Recurrent cerebellar loops simplify adaptive control of redundant and non-linear motor systems

John Porrill & Paul Dean
Centre for Signal Processing in Neural Imaging & Systems Neuroscience
University of Sheffield

We have described elsewhere an adaptive filter model of cerebellar learning in which the cerebellar microcircuit acts to decorrelate motor commands from their sensory consequences (Dean, Porrill & Stone, Proc Roy Soc B, 2002). We found that stable learning for generic motor plants required the cerebellar microcircuit to be embedded in a recurrent loop; this arrangement lead to a simple and modular adaptive control architecture when applied to the linearised 3D vestibular ocular reflex (Porrill, Dean & Stone, Proc Roy Soc B, 2004). Here we investigate the properties of recurrent loop connectivity for the case of redundant and non-linear motor systems using the example of kinematic control of a robot arm. We will demonstrate that

- the learning rule does not depend on the availability of 'motor error' signals or require complex neural 'reference structures' to estimate such signals. It is thus a genuinely local learning architecture in which proximal rather than distal error can be used for learning
- the form of the learning rule is biologically plausible, independent of the details of the controlled plant and in particular is unaffected by motor plant non-linearities
- control of redundant systems is not subject to the 'convexity problem' in which incorrect average motor commands are learnt for end-effector positions which can be accessed by more than one arm configuration
- cerebellar connectivity is intrinsically modular and its complexity scales linearly with the dimensionality *N* of output space rather than with the product of *N* and the (for highly redundant biological motor systems usually much higher) dimensionality of motor command space.

These properties suggest that the highly specific closed cerebellar loops which have been shown to be ubiquitous in motor systems (Kelly & Strick, 2003) may play a central role in simplifying the adaptive control problem.

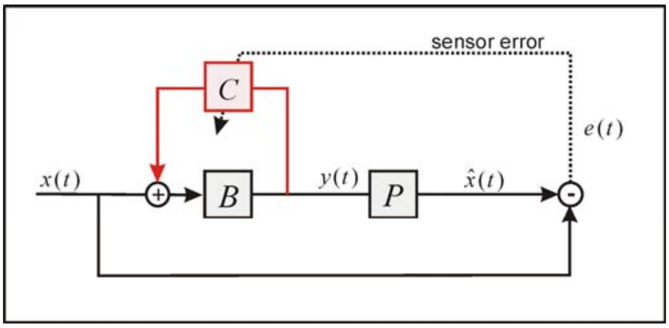

Figure 1. *Recurrent cerebellar architecture applied to kinematic control of a robot arm.* Robot arm motor plant (forward kinematics) is denoted by P. Desired end effector positions $x(t)$ are processed by a recurrent loop containing a fixed approximation B to the inverse kinematics P^{-1} and an adaptive element $C = \Sigma w_i G_i$ representing the cerebellar microcircuit. Reaching error $e = \hat{x} - x$ is used to train the weights w_i using the local learning rule $\dot{w}_i = -\beta e G_i(y)$. Convergence follows since sum square synaptic error $V = 1/2\Sigma(w_i - w_i^{\text{opt}})^2$ is a Lyapounov function with $\dot{V} = -\beta e^2$. This behaviour contrasts with conventional forward loop architectures where an approximation to the Jacobian of P is required to enable learning.

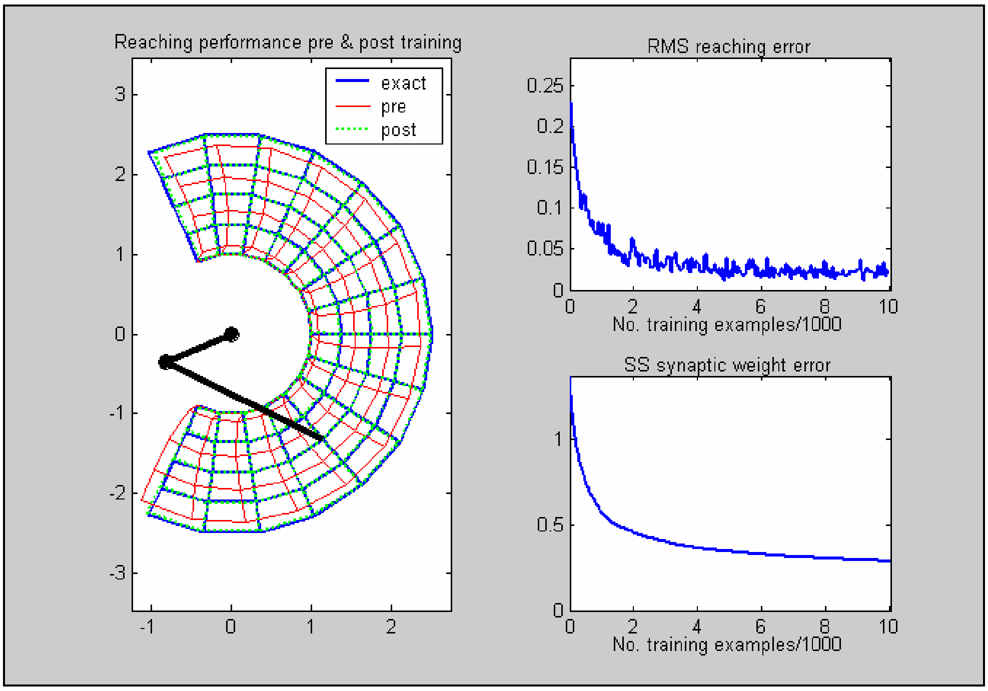

Figure 2. *Adaptive re-calibration of a two joint robot arm.* Left panel shows simulated two joint arm. Initially B (see Fig 1) is an exact inverse kinematics for arm lengths $L_1 = 1$, $L_2 = 2$ and C is an RBF net with weights $w_i = 0$: in this case the arm reaches accurately to points within the workspace outlined by the blue grid. If arm lengths are changed to $L_1 = 0.9$, $L_2 = 2.2$ the arm reaches to inaccurate positions shown by the distorted pre-training (red) grid. Performance after 10^4 training examples is shown by the post-training (dotted green) grid. Top right panel shows stochastic decrease of RMS reaching error during training. Bottom right panel shows monotonic decrease in V (exactly as predicted by the Lyapounov analysis). Note that much faster convergence can be obtained using optimised bases G_i.