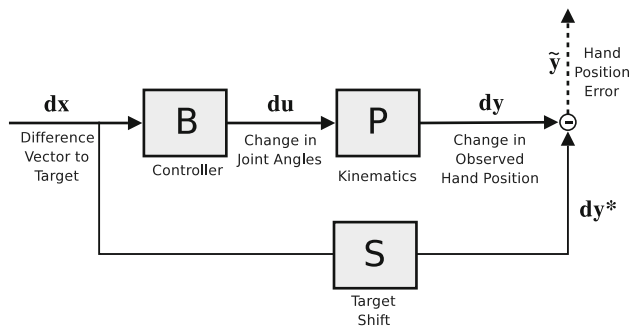


Fig. 6 Schematic of human reaching model

the change in joint angles \mathbf{du} to a change in observed hand position \mathbf{dy} ,

$$\mathbf{dy} = P(\mathbf{du}). \tag{33}$$

Note that this P includes both the forward kinematics of the arm, which map joint angles into hand position, and any distortions of visual feedback such as a rotation about the initial hand position.

Meanwhile, the target may be shifted mid-reach with the shift S determining the difference vector \mathbf{dy}^* of the shifted target (and therefore, the desired observed change in hand position),

$$\mathbf{dy}^* = S(\mathbf{dx}). \tag{34}$$

The error $\tilde{\mathbf{y}}$ in the observed hand position is used to guide adaptation of the controller and is given by the difference between the the actual and desired change in observed hand position,

$$\tilde{\mathbf{y}} = \mathbf{dy}^* - \mathbf{dy}. \tag{35}$$

Adaptation in our model is mediated by the cerebellar pathway which may be connected with the motor cortex B through either a feedforward or a recurrent architecture. We model the cerebellar output as being constructed from a weighted sum of radial basis functions defined over the input space. The combination weights are learnt using the same gradient-descent learning rule as in previous sections. The error in the cerebellar output was estimated by transforming the error in hand position appropriately according to the architecture used, as described in Sect. 2. For the feedforward architecture, this corresponded to transforming the output error through a linearization of the original inverse kinematics mapping (without any rotation of visual feedback taken into account). For the recurrent architecture, this simply corresponded to using the raw observed output error, since the initial stimulus-desired outcome mapping S is simply the identity. More technical details are given in Appendix 7.

4.1.2 Experiment

We set up a 10×10 square grid of targets around the initial position of the hand (given by the initial joint-angles θ_0 (see Fig. 7a). The goal was to find a suitable change of joint angles \mathbf{du} such that the resulting change in hand position brought the hand in line with the target.

Two separate conditions were chosen requiring adaptation, the first corresponding to a change in the stimulus-desired outcome mapping, the second a change in the motor command-outcome mapping.

In the first condition, we implemented a target-shifting paradigm along the lines of the experiment presented in (Lurito et al. 1991; Magescas and Prablanc 2006). During the reach, the position of each target was shifted by rotating it about the initial hand location by an angle $-\psi$.

In the second condition, we implemented a rotation of the visual field by angle ϕ about the initial hand location. This led to a rotation of both the stimulus \mathbf{dx} and the estimated error $\tilde{\mathbf{y}}$. To ensure that the set of stimuli used (i.e. visually estimated difference vectors \mathbf{dx}) was the same in both conditions, the grid of targets was rotated by angle $-\phi$ during the visual field rotation. The mapping to be learnt by the cerebellum C and the set of stimuli used were the same in either condition provided $\phi = \psi$. The only difference between conditions, therefore, was in the nature of the error signal and how it related to the error in the cerebellar output.

A sequence of 200 reaching movements was simulated to a random sequence of the 100 targets and this was repeated 100 times with different randomly selected target sequences. The same target sequences were used for each architecture and for each condition. Figure 7 illustrates the results of learning under the different architectures and across the two different conditions. In particular, (d) and (e) show the normalized mean squared global test error (nMSE), averaged over all targets and all sessions, as a function of the number of reach trials performed for visual and target rotations of 45° . For the feedforward architecture, it is particularly clear that performance is impaired under the visual rotation condition relative to the target rotation condition, in accordance with the theory. For the recurrent architecture, there was less difference in performance between conditions with marginally better asymptotic average performance under the visual rotation condition.

To demonstrate more clearly the effect that different kinds of transformation had on the quality of learning for different architectures, we examined the trend in reach errors during learning for a small subset of targets (marked with dots in the grid in Fig. 7a). We sampled the final hand position obtained during test trials to these targets after every 10 training trials. Because of the strong effect of the order of training targets on learning, we averaged these positions

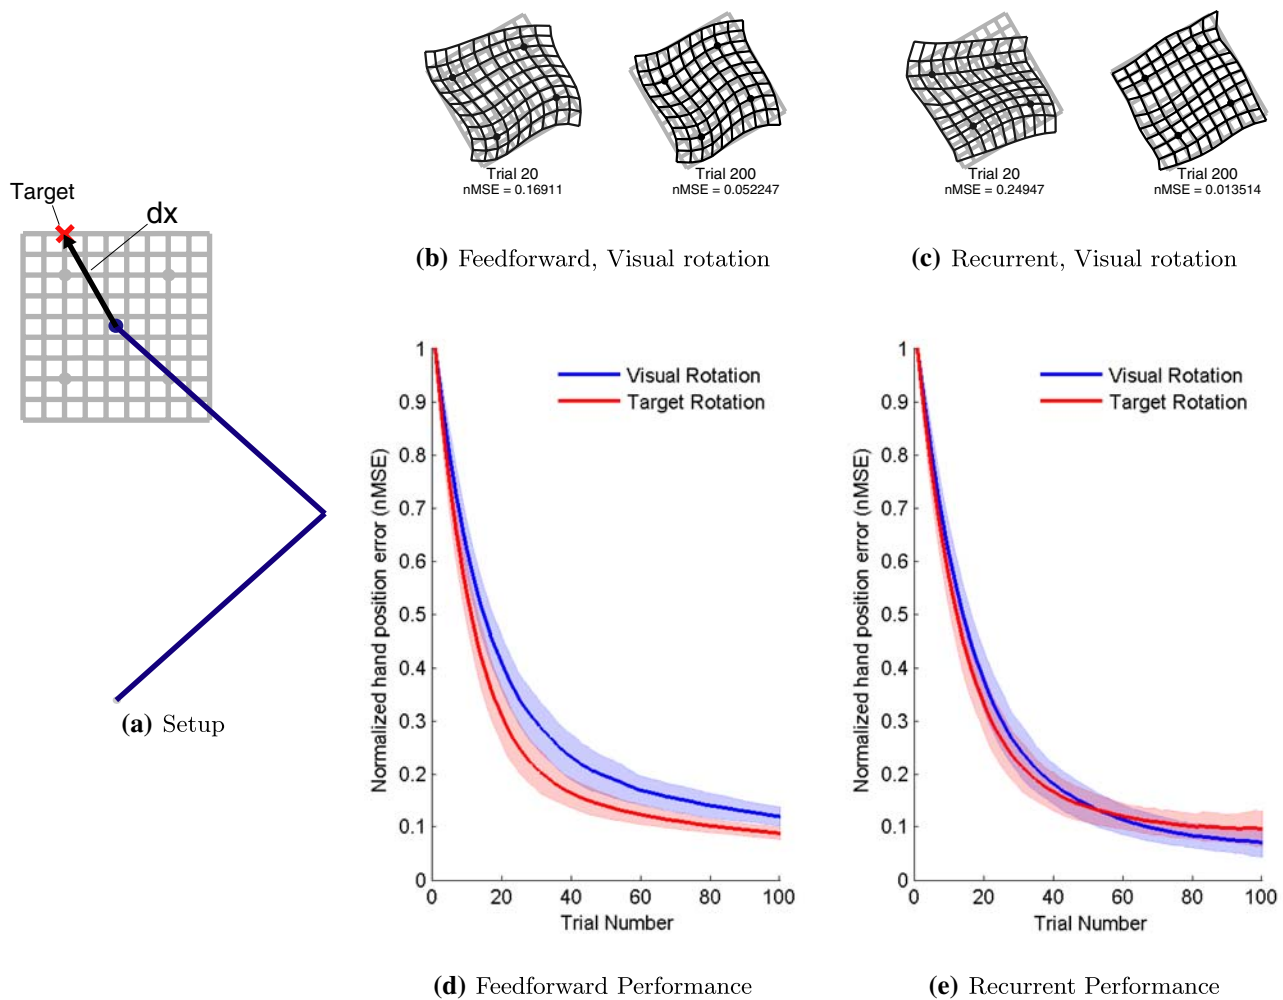


Fig. 7 Learning kinematic control of a two-link planar arm. **a** Experimental setup of arm and grid of observed targets. Difference vector $\mathbf{d}\mathbf{v}$ is estimated from seen target positions (which may be rotated from the actual positions) and this constitutes the stimulus. **b–c** Grid of learnt hand positions following 30° visual rotation for feedforward and recurrent architectures after 20 and 200 trials. Light grey grid shows the actual (rather than seen) final location of targets. *Dark*

grid shows the grid of hand positions attained while testing reaching to all targets following the indicated number of training trials. **d–e** Normalized global mean squared hand position error as a function of the number of trials. *Shaded regions* indicate standard deviation of error across multiple training runs with training targets presented in different (random) sequences. Visual rotation and target rotation are compared for each architecture

over 100 different training runs with different randomly selected target sequences to obtain an impression of the general trends of learning in each condition. The paths of these average positions as training proceeds are plotted in Fig. 8. The actual target locations are marked by a ‘ \times ’. For the feedforward architecture, under the target-shift condition (Fig. 8b), the error estimate used for training is almost equal to the true error and, consequently, the improvement from trial-to-trial comes close to following a straight line in task space. Under the visual rotation condition (Fig. 8a), however, the effect of using a poorer estimate of the error is clearly seen in the fact that the trial-to-trial trend in reach errors does not follow a straight line but an indirect, curved one.

Under the recurrent architecture, there is still a clear difference between the two conditions. However in the visual rotation condition (Fig. 8c), where we expect the cerebellar output error estimate to be correct, the trial-to-trial trend in final hand position is not straight, as in the feedforward/target shift combination, but has slight curvature. This is due to the fact that there is a nonlinear relationship between the improvement in the cerebellar weights and improvements in task-space performance. So although the improvement takes the shortest path in cerebellar weight-space, this does not necessarily correspond to the shortest path of improvement in task space in the same way as it does for the feedforward architecture. In the target shift condition (Fig. 8d), the trend, although different, doesn’t

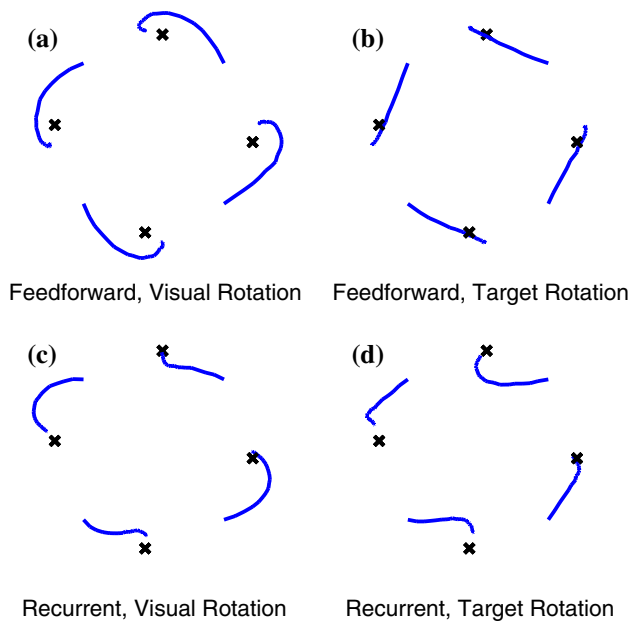


Fig. 8 Time-course of learning for selected targets

appear to be much worse than in the visual rotation condition.

To highlight the fact that learning in the recurrent architecture really was ‘better’ for the visual rotation condition than for the target rotation conditions, we quantified the quality of the estimate of the cerebellar output error in each condition by computing the overlap between the estimated and true error, given by $\hat{\tilde{c}}^T \tilde{c} / \tilde{c}^T \tilde{c}$. An overlap of 1 or close to 1 indicates a good approximation while an overlap of 0 indicates that the estimated error and the true error are perpendicular. Table 1 shows how this overlap varies with increasing ψ and ϕ . As expected, the estimated error was close to perfect in the case of the feedforward architecture/target shift combination and the recurrent architecture/visual rotation combination, independent of the magnitude of the transformation. For the other two conditions (feedforward with visual rotation and recurrent with target rotation), the quality of the estimate clearly diminished with increasing extent of visual field or target rotation.

As in the case of the VOR, the recurrent architecture was unstable in some circumstances with divergence of the motor command during iteration of the recurrent loop. For

transformations with $\psi > 60^\circ$ or $\phi > 60^\circ$, this led to a total breakdown of learning for all targets (indicated by ‘—’ in the table). For transformations of magnitude $\psi = 45^\circ$ or $\phi = 45^\circ$ the divergence of the motor command was only apparent for some targets and not for others. This caused a large discrepancy between the estimated cerebellar output and the error in hand position for these particular targets, leading to the large variation in overlap recorded in Table 1.

5 Conclusion

We have compared and contrasted two previously proposed architectures for cerebellar-based motor adaptation—one feedforward and one recurrent—initially in quite general theoretical terms and then in the context of two specific behaviour: vestibulo-ocular reflex adaptation and reach adaptation.

The central computational difference between the two cerebellar architectures lies in their relative capabilities in adapting to two broad classes of sensorimotor disturbance: (i) changes in the mapping between stimulus and desired outcome and (ii) changes in the mapping between motor command and movement outcome. Because of differences across architectures in how the distal error signal is related to cerebellar output error, the resulting adaptation patterns are qualitatively different. Learning under the recurrent architecture is impaired in response to the first kind of disturbance but not the second, while for the feedforward architecture, this pattern is reversed.

In order to illustrate and validate the theoretical argument, we have simulated learning under both architectures on two biological motor adaptation problems—the VOR and reaching. In each case we have highlighted examples where either the stimulus-desired outcome relationship or the motor command-outcome relationship may be subject to change either naturally or through experimental manipulation. Our simulations confirm our theoretical predictions. However the results also largely indicate that error-driven learning is quite robust with successful, if slow, learning even when the approximating error signal is not that accurate. Only in extreme and generally biologically implausible cases,

Table 1 Quality of cerebellar output error estimates under different conditions

Architecture	Condition	Average $\hat{\tilde{c}}^T \tilde{c} / \tilde{c}^T \tilde{c} \pm$ (SD)					
		15°	30°	45°	60°	75°	90°
Feedforward	Vis. rotation	0.96 ± 0.17	0.87 ± 0.33	0.72 ± 0.48	0.49 ± 0.61	-0.03 ± 0.15	-0.04 ± 0.17
	Target shift	1.00 ± 0.01	1.00 ± 0.01	1.00 ± 0.02	1.00 ± 0.02	1.00 ± 0.03	1.00 ± 0.03
Recurrent	Vis. rotation	1.00 ± 0.00	0.99 ± 0.11	1.05 ± 1.83	—	—	—
	Target shift	0.97 ± 0.00	0.85 ± 0.11	0.82 ± 2.24	—	—	—

when the approximated cerebellar output error is either not correlated or negatively correlated with the true error, is learning not possible at all.

A major problem identified with the recurrent architecture, however, is the threat of instability in the recurrent loop. As a result of this instability, learning was more likely to fail under the recurrent architecture than the feedforward architecture, even under changes in the motor command–outcome mapping, where it is expected to perform well. It is possible that these instability problems may be rectifiable in a biologically-plausible manner through transfer of learning from the cerebellum to the brainstem (Porrill and Dean 2007a).

6 Discussion

One of our main contributions in this work has been to explicitly distinguish two important classes of sensorimotor disturbance: changes in the relationship between the stimulus and the desired outcome, and changes in the relationship between the motor command and the resulting movement outcome. Furthermore, we have shown that these different kinds of disturbance can have different effects on the adaptive response, even though they may initially produce a similar pattern of errors. While the distinction between these disturbance classes has been noted previously (Jordan and Rumelhart 1992; Shadmehr and Wise 2005), here we have highlighted specific examples in the context of the VOR and reaching and examined the impact on learning.

In addition to the VOR and saccades, this distinction is likely to be relevant for many other motor behaviours. In particular, the themes of this paper are highly relevant to saccade adaptation, which is also known to be cerebellar-dependent (Optican and Robinson 1980). Saccade adaptation can be induced by physically adjusting the properties of the oculomotor plant through surgically weakening one or more of the extraocular muscles (Scudder et al. 1998)—a change in the motor command–outcome mapping. A much more common approach, however, is the (non-invasive) target shift paradigm where the target of the saccade is shifted mid-movement (McLaughlin 1967), leading to an apparent overshoot or undershoot. This corresponds to a change in the relationship between the stimulus (initial target position) and the desired outcome (final eye position). It has been noted that these two methods for eliciting saccadic errors result in quite different patterns of adaptation even though initial errors are quite similar (Scudder et al. 1998). While it is possible that additional cues besides visual error are responsible for these differing adaptation patterns, we propose that the difference in learning is at least partly due to the fact that the different disturbances have different effects on the quality of cerebellar learning.

It is difficult, however, to draw any firm conclusions on what architecture the brain actually uses based purely on behavioural data. Actual learning rules in the cerebellum are considerably more complex than the simplified versions we have assumed here and the details of these learning rules are likely to have as much impact on learning as the architecture-dependent effects we have highlighted here. Moreover, it appears that, in reaching tasks, subjects respond fundamentally differently to the two kinds of disturbances, as can be seen through different generalization patterns (Magescas and Prablanc 2006) and different brain activity (Diedrichsen et al. 2005).

Although the example we have given here of a stimulus-desired outcome disturbance during reaching movements may only be plausible in a laboratory setting, it is not too difficult to imagine circumstances in which the ability to learn mappings of this kind may be useful in general. For instance, when catching a ball, a catcher might learn to predict the trajectory of the ball from its mid-flight state and choose a suitable position along this trajectory at which to catch the ball. The movement of the hand to this intermediately calculated position can then be made using the existing reaching circuitry. It may, however, be beneficial to learn a direct relationship between the mid-flight state of the ball and the arm movements which will result in a successful catch. Such a direct strategy would lead to lower latency and introduce less noise than having to maintain an intermediate representation of a desired hand location.

In the case of this catching example, any change in ball dynamics—e.g. caused by wind or a change in elasticity if the ball is to bounce, will affect the relationship between the mid-flight state of the ball (the ‘stimulus’) and the required hand position. Any change in arm kinematics or dynamics, for instance if trying to catch the ball using a net on the end of a heavy pole (note that this would affect both kinematics and dynamics), will alter the relationship between motor command and resulting end-effector motion.

The alternative composite strategy of maintaining and adapting a prediction of the ball’s future trajectory alongside a separate model of arm dynamics can be described in more general terms as a *target-substitution* strategy (Shadmehr and Wise 2005). This strategy has analogs in VOR adaptation (pre-processing of the vestibular signal) and saccadic gain adaptation (re-scaling of the estimated difference vector). How the intermediate target-substitution mapping might be learnt, however, and particularly how changes in the stimulus-desired outcome mapping might be distinguished from changes in the motor-command outcome mapping is unclear since both can give rise to identical errors.

In reality it is quite possible that a combination of target-substitution and direct mapping strategies, as well as feedforward and recurrent strategies are employed by the nervous system. Some behaviours are likely to be more susceptible

to certain kinds of disturbances than others, rendering either a feedforward or recurrent architecture more appropriate.

Biomimetic robotics often also employ cerebellar-like adaptive control strategies, (e.g. [Shibata and Schaal \(2001\)](#)). Understanding the nature of the disturbances which are likely to be encountered should similarly guide selection of which style of architecture may be more appropriate. To this end we have shown elsewhere ([Haith and Vijayakumar 2007](#)) that in certain cases both architectures can be employed concurrently in a feedback-error learning framework to ensure robust adaptation to both kinds of change, combining the advantages of each individual architecture to maximum effect.

Acknowledgments This work was funded in part by the UK EPSRC/MRC through the Neuroinformatics Doctoral Training Centre, University of Edinburgh and by the EU FP6 Integrated Project SENSOPAC. We thank John Porrill and Paul Dean for helpful discussions.

7 Appendix: Implementation details

The same pattern of training was used and very similar control and learning algorithms were employed for simulating each of the behaviours discussed here. Table 2 outlines the basic algorithm underlying all of the simulations.

In each case, an initial motor command–outcome mapping P_0 and an initial stimulus–desired outcome mapping S_0 were specified. The algorithm then simulated cerebellar-based adaptation to a new pair of mappings P_1 and S_1 (in practice only one was varied at a time) using either a feedforward (FF) or recurrent (REC) architecture.

A sequence of stimuli $\mathbf{x}_{1:T}$ was selected. For the VOR, this \mathbf{x}_t represented a discrete-time series of head velocity measurements with a discretization timestep of .01s. For reaching, each \mathbf{x}_t represented a difference vector movement plan for a single trial.

The fixed controller B generates motor commands which are optimal under the initial conditions P_0 and S_0 , i.e.

$$B(\mathbf{x}_t) = P_0^{-1}(S_0(\mathbf{x}_t)). \tag{36}$$

The input to the cerebellum, which we denote by \mathbf{z}_t , varied depending on the architecture employed and the task. For linear P and S (i.e. in the case of VOR), under the forward architecture this was equal to the stimulus \mathbf{x}_t , while under the recurrent architecture this was equal to the motor command \mathbf{u}_t .

For nonlinear P and S (i.e. reaching), \mathbf{z}_t was given by a set of non-linear basis functions Φ defined over the same input space, i.e. $\mathbf{z}_t = \Phi(\mathbf{x}_t)$ for the feedforward architecture and $\mathbf{z} = \Phi(\mathbf{u}_t)$ for the recurrent architecture. The basis functions Φ were Gaussians given by

$$\Phi_i(\mathbf{o}) = e^{(\mathbf{o}-\mathbf{r}_i)^T \Sigma (\mathbf{o}-\mathbf{r}_i)}, \tag{37}$$

Table 2 Pseudocode summary of algorithm used for all simulating all behaviours simulated and for both architectures (Exceptions for particular behaviours are given in parentheses)

Initialize:
 Define stimulus sequence $\mathbf{x}_{1:T}$
 $W_1 = 0$
 ($\mathbf{u}_0 = 0$ for VOR)

Run
 For $t = 1:T$

1. **Generate motor command \mathbf{u}_t**
 if (FF)
 $\mathbf{z}_t = \mathbf{x}_t$ (or $\mathbf{z}_t = \Phi(\mathbf{x}_t)$ for reaching)
 $\mathbf{c}_t = W_t \mathbf{z}_t$
 if (REC)
 $\mathbf{u}_t = 0$
 iterate
 $\mathbf{z}_t = \mathbf{u}_t$ (or $\mathbf{z}_t = \Phi(\mathbf{u}_t)$ for reaching)
 $\mathbf{c}_t = W_t \mathbf{z}_t$
 $\mathbf{u}_t = B(\mathbf{x}_t + \mathbf{c}_t)$
 until convergence of \mathbf{u}_t :
 (or $\mathbf{u}_t = B(\mathbf{x}_t + \mathbf{c}_{t-1})$ for VOR)
2. **Calculate outcome \mathbf{y}_t and observed error $\tilde{\mathbf{y}}_t$**
 $\mathbf{y}_t = P_1(\mathbf{u}_t)$
 $\mathbf{y}_t^* = S_1(\mathbf{u}_t)$
 $\tilde{\mathbf{y}}_t = \mathbf{y}_t^* - \mathbf{y}_t$
3. **Estimate cerebellar error $\tilde{\mathbf{c}}_t$**
 if (FF)
 $\hat{\mathbf{c}}_t = P_0^{-1}(\tilde{\mathbf{y}}_t)$ (or $\hat{\mathbf{c}}_t = J_{P_0^{-1}} \tilde{\mathbf{y}}_t$ for reaching)
 if (REC)
 $\hat{\mathbf{c}}_t = S_0^{-1}(\tilde{\mathbf{y}}_t)$
4. **Update cerebellar weights**
 $W_{t+1} = W_t + \beta \hat{\mathbf{c}}_t \mathbf{z}_t^T$

End

where \mathbf{o} represents the appropriate input (\mathbf{x}_t or \mathbf{u}_t) depending on the architecture. The function centres \mathbf{r}_i were distributed on a uniform square grid in the input space and the metric Σ was chosen so that the width of each tuning function along each dimension was equal to twice the separation $\Delta \mathbf{r}$ between functions:

$$\Sigma = 2 \begin{pmatrix} \frac{1}{\Delta r_1} & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \frac{1}{\Delta r_n} \end{pmatrix}. \tag{38}$$

Note that \mathbf{r} and Σ were different between the two architectures due to different distributions of inputs. A total of 16 basis functions in a 4×4 grid was used in each case.

The cerebellar output \mathbf{c}_t was then given by multiplying the input \mathbf{z}_t by the learnt cerebellar weight matrix W_t ,

Table 3 Summary of model details for different behaviours

Model component	Notation	Description		
		VOR	Saccades	Reaching
Stimulus	\mathbf{x}_t	Head velocity	Target difference vector	Target difference vector
Outcome	\mathbf{y}_t	Eye velocity	Eye displacement	Hand displacement
Motor command	\mathbf{u}_t	Oculomotor torque	Motor amplitude	Change in joint angle
Motor command–outcome mapping	P	Oculomotor dynamics	Oculomotor dynamics	Visual rotation
Stimulus–desired outcome mapping	S	Visual rotation	Target shift	Target shift

$$\mathbf{c}_t = W_t \mathbf{z}_t. \quad (39)$$

The motor command \mathbf{u} was constructed differently for different architectures. For the feedforward architecture, it was given directly by the sum of the cerebellar and brainstem outputs

$$\mathbf{u}_t = B(\mathbf{x}_t) + C(\mathbf{x}_t). \quad (40)$$

For the recurrent architecture, when approximating continuous time dynamics, as in the VOR, the motor command was calculated as

$$\mathbf{u}_t = B(\mathbf{x}_t + C(\mathbf{u}_{t-1})). \quad (41)$$

For simulating single trials of reaching, the motor command was determined by iterating the equation

$$\mathbf{u}_t = B(\mathbf{x}_t + C(\mathbf{u}_t)) \quad (42)$$

until the difference in \mathbf{u} between successive iterations was less than 0.1 %.

In some cases the recurrency led to divergence of \mathbf{u}_t either while iterating within a single trial (reaching) or over time (VOR), in which case the recurrent architecture was unstable and unable to learn the task. However, this was typically only an issue for transformations of moderate to large magnitude and for less severe transformations \mathbf{u}_t converged within 10–20 iterations.

The motor command \mathbf{u}_t was then transformed into an observed output \mathbf{y} via the (transformed) plant dynamics P_1 ,

$$\mathbf{y}_t = P_1(\mathbf{u}_t). \quad (43)$$

For reaching, the converged value of \mathbf{u}_t was used for this.

$$\mathbf{y}_t = P_1(\mathbf{u}_t(1 + \epsilon_t)); \quad \epsilon_t \sim N(0, 0.5^2) \quad (44)$$

The desired outcome at each timestep or trial, \mathbf{y}_t^* , was calculated separately according to the (transformed) stimulus–desired outcome relationship,

$$\mathbf{y}_t^* = S_1(\mathbf{x}_t). \quad (45)$$

The observation error was then calculated as

$$\tilde{\mathbf{y}}_t = \mathbf{y}_t^* - \mathbf{y}_t. \quad (46)$$

To estimate the error in the cerebellar output, the initial mappings between motor command and observed outcome P_0 and between stimulus and desired outcome S_0 were used, according to the theory presented in Sect. 2, i.e. for the recurrent architecture,

$$\hat{\mathbf{c}}_t = \tilde{\mathbf{y}}_t, \quad (47)$$

and for the feedforward architecture,

$$\hat{\mathbf{c}}_t = P_0^{-1} \tilde{\mathbf{y}}_t. \quad (48)$$

Where this mapping was nonlinear (i.e. for reaching) the error was approximated to first order using a Taylor expansion,

$$\hat{\mathbf{c}}_t = J_{P_0^{-1}}(\mathbf{y}_t) \tilde{\mathbf{y}}_t \quad (49)$$

where $J_{P_0^{-1}}(\mathbf{y}_t)$ is the Jacobian of P_0^{-1} at \mathbf{y}_t . This was estimated numerically by finite differences. Note that in all simulations S was linear, although the same principle could be used for approximating $\hat{\mathbf{c}}$ in the recurrent architecture if it were nonlinear.

Finally, the cerebellar weights were updated at each time step using a discrete-time analog of the gradient learning rule stated in Sect. 2.1

$$W_t = W_{t-1} + \beta \tilde{\mathbf{c}}_t \mathbf{p}_t^T, \quad (50)$$

where $\tilde{\mathbf{c}}_t$ is the *estimated* cerebellar output error.

The learning rate β was different in each case and chosen to give approximately realistic timescales of adaptation in comparison to experimental data. The same value of β was always used for both architectures.

References

- Baizer J, Kralj-Hans I, Glickstein M (1999) Cerebellar lesions and prism adaptation in macaque monkeys. *J Neurophysiol* 81(4):1960–1965
- Bastian A, Martin T, Keating J, Thach WT (1996) Cerebellar ataxia: abnormal control of interaction torques across multiple joints. *J Neurophysiol* 76:492–509
- Boyden ES, Katoh A, Raymond JL (2004) Cerebellum-dependent learning: the role of multiple plasticity mechanisms. *Annu Rev Neurosci* 27:581–609

- Coenen OJ, Sejnowski TJ (1996) A dynamical model of context dependencies for the vestibulo-ocular reflex. In: *Advances in neural information processing systems*, vol 8. MIT Press, Cambridge
- Diedrichsen J, Hashambhoy Y, Rane T, Shadmehr R (2005) Neural correlates of reach errors. *J Neurosci* 25:9919–9931
- Gomi H, Kawato M (1990) Learning control for a closed loop system using feedback-error-learning. In: *Proceedings of the 29th conference on decision and control*, pp 3289–3294
- Haith A, Vijayakumar S (2007) Robustness of VOR and OKR adaptation under kinematics and dynamics transformations. In: *Proceedings of 6th IEEE international conference on development and learning (ICDL '07)*, London (2007)
- Hirata Y, Highstein SM (2001) Acute adaptation of the vestibuloocular reflex: signal processing by floccular and ventral parafloccular Purkinje cells. *J Neurophysiol* 85(5):2267–2288
- Imamizu H, Miyauchi S, Tamada T, Sasaki Y, Takino R, Putz B, Yoshioka T, Kawato M (2000) Human cerebellar activity reflecting an acquired internal model of a new tool. *Nature* 403(6766):192–195
- Ito M (2000) Mechanisms of motor learning in the cerebellum. *Brain Res* 886(1–2):237–245
- Jordan MI, Rumelhart DE (1992) Forward models: supervised learning with a distal teacher. *Cognit Sci* 16:307–354
- Kawato M, Gomi H (1992) The cerebellum and VOR/OKR learning models. *Trends Neurosci* 15(11):445–453
- Krakauer JW, Pine ZM, Ghilardi MF, Ghez C (2000) Learning of visuomotor transformations for vectorial planning of reaching trajectories. *J Neurosci* 20(23):8916–8924 (clinical trial)
- Lurito JT, Georgakopoulos T, Georgopoulos AP (1991) Cognitive spatial-motor processes. 7. The making of movements at an angle from a stimulus direction: studies of motor cortical activity at the single cell and population levels. *Exp Brain Res* 87(3):562–580
- Magescas F, Prablanc C (2006) Automatic drive of limb motor plasticity. *J Cogn Neurosci* 18:75–83
- McLaughlin S (1967) Parametric adjustment in saccadic eye movement. *Percept Psychophys* 2:359–362
- Optican LM, Robinson DA (1980) Cerebellar-dependent adaptive control of primate saccadic system. *J Neurophysiol* 44(6):1058–1076
- Porrill J, Dean P (2007a) Cerebellar motor learning: when is cortical plasticity not enough? *PLoS Comput Biol* 3:1935–1950
- Porrill J, Dean P (2007b) Recurrent cerebellar loops simplify adaptive control of redundant and nonlinear motor systems. *Neural Comput* 19(1):170–193
- Porrill J, Dean P, Stone JV (2004) Recurrent cerebellar architecture solves the motor-error problem. *Proc Biol Sci* 271(1541):789–796
- Robinson D (1964) The mechanics of human saccadic eye movement. *J Physiol* 174:245–264
- Schweighofer N, Arbib MA, Kawato M (1998) Role of the cerebellum in reaching movements in humans. I. Distributed inverse dynamics control. *Eur J Neurosci* 10(1):86–94
- Scudder CA, Baturina EY, Tunder GS (1998) Comparison of two methods of producing adaptation of saccade size and implications for the site of plasticity. *J Neurophysiol* 79(2):704–715 (comparative study)
- Shadmehr R, Mussa-Ivaldi FA (1994) Adaptive representation of dynamics during learning of a motor task. *J Neurosci* 14(5 Pt 2):3208–3224
- Shadmehr R, Wise SP (2005) *The computational neurobiology of reaching and pointing: a foundation for motor learning*. MIT Press, New York
- Shibata T, Schaal S (2001) Biomimetic gaze stabilization based on feedback-error-learning with nonparametric regression networks. *Neural Netw* 14(2):201–216
- Simpson J, Wylie D, De Zeeuw C (1996) On climbing fiber signals and their consequence(s). *Behav Brain Sci* 19(3):339–527
- Smith M, Shadmehr R (2005) Intact ability to learn internal models of arm dynamics in huntington's but not cerebellar degeneration. *J Neurophysiol* 93:2809–2821
- Todorov E (2000) Direct cortical control of muscle activation in voluntary arm movements: a model. *Nat Neurosci* 3(4):391–398
- Wolpert DM, Miall RC, Kawato M (1998) Internal models in the cerebellum. *Trends Cognit Sci* 2:338–347
- Yutaka H, Akimasa Y, Blazquez PM, Highstein SM (2005) Evaluation of the inverse dynamic model in cerebellum during visual-vestibular interactions at different VOR gains in squirrel monkeys. *Neurocomputing* 65–66:709–717