

Robustness of VOR and OKR adaptation under kinematics and dynamics transformations

Adrian Haith and Sethu Vijayakumar

Institute of Perception, Action and Behaviour, School of Informatics, University of Edinburgh
 {adrian.haith, sethu.vijayakumar}@ed.ac.uk

Abstract—Many computational models of vestibulo-ocular reflex (VOR) adaptation have been proposed, however none of these models have explicitly highlighted the distinction between adaptation to dynamics transformations, in which the intrinsic properties of the oculomotor plant change, and kinematic transformations, in which the extrinsic relationship between head velocity and desired eye velocity changes (most VOR adaptation experiments use kinematic transformations to manipulate the desired response). We show that whether a transformation is kinematic or dynamic in nature has a strong impact upon the speed and stability of learning for different control architectures. Specifically, models based on a purely feedforward control architecture, as is commonly used in feedback-error learning (FEL), are guaranteed to be stable under kinematic transformations, but are susceptible to slow convergence and instability under dynamics transformations. On the other hand, models based on a recurrent cerebellar architecture [7] perform well under dynamics but not kinematics transformations. We apply this insight to derive a new model of the VOR/OKR system which is stable against transformations of both the plant dynamics and the task kinematics.

I. INTRODUCTION

The vestibulo-ocular reflex in mammals acts during head rotations to stabilize gaze by counter-rotating the eyes. The characteristics of this reflex are not fixed, but can be modified through experience. For instance, if lenses are worn which magnify the image of the external environment on the retina, the required eye movement to stabilize gaze following a given head movement changes. 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 [2].

II. COMPUTATIONAL MODELS OF VOR ADAPTATION

Computationally, VOR adaptation is often viewed as the process of learning an inverse dynamics model of the oculomotor plant, which may change through growth, ageing or disease. The role of the task kinematics is often overlooked. Fig. 1 gives a schematic illustration of the vestibulo-ocular system. P represents the oculomotor plant, B represents the neural control circuitry, and K represents the task kinematics which dictate the required eye movement following a given head movement.

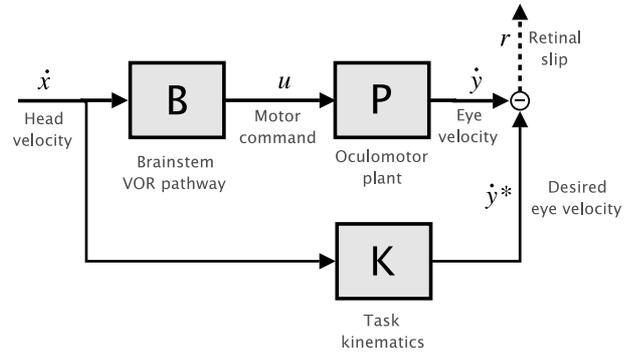


Fig. 1. Basic framework of the VOR model.

Since the inertia of the oculomotor plant is very low, we will assume a first order oculomotor plant for which the forward dynamics is given by

$$P : (y, u) \rightarrow \dot{y}, \quad (1)$$

where y is the current eye position, u is the motor command supplied to the plant and \dot{y} is the resulting eye velocity. The inverse dynamics model is correspondingly defined as

$$P^{-1} : (y, \dot{y}^*) \rightarrow u. \quad (2)$$

P^{-1} maps the current eye position y and some desired eye velocity \dot{y}^* to a motor command u which achieves this.

Let us define retinal slip r as the velocity of the image on the retina in degrees/sec. In the context of the VOR and OKR, the desired eye velocity \dot{y}^* is defined as the unique eye velocity that results in zero retinal slip. \dot{y}^* depends primarily on the angular velocity of the head, \dot{x} , which is measured by the vestibular system (\dot{y}^* may also depend on eye's angular position y [8]). u^* is defined as the motor command which achieves the desired eye velocity \dot{y}^* when supplied to the plant,

$$u^* = P^{-1}(y, \dot{y}^*). \quad (3)$$

In most VOR models, desired eye velocity is taken as equal and opposite to head velocity, i.e. $\dot{y}^* = -\dot{x}$. 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 in phase with head movements.

VOR adaptation therefore cannot be regarded as simply a process of learning the inverse dynamics of the oculomotor plant. Fundamental to VOR adaptation is the need to adapt to changes in the kinematics of the task, that is, the relationship between head movement and required eye movement, as illustrated in Fig. 1.

More formally, we can define a forward kinematics model

$$K : (y, \dot{x}) \rightarrow \dot{y}^* \quad (4)$$

and corresponding inverse kinematics model

$$K^{-1} : (y, \dot{y}) \rightarrow \dot{x}. \quad (5)$$

K gives the desired (gaze-stabilizing) eye movement as a function of head velocity \dot{x} and eye position y , while K^{-1} gives the head rotation that would have required an eye movement \dot{y}^* to stabilize gaze.

In general, the kinematic mapping K incorporates effects due to distortion of the visual image before reaching the eye, movements of the external scene which are correlated with head movement (visual-vestibular mismatch), off-axis effects and inaccuracies or nonlinearities in the vestibular measurements of \dot{x} . VOR adaptation must, either explicitly or implicitly, reflect learning of this kinematic mapping as well as the plant dynamics.

Hence from (2) and (4) the overall mapping which must be learnt is a composite function

$$u = P^{-1}(y, K(y, \dot{x})). \quad (6)$$

This distinction between kinematics and dynamics and their different roles in the motor task *must* be taken into account in modelling the VOR. As we shall see, the failure of previous models to do so means they are susceptible to instability following changes in either the kinematics or the dynamics.

III. CEREBELLAR LEARNING OF INTERNAL MODELS

It is well known that the cerebellum plays a crucial role in VOR adaptation. The VOR is comprised of a fixed, direct pathway through the brainstem, which is augmented by an adaptive pathway through the cerebellum. Physiological studies have shown that as well as receiving the same vestibular input as the brainstem, the cerebellum receives an efferent copy of the final motor command [4]. Most VOR models have tended to disregard these recurrent connections, modelling the brainstem and cerebellar pathways as having a purely feedforward architecture [3], [6], [8] (see Fig. 2(a)). More recently, Porrill and Dean [7] have advocated a recurrent cerebellar architecture model which emphasizes the efferent motor command inputs to the cerebellum (Fig. 2(b)). The

choice of architecture (feedforward or recurrent) turns out to have a significant impact upon learning, as we shall see.

For the purpose of illustration, we will assume the standard Marr-Albus-Ito model of cerebellar learning [5] which essentially models the cerebellum as a single layer neural network. The cerebellar output is given by

$$c(t) = \sum_i w_i g_i(t),$$

where $g_i(t)$ is the activity of the i th parallel fibre and w_i is the strength of the corresponding Parallel fibre-Purkinje cell synapse.

Learning is driven by the climbing fiber signal $e(t)$ which corresponds to the error in $c(t)$. A simple learning rule which approximately captures the known plasticity laws at the synapses is

$$\dot{w}_i = -\beta g_i(t) e(t). \quad (7)$$

This is equivalent to performing gradient descent on the squared error, *provided* $e(t)$ 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 retinal slip r and the error in the cerebellar output e_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 have only taken dynamics transformations into account. In this case, (i.e. ignoring kinematics for the moment) under the feedforward architecture, the error in the cerebellar output is equal to the error in the motor command. Porrill and Dean [7] have shown that under the recurrent architecture, the error in the cerebellar output is in fact equal to the raw retinal slip signal, circumventing the distal learning problem. However, both of these results were obtained by assuming that the desired eye velocity is exactly equal to the current head velocity (i.e. trivial kinematics).

We will now examine learning within each of these architectures when both the kinematics and the dynamics are subject to change. In both cases, we require an expression for the error in the cerebellar output, e_C , in terms of the observed retinal slip, r . For simplicity we will assume linear plant, cerebellum and brainstem dynamics. We will point out in the derivations exactly where linearity is exploited.

A. Feedforward Architecture

Let us refer to the feedforward architecture shown in Fig. 2(a). The inputs to the cerebellum are head velocity \dot{x} and head position x (omitted from the figure for clarity). We assume that an optimal cerebellar model C^* exists. The error

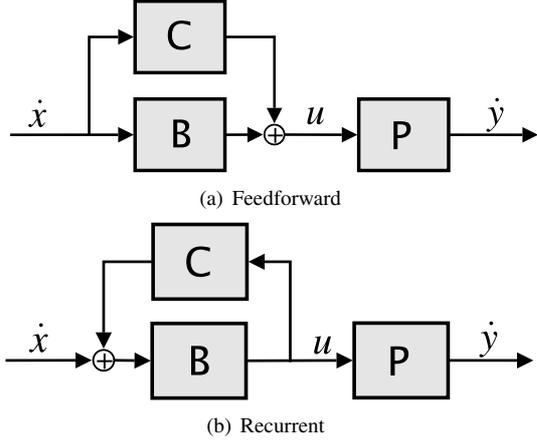


Fig. 2. Schematic of feedforward and recurrent architectures

in the cerebellar output is then defined as

$$e_C(x, \dot{x}) = C^*(x, \dot{x}) - C(x, \dot{x}). \quad (8)$$

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

$$u = C(x, \dot{x}) + B(x, \dot{x}), \quad (9)$$

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

$$u^* = C^*(x, \dot{x}) + B(x, \dot{x}). \quad (10)$$

Taking the difference between (9) and (10) gives us an expression for the error in the cerebellar output:

$$e_C = u^* - u, \quad (11)$$

that is, e_C is equal to the *motor error*. Rewriting the right hand side of (11) in terms of the inverse plant model, exploiting linearity of the plant to combine the two terms and using $r = \dot{y}^* - \dot{y}$ for the retinal slip, we have

$$e_C = P^{-1}(y, r). \quad (12)$$

Equation (12) shows that the error in the cerebellar output can be calculated from the retinal slip via the inverse dynamics of the plant. We assume that some internal model is available to compute this, but if the plant dynamics change then we can no longer be confident that our estimate of the cerebellar output is accurate.

The required training signal is, however, independent of the kinematics. This is an important but usually overlooked advantage of employing a purely feedforward architecture.

So in general we expect learning under the feedforward architecture to be impaired (converge more slowly) or even made entirely unstable (not converge at all) following a change in the dynamics. However, we expect learning to be unaffected by a change in the kinematics. In Section V, we verify these theoretical claims with a simulation of a 2 degree-of-freedom oculomotor plant.

B. Recurrent architecture

We wish to derive an expression for the error in the cerebellar output in terms of the measured retinal slip for the recurrent architecture (Fig. 2(b)). We assume that the inputs to the cerebellum are the head position x and the afferent motor command u , the latter being the defining characteristic of the recurrent architecture (the head-position input is omitted in Fig. 2(b) for clarity).

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

$$C(x, u) + \dot{x} = B^{-1}(x, u). \quad (13)$$

Again, as in the feedforward case, we assume there exists an optimal cerebellar model C^* which yields exactly the desired motor command u^* . This will obey

$$C^*(x, u^*) + \dot{x} = B^{-1}(x, u^*). \quad (14)$$

Note that since $\dot{x} = K^{-1}(y, \dot{y}^*)$ and $\dot{y}^* = P(y, u^*)$, we can expand (14) to obtain an identity in terms of u^* :

$$B^{-1}(x, u^*) - C^*(x, u^*) = K^{-1}(y, P(y, u^*)) \quad (15)$$

which we will make use of later. Now, taking the difference between (13) and (14), we have (by linearity of B^{-1})

$$B^{-1}(x, u^* - u) = (C^*(x, u^*) - C(x, u)) \quad (16)$$

We can, however, use the linearity of C^* to re-express $C^*(u^*) = C^*(u) + C^*(e_u)$ (where $e_u \equiv u^* - u$). Substituting this into (16) and applying (8) gives us an expression for e_C in terms of the motor error e_u .

$$e_C(x, u) = B^{-1}(x, e_u) - C^*(x, e_u). \quad (17)$$

We can now apply (15) to the right-hand side of (17) and obtain

$$e_C(x, u) = K^{-1}(y, P(y, e_u)) = K^{-1}(y, r). \quad (18)$$

Equation (18) shows that the error in the cerebellar output is given by the retinal slip transformed via the inverse kinematics. This can be thought of as an error in the original vestibular signal \dot{x} , given fixed cerebellum and brainstem.

We expect learning under the recurrent architecture to be impaired under kinematics transformations, but to be unaffected by dynamics transformations. Again, in Section V, we verify these theoretical claims with simulations.

This reveals a duality between the feedforward and recurrent architectures for the cerebellar pathway. The (in)stability properties of learning in the feedforward architecture under dynamics transformations are mirrored by those of learning in the recurrent architecture under kinematics transformations.

In the case of nonlinear dynamics, a Taylor expansion argument can be used to show that the critical observation of learning in the feedforward architecture being independent of the kinematics and learning in the recurrent architecture being independent of the dynamics still holds.

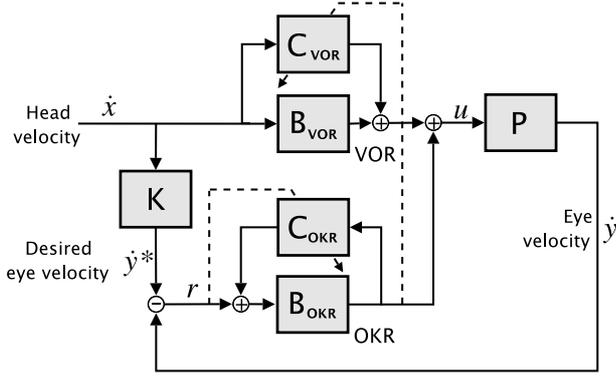


Fig. 3. Composite model of VOR/OKR adaptation.

Stability of recurrent loop

A major problem with the recurrent architecture is the threat of instability in the brainstem-cerebellum loop. In the case of linear dynamics, the motor command u will grow exponentially if any of the eigenvalues λ of the matrix BC has magnitude $|\lambda| > 1$, since the associated eigenvector will increase in magnitude each time it traverses the loop.

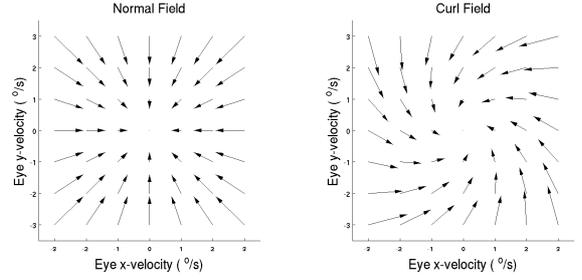
It may, however, be possible to avoid entering into unstable regions of the parameter space by adapting B , using C as a training signal. This learning transfer from C to B would steer the loop away from regions of instability by ensuring that C^* (which would now depend on B) would tend asymptotically to 0. Learning transfer of this kind is supported by physiological evidence [2] and Dean and Porrill [1] have suggested it may be used as a mechanism to enhance VOR response at high frequencies.

IV. FEEDBACK-ERROR LEARNING WITH ADAPTIVE OKR

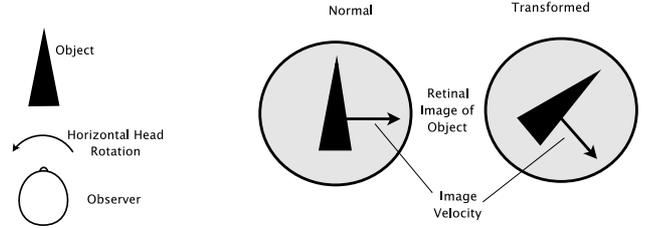
We have seen two alternative models of vestibulo-ocular reflex adaptation, each of which have distinct shortcomings which are complimentary to one another. The feedforward architecture is susceptible to instability under dynamics transformations, whereas the recurrent architecture is susceptible to instability under kinematics transformations.

Shibata and Schaal [8] implemented a feedback-error-learning (FEL) model [3] of the VOR. Here, the output of an existing feedback controller is used as an approximation to the motor error e_u for training the cerebellum. The optokinetic response (OKR) is a reflex which generates compensatory eye movements in response to retinal slip. Shibata and Schaal likened this to a conventional feedback controller and used the output to train the cerebellum according to the standard feedback-error-learning framework.

However, this framework inherits all the problems of the feedforward architecture. The feedback controller implicitly contains a model of the plant dynamics and if these dynamics change, the feedback controller may no longer generate suitable motor commands or training signals. One



(a) Dynamics transformation



(b) Kinematics transformation

Fig. 4. Example of dynamic and kinematic transformations. 4(a) illustrates a viscous dynamics transformation. The original force field (as a function of eye angular velocity) is plotted on the left, and the transformed viscous curl field is plotted on the left. 4(b) illustrates the consequences of a rotation of the visual field by 45° .

solution to this problem is to independently adapt the OKR in addition to the VOR.

In Section II, we described how, in the VOR, learning the kinematic relationship between head velocity and desired eye velocity is a central component of adaptation. In the OKR, the analog of head velocity is retinal slip. Unlike head velocity, however, the relationship between retinal slip and desired eye velocity is veridical. Kinematics, therefore, do not play a role in the OKR and adaptation of the OKR *can* be described purely in terms of learning the inverse dynamics. Consequently, the recurrent architecture is ideally suited for adaptation of the OKR, since, as outlined above, it is robust against dynamics transformations.

We therefore propose a joint VOR/OKR model based on feedback-error-learning in which the OKR adapts to dynamics transformations using a recurrent cerebellar architecture framework. The output of the OKR then also acts as a training signal for the VOR, which employs the feedforward architecture just as in Shibata and Schaal's model [8]. This full joint adaptive FEL model of the VOR and OKR is illustrated in Fig. 3.

V. SIMULATION RESULTS

We tested the ability of each of the three proposed VOR architectures to adapt to a range of changes in the kinematics and the dynamics. We simulated a 2 degree-of-freedom oculomotor plant with linear, viscous dynamics given by

$$\mathbf{u} = P^{-1}\mathbf{y}. \quad (19)$$

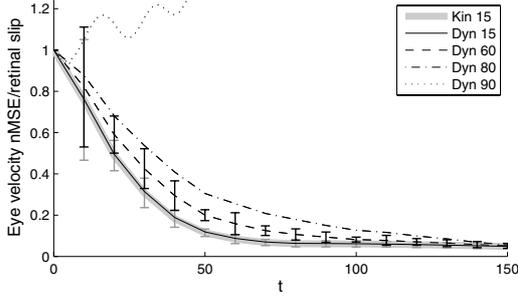


Fig. 5. Feedforward architecture performance under a viscous curl field. Different traces plot the nMSE over time for different angles of viscous force direction (i.e. different values of θ in (21)).

The task kinematics were given by

$$\dot{\mathbf{y}}^* = K\dot{\mathbf{x}}. \quad (20)$$

The dynamics transformation we employed was to change the dynamics from an ordinary resistive viscosity field ($P^{-1} = I$) to a viscous curl field (Fig. 4(a)) in which there is an angle θ between the eye velocity and the force, i.e.

$$P^{-1} = \begin{pmatrix} \cos \theta & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{pmatrix}. \quad (21)$$

The kinematics transformation we employed was a rotation of the visual field (Fig. 4(b)) by angle ψ . In the initial case, we assume that eye velocity is exactly equal to head velocity, i.e. $\dot{\mathbf{y}}^* = \dot{\mathbf{x}}$ ($K = I$). Following the transformation, the desired eye velocity is rotated by angle ψ relative to the head velocity ($K = (\cos \psi, \sin \psi; -\sin \psi, \cos \psi)$). The brainstem controller was optimal for the initial dynamics and kinematics. An OKR feedback controller was also employed, which generated appropriate feedback commands given the initial plant dynamics, with a conservative gain of 0.1

$$u_{OKR} = 0.1P^{-1}(\mathbf{y}, \mathbf{r}). \quad (22)$$

In all experiments, the head position repeatedly traced out a figure-of-eight, i.e. $\mathbf{x}(t) = [\sin(0.1t), \sin(0.2t)]^T$. All experiments were run 10 times with different initial positions around the figure-of-eight and simulation timestep of $dt = 0.01s$.

Firstly, we tested the performance of the feedforward architecture in adapting to the visuomotor rotation. Analyzing the normalised mean-squared velocity error (nMSE) or retinal slip, we found no significant difference between different magnitudes of rotation (values of ψ). Fig. 5 plots the average nMSE over time for $\theta = 15^\circ$ which is representative of all values of θ . 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. Fig. 5 plots the evolution of the nMSE over time for different values of

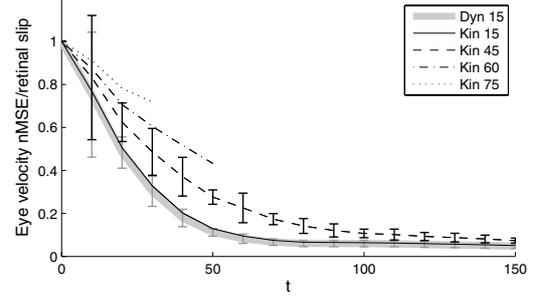


Fig. 6. Recurrent architecture performance under visual field rotations of varying magnitudes (i.e. different values of ψ).

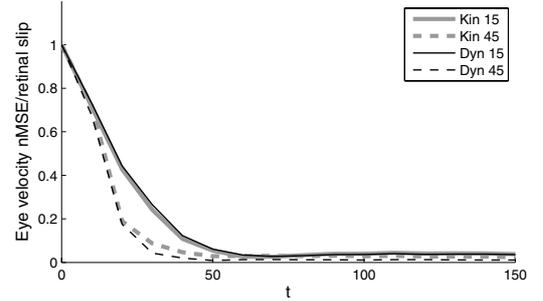
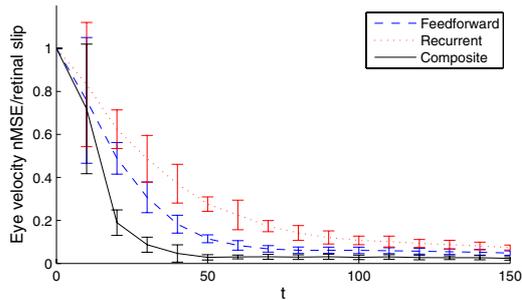


Fig. 7. Composite architecture performance under dynamics and kinematic transformations

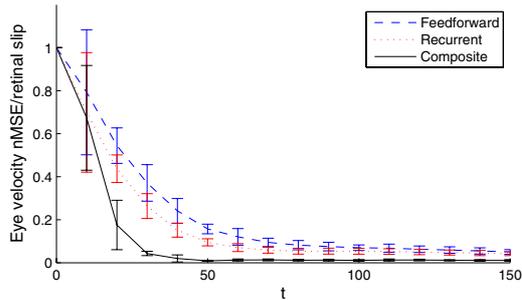
θ . For $\theta = 15^\circ$, performance is the same as under the kinematics transformation. As θ increases, however, the rate of improvement drops. Error bars are plotted for the $\theta = 60^\circ$ case to show that this difference is significant (error bars on other plots are omitted for clarity). At $\theta = 90^\circ$, the VOR no longer converges. In this instance, data from a single representative trial is plotted.

For the recurrent architecture, we first tested the performance under the dynamics transformation. For $\theta < 60^\circ$, we found no significant difference in performance between different values of θ . For other values of θ , however, the recurrent loop tended to become unstable after a period of initial improvement. Fig. 6 shows the nMSE over time for $\theta = 15^\circ$, which was representative of all trials for $\theta < 60^\circ$.

Also plotted in Fig. 6 are the results for the recurrent architecture adapting under the kinematics transformation. Again, for clarity, error bars are only plotted for representative transformations. For $\psi = 15^\circ$, performance is similar to that under the dynamics transformation. For $\psi = 45^\circ$, however, the adaptation is significantly slower. For $\psi \geq 60^\circ$, the recurrent loop tended to become unstable. Once the loop became unstable, the error tended to explode over very few timesteps and the plots are therefore curtailed at this point. Before hitting instability of the recurrent loop, the initial rate of improvement can be seen to diminish with increasing ψ .



(a) Dynamics transformation: Viscous curl field ($\theta = 45^\circ$).



(b) Kinematics transformation: Visual field rotation ($\psi = 45^\circ$).

Fig. 8. Comparative control performance of all three architectures under similar and typical (a) dynamic and (b) kinematic transformations.

Finally, we evaluated the performance of the composite architecture under the kinematics and dynamics transformations. Fig. 7 plots the nMSE over time for two kinematics and dynamics transformations (error bars are omitted for clarity). In contrast to the other two architectures, changes in the magnitude of the transformation had a similar effect, whether kinematic or dynamic. Interestingly, the rate of improvement actually increased for larger magnitude transformations. This was due to the *gain* of the OKR increasing during adaptation and the effect was greater for more severe transformations. Effectively, this constituted an increase in the learning rate for the VOR. The composite architecture also inherited the problems of instability in the recurrent loop from the recurrent architecture and became unstable under the dynamics transformation for $\theta \geq 60^\circ$.

Fig. 8 compares the performance of all three architectures under both dynamics and kinematics transformations. The composite architecture outperforms both the feedforward and recurrent architectures under both kinematic and dynamic transformations of the plant.

VI. CONCLUSION

We have examined existing control-theoretic models of VOR adaptation with respect to their stability under dynamics and kinematic transformations of the task. We found that the feedforward architecture typically employed in feedback-error learning is not well-suited to situations where the

dynamics are subject to change, while the alternative recurrent architecture is not suited to situations in which the kinematics are subject to change. We applied this insight to propose a new model of VOR adaptation in which, as in the feedback-error learning model of the VOR, adaptation is driven by the output of the OKR. In our model, however, the OKR itself is adaptive, employing the recurrent cerebellar architecture to ensure robustness against changes in plant dynamics while exploiting the fact that the OKR is not susceptible to kinematic variability. The overall composite system, therefore, has the potential to learn stably under both kinematic/dynamic transformations. To our knowledge, no other previously published model of VOR adaptation can claim robustness in this way.

The threat of instability of the recurrent loop is a major issue with the recurrent architecture. It may, however, be possible to avoid this through a combination of constraining the magnitude of the control signals and/or the cerebellar output and transfer of learning from the cerebellum to the brainstem. If the recurrent loop can be stabilized in this way, then the learning in the composite architecture should be stable across all conditions in which either one of the forward or recurrent architecture can learn stably in isolation. Finally, we chose to simulate a simple linear, viscous plant here to highlight the fact even in this simplest of cases, the distinction between kinematics and dynamics transformations is a crucial one. It should, however, be possible to extend these results to higher-order and non-linear plants.

ACKNOWLEDGEMENTS

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.

REFERENCES

- [1] Paul Dean and John Porrill. Plant compensation in vestibulo-ocular reflex: Computational analysis points to multiple points of plasticity. Poster, Society for Neuroscience meeting, 2004.
- [2] Lisberger S. G., Bronte-Stewart H.M. Pavelko T. A., and Stone L.S. Neural basis for motor learning in the vestibuloocular reflex of primates. *Journal of Neurophysiology*, 72(2):954–973, 1994.
- [3] Hiroaki Gomi and Mitsuo Kawato. Learning control for a closed loop system using feedback-error-learning. In *Proceedings of the 29th conference on decision and control*, pages 3289–3294, Dec 1990.
- [4] Y Hirata and S M Highstein. Acute adaptation of the vestibuloocular reflex: signal processing by floccular and ventral parafloccular Purkinje cells. *J Neurophysiol*, 85(5):2267–2288, May 2001.
- [5] M Ito. Mechanisms of motor learning in the cerebellum. *Brain Research*, 886(1-2):237–245, Dec 2000.
- [6] M Kawato and H Gomi. The cerebellum and VOR/OKR learning models. *Trends in Neuroscience*, 15(11):445–453, Nov 1992.
- [7] John Porrill, Paul Dean, and James V Stone. Recurrent cerebellar architecture solves the motor-error problem. *Proceedings in Biological Sciences*, 271(1541):789–796, Apr 2004.
- [8] T Shibata and S Schaal. Biomimetic gaze stabilization based on feedback-error-learning with nonparametric regression networks. *Neural Networks*, 14(2):201–216, Mar 2001.