



UNSUPERVISED LEARNING OF GRANULE CELL SPARSE CODES ENHANCES CEREBELLAR ADAPTIVE CONTROL

N. SCHWEIGHOFER,^{a,b,*} K. DOYA^a and F. LAY^c

^aERATO Japan Science and Technology Corporation, 2-2, Hikaridai, Seika-cho, Soraku-gun, Kyoto 619-0288, Japan

^bLearning Curve, K.K., 2F Fuji Building 40, 15-14 Sakuragaoka-cho, Shibuya-ku, Tokyo 150-0031, Japan

^cInstitut National des Telecommunications, 91 Evry, France

Abstract—Marr [*J. Physiol.* (1969) **202**, 437–470] and Albus [*Math. Biosci.* (1971) **10**, 25–61] hypothesized that cerebellar learning is facilitated by a granule cell sparse code, i.e. a neural code in which the fraction of active neurons is low at any one time. In this paper, we re-examine this hypothesis in light of recent experimental and theoretical findings. We argue that cerebellar motor learning is enhanced by a sparse code that simultaneously maximizes information transfer between mossy fibers and granule cells, minimizes redundancies between granule cell discharges, and re-codes the mossy fiber inputs with an adaptive resolution such that inputs corresponding to large errors are finely encoded. We then propose that a set of biologically plausible unsupervised learning rules can produce such a code. To maintain a low mean firing rate compatible with a sparse code, an activity-dependent homeostatic mechanism sets the cells' thresholds. Then, to maximize information transfer, the mossy fiber–granule cell synapses are adjusted by a Hebbian rule. Furthermore, to minimize redundancies between granule cell discharges, the inhibitory Golgi cell–granule cell synapses are tuned by an anti-Hebbian rule. Finally, to allow adaptive resolution, a performance-based neuromodulator-like signal gates these three plastic processes. We integrate these gated learning rules into a simplified model of the cerebellum for arm movement control, and show that unsupervised learning of granule cell sparse codes greatly improves cerebellar adaptive motor control in comparison to a “fixed” Marr–Albus-type model.

Until recently, activity-dependent cerebellar plasticity was thought to be largely confined to the granule cell–Purkinje cell synapses. This static view of the cerebellum is, however, quickly being replaced by an extremely dynamic view in which plasticity is omnipresent. The present theoretical study shows how several forms of plasticity in the granular layer of the cerebellum can produce fast, accurate and stable cerebellar learning. © 2001 IBRO. Published by Elsevier Science Ltd. All rights reserved.

Key words: anti-Hebbian learning, cerebellum, Golgi cells, granule cells, Hebbian learning, sparse code.

Although it is well established that the cerebellum is involved in motor coordination, information processing in the cerebellar cortex is still unclear. In particular, it is not understood how the granule cells (GCs) transform the mossy fiber (MF) inputs for further processing by the Purkinje cell (PC) layer.

The cerebellar input is characterized by its divergence from the MFs to the GCs: there are more than 100 times as many GCs as MFs. Each GC receives an average of four MF inputs and gives excitatory projections to the PC and to the Golgi cells (GOs), which in turn inhibit the GCs extensively. Marr³⁷ and Albus² proposed that the

cerebellum is an associative learning device. In this theory, the GCs provide a sensorimotor context to the PCs, while the climbing fibers, the axons of the inferior olive (IO) neurons, carry the error signals that modify the GC–PC synapses in a supervised manner (see Fig. 1A). They further speculated that the GCs provide a sparse code, i.e. a code with only a small fraction (less than 10%) of cells active at any one time (see Fig. 1C). In this GC “expansion recoder” theory, supervised learning is facilitated because different MF inputs create highly dissimilar sparse GC activity patterns, which are easily recognizable by the PCs. The Marr and Albus supervised cerebellar learning models are supported by experiments showing that coincident activation of GC and climbing fiber input induces long-term depression (LTD) of the GC–PC synapses,²⁹ and that climbing fibers carry error signals.^{33,46} However, although the divergence from the MFs to the GCs and the existence of inhibitory feedback pathways are compatible with the transformation of dense MF codes (see data of van Kan *et al.*)⁵³ into GC sparse codes, it is not yet clear how GC sparse codes could actually be created.

LTD is only one of several types of synaptic plasticity in the cerebellum: recent experimental evidence points to

*Correspondence to: N. Schweighofer, Learning Curve, K.K., 2F Fuji Building 40, 15-14 Sakuragaoka-cho, Shibuya-ku, Tokyo 150-0031, Japan.

E-mail address: nicolas@neurotek.co.jp (N. Schweighofer).

Abbreviations: GC, granule cell; GO, Golgi cell; GT, gated threshold adaptation; GTA, gated threshold adaptation and gated anti-Hebbian learning; GTAH, gated threshold adaptation, gated anti-Hebbian and gated Hebbian learning; GTH, gated threshold adaptation and gated Hebbian learning; IO, inferior olive; LTD, long-term depression; MF, mossy fiber; MSE, mean square error; PC, Purkinje cell; T, threshold adaptation; TAH, non-gated threshold adaptation, anti-Hebbian and Hebbian learning.

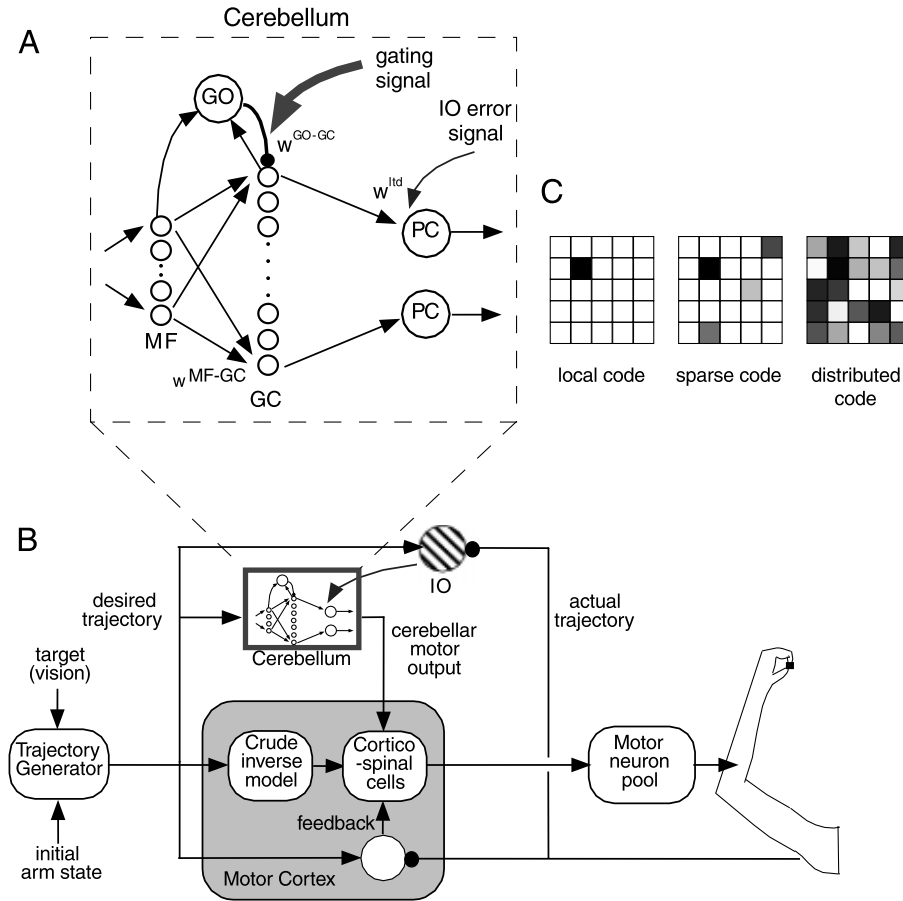


Fig. 1. The cerebellar model. (A) General, simplified circuitry of the cerebellum. Each circle represents a group of cells, the black dots inhibitory synapses and the arrows excitatory synapses. Three types of synapses are modifiable: the MF–GC excitatory synapses w^{MF-GC} are Hebbian, the GO–GC inhibitory synapses w^{GO-GC} are anti-Hebbian and the GC–PC synapses w^{td} are the LTD synapses. The direct MF–GO synapses and the GC–GO synapses are fixed. (B) Functional diagram of the (simplified) model for cerebellar control of arm movements. In this model, the cerebellum learns how to compensate for interaction torques that occur during reaching movements (when a multi-joint limb joint is accelerating, movement in one joint causes disruptive motion in all other joints). The visual system locates the target and an inverse kinematics model provides the trajectory generator with the target position in joint coordinates. Three subsystems generate motor commands. First, one area of the motor cortex generates a basic feedforward motor command from desired kinematic variables using only simple joint-wise computation. Second, another area of the motor cortex computes feedback motor synergy command by comparing desired and actual kinematic variables. Third, the intermediate cerebellum provides the corticospinal cells with motor commands compensating for the interaction torques acting on a joint when the other joint is moving. The three types of motor commands (feedforward, feedback and compensatory) are summed in corticospinal cells, which project to the spinal reaching system. The spinal cord transforms descending commands into individual muscle commands. IO cells compute torque-like errors in performance by subtracting sensory representations of actual trajectories from efference copies of desired movements. Because the cerebellar network is embedded in the control model (it resides in the box “intermediate cerebellum”), realistic cerebellar inputs can be derived. (C) Illustration of the different possible types of neural codes. The grids represent hypothetical arrays of granule cells, whose activities at some time t are shown by varying gray levels. In a local code, only one cell is active at any one time. In a sparse code there are only a small fraction of cells active at any one time. In a fully distributed code most of the cells are active at any one time.

several plastic processes in the granular layer. First, recent data show that mean neural firing rates can be set by activity-dependent regulation of neuronal electrical properties.^{1,18,51} In particular, potentiation of intrinsic excitability has been reported recently in cerebellar granule neurons.⁴ Second, D’Angelo *et al.*¹⁷ showed that MF high-frequency stimulation paired with GC membrane potential depolarization leads to potentiation of the MF–GC synaptic efficacies. Third, indirect evidence points to activity-dependent GO–GC synaptic plasticity: Zhu *et al.*⁵⁹ showed that GABA_A receptors on GCs alter their composition in response to excitatory afferent receptor stimulation.

Console-Bram *et al.*¹⁵ demonstrated that 43-kDa growth-associated protein mRNA, believed to be involved in experience-dependent plasticity, is modulated by GABA and glutamate in cultured GCs. Finally, as found in cortical neurons¹¹ and in the amygdala,⁵⁵ the serotonergic input from the raphe nucleus to the cerebellum⁵⁸ seems to play a role in cerebellar plasticity, as it modulates both cerebellar GC activity and the strength of GO inhibition *in vitro*.⁶ Furthermore, serotonin and norepinephrine have been shown to modulate GABA-mediated inhibitory transmission in the cerebellar cortex.³⁸

What could be the functional role of these plastic

processes? Here, we use a computational model to test the hypothesis that unsupervised learning can create “good” GC sparse codes that enhance cerebellar supervised learning. We argue that if (i) the mean GC firing is constrained to be low, (ii) the information transfer between MFs and GCs is maximized, (iii) the redundancy between GC discharges is minimized, and (iv) the resolution of the code is adaptive so that inputs corresponding to large errors in performance are finely encoded, then supervised learning in PCs is greatly improved. To create such GC sparse codes, we augment a Marr–Albus-type cerebellar model with a combination of three activity-dependent plastic processes that are all gated by a performance-dependent neuromodulator-like signal. These plastic processes are: GC threshold adaptation, Hebbian MF–GC synapses and anti-Hebbian GO–GC synapses. Our simulations show that integrating these biologically plausible unsupervised learning rules in a cerebellar model for arm movement control leads to accurate, fast and stable adaptive cerebellar control.

EXPERIMENTAL PROCEDURES

Theoretical background

According to the Marr and Albus theory, the GC layer pre-processes MF inputs for supervised learning in the output PC layer. Several authors have shown that information theory is helpful to analyze three-layered neural networks such as the cerebellar network (see, for instance, Linsker,³⁴ Attick,⁷ Plumbley,⁴¹ Field²² and Harpur²⁶). In particular, for efficient supervised learning in the output layer, information theory suggests that the neural code of the intermediate, pre-processing layer should have the following three properties. First, the code should be maximally informative of the input. Second, it should have low redundancy between neurons. Finally, it should be sparse.

The concept of information maximization is intuitively clear: the PCs should receive all the information they need to compute the desired output, i.e. there should be only minimal information loss between the MFs and the GCs. From our hypothetical three-layer cerebellar circuit, we can record a set of MF and GC activity patterns, given by the vectors \mathbf{MF} and \mathbf{GC} . We can then estimate the overall entropy $H(\mathbf{GC})$, which is a (positive) measure of the uncertainty, or diversity, conveyed by the \mathbf{GC} patterns [see Eq. (8)]. Likewise, the entropy $H(\mathbf{MF})$ can be estimated from the set of patterns \mathbf{MF} . The mutual information, defined as the difference in the uncertainty of the GC output without and with the observation of the MF input, is given by:

$$I(\mathbf{MF}, \mathbf{GC}) = H(\mathbf{GC}) - H(\mathbf{GC}|\mathbf{MF}). \quad (1)$$

If the randomness of GC activity is independent of the MF input, $H(\mathbf{GC}|\mathbf{MF})$ is constant, and maximizing $I(\mathbf{MF}, \mathbf{GC})$ is equivalent to maximizing $H(\mathbf{GC})$.

Redundancies, or correlations, between GCs lower the diversity of the firing patterns, i.e. the entropy $H(\mathbf{GC})$; this is intuitively satisfying because correlated neurons carry the same information. A measure of the redundancy is the mutual information between neurons given by:²⁶

$$I(GC_1, \dots, GC_N) = \sum_{k=1}^N H(GC_k) - H(\mathbf{GC}) \geq 0. \quad (2)$$

Thus, to minimize the redundancies, we should minimize the entropy carried by individual cell $H(GC_k)$. Because flat distributions of activities have maximal entropy and peaky distributions of activities minimal entropy, the entropy of an individual GC, $H(GC_k)$, can be reduced by enforcing a distribution of firing rates with a sharp peak. In other words, after recording the activity from all the cells during an experiment, and plotting a histogram

of the frequencies, the histogram should exhibit a single, sharp peak.

Whereas the sharp peak of the firing rate distribution of all the neurons could, in principle, be centered at any firing rate, the special case for which the probability distribution of firing peaks near 0 spikes/s corresponds to a sparse code. That is, the peak in the histogram of all the frequencies recorded should be near zero. The interest of sparse codes resides in that they provide desirable inputs for a subsequent neural layer.

For supervised learning in a subsequent neural layer, sparse codes combine the advantages of both local and distributed codes without their drawbacks.²² In a local code, only one cell is active at any time (see Fig. 1C); such a code leads to fast supervised learning, but requires many cells and generalizes poorly. In a fully distributed code, most of the cells are active at any one time (see Fig. 1C); such a code requires only few neurons and generalizes well, but leads to extremely slow learning. The best compromise between the two extremes is a sparse code, in which an appropriate percentage of cells is active at any one time. In Appendix A, we demonstrate that sparse GC codes allow large learning rates at the GC–PC synapses: this results in considerable speed-up of cerebellar learning.

In summary, from an information theoretic viewpoint, the GC code should simultaneously maximize the mutual information $I(\mathbf{MF}, \mathbf{GC})$ [Eq. (1)], i.e. increase $H(\mathbf{GC})$, minimize the redundancy [Eq. (2)], i.e. decrease $\sum_{k=1}^N H(GC_k)$, and enforce a low mean firing rate. Note that because $H(\mathbf{GC})$ appears in both Eqs. (1) and (2), maximizing the mutual information will also decrease the redundancies.

A limitation of information theory, however, is that it cannot distinguish useful from useless information. Because cerebellar LTD is the result of correlation between GC inputs and climbing fiber inputs, the value of the GC population is determined by its ability in “preparing” useful, goal-directed inputs to the PCs. In other words, the GC sparse code should depend on the performance errors (carried by the climbing fibers) of the subsystem in which the GCs are embedded. Consequently, the GC code should possess an adaptive resolution, so that the regions of the input space corresponding to large errors are finely encoded, thereby providing the PCs with fine inputs when they need to learn to compensate for large movement errors.

Algorithms for unsupervised learning of neural sparse codes

Learning algorithms derived from the above-mentioned theoretical arguments^{26,40} include Hebbian plasticity to maximize mutual information between the input and the output [Eq. (1)], anti-Hebbian learning to minimize redundancies [Eq. (2)], and a constraint on the firing rate to enforce sparse coding. Although these theoretical learning algorithms provide optimal solutions, they are not plausible biologically, however, because they assume symmetric input and feedback connections.

Other, less optimal but more plausible network models that address information maximization and redundancy reduction point to similar learning mechanisms. First, Hebbian learning between inputs and a layer of non-linear cells creates event detectors, i.e. each cell responds solely to a certain subset of input patterns.^{25,36} Second, anti-Hebbian synapses between non-linear cells can create independent output patterns.⁴ Third, an additional threshold adjusting rule that forces sparse coding in a network consisting of a layer of non-linear Hebbian units connected by anti-Hebbian feedback connections allows for high information transfer and low redundancy.²³ Note that these inhibitory connections are not plausible because they are to, and from, all the neurons of the layer. However, this does not have to be the case: more realistic architectures in which input neurons are connected to inhibitory interneurons with Hebbian and anti-Hebbian synapses achieved similar results.^{30,41}

Furthermore, as discussed above, sparse codes should also depend on performance errors. The adaptive topology encoding scheme²¹ implemented in the Cerebellar Model Arithmetic Computer³ aims both at using fewer encoding units and at increasing performance. This is achieved by moving the receptive fields of the input units into those parts of the input space where

Table 1. Parameter values for the four neuron models

	MFs	GCs	GOs	PCs
Time constant τ (ms)	5.0	5.0	5.0	10.0
Initial synaptic efficacies ranges	{-1; 0; 1}	MF \rightarrow GC [0.4, 0.6] GO \rightarrow GC [0.8, 1.2]	MF \rightarrow GO [0.1, 0.3] GC \rightarrow GO [0.01, 0.03]	[-0.01, 0.01]
Potential at which the activity is half of its maximum B	-1.0	120 (fixed model)	300 (fixed model)	0
Slope G	1	2.0	0.3	1
Max. firing (spikes/s)	100	100	100	Inf.

The cells interact via their activity, which depends on their membrane potential. The membrane potential follows a leaky integrator equation with time constant τ . Changes in membrane potential are brought about by synaptic inputs, whose initial efficacies are given in the second line of the table. Then, neuronal activities are computed by passing the membrane potential through sigmoid functions, whose slopes are given by G and by B , the potential at which the activity is half of its maximum.

they are really needed instead of distributing them randomly. Central to this scheme is a diffuse “gating” signal, the absolute value of the error in performance in Eldracher *et al.*’s scheme,²¹ that modulates the unsupervised learning rules.

Cerebellar unsupervised learning model

Functional and neural models. Testing the unsupervised learning of the GC sparse code in cerebellar adaptive control requires a task that is sufficiently difficult to learn but still computationally tractable. Learning to control a six-muscle two-joint arm has this proper balance.

Schweighofer *et al.*^{44,45} investigated the role of the cerebellum in adaptive movement control by inserting a detailed cerebellar model in a basic control system model (cerebral cortex and spinal network) that fed crude motor commands to a six-muscle two-joint arm. This model, based on feedback error learning,³¹ learned how to compensate for interaction torques occurring during visually guided reaching movements.⁹ Here, we insert a simplified version of the Schweighofer *et al.* cerebellar neural model, together with its sets of unsupervised learning rules (Fig. 1A), into the basic control system model (Fig. 1A, B). The inputs to the MF are kinematic variables, the IO carries feedback errors and, after learning the inverse dynamics of the arm, the outputs of the cerebellar neural network are the interaction torques necessary for precise control. The cerebellar neural circuit, which comprises 20 MFs, 100 GCs, one GO and two PCs, is described in more detail in Appendix B, and the parameters of the different neurons are given in Table 1.

Learning model. We propose here a novel unsupervised learning cerebellar model. The model learns GC sparse codes and is compatible with cerebellar anatomy, physiology and the new findings in cerebellar cortex plasticity.

First, a homeostatic mechanism, in the form of an activity-dependent threshold adaptation rule, maintains appropriate low mean GC firing rate compatible with a sparse code (as well as the mean GO activity).

Second, the MF–GC synaptic efficacies vary according to Hebbian rules (see Chauvet).¹⁴ Hebbian learning at the MF–GC synapses creates event detectors, i.e. the GCs respond more often to previously seen patterns than to unknown patterns.

Third, inhibitory GO–GC synaptic efficacies vary according to anti-Hebbian learning rules. This results in an increase in diversity (lower redundancy) of the GC firing rates. Note that, because they receive inputs from numerous GCs, the GOs compute an approximation of the local average of the GC activity. Because anti-Hebbian learning at the GO–GC synapses decreases the correlation between each GC response and the GO response, GC responses are driven away from their average, increasing diversity in firing rates.

Fourth, we introduce a diffuse, neuromodulator-like, gating signal that allows threshold adaptation, Hebbian and anti-Hebbian learning during periods of large performance errors and blocking it during periods of low errors (or inaction). We see next how these rules are actually implemented.

The gated activity-dependent adaptation rule modifies the potential value for which the firing rate is half of its maximum value (which we call “threshold” for convenience) B_i^{gc} according to:

$$\Delta B_i^{gc} = \beta g \left(\frac{GC_i}{GC^d} - 1 \right), \quad (3)$$

where β is a learning rate, GC_i the firing rate of the i th GC [see Eq. (B2) in Appendix B] and GC^d is the desired (low) average GC firing rate. The global (as opposed to local to each neuron) “gating” term g , which enables threshold adaptation, is a function of the error in performance [see Eq. (B6) in Appendix B].

We model Hebbian learning with the following gated version of Oja’s Hebbian rule:³⁹

$$\Delta w_{ij}^{MF-GC} = \alpha_1 g (MF_i gc_j - w_{ij}^{MF-GC} gc_j^2), \quad (4)$$

where MF_i is the firing rate of the i th MF, gc_j the membrane potential of the j th GC [see Eq. (B2) in Appendix B] and $\alpha_1 = 3 \times 10^{-8}$ a learning rate. The activity-dependent decay term $-w_{ij}^{MF-GC} gc_j^2$ is necessary because of the inherent instability of the Hebbian rule. Moreover, we added the constraints $w_{ij}^{MF-GC} \geq 0$ so that the MF–GC synapses are always excitatory.

The GO–GC inhibitory synaptic efficacies vary according to a gated anti-Hebbian rule:

$$\Delta w_j^{GO-GC} = \alpha_2 g (GO gc_j - \alpha_3 w_j^{GO-GC} gc_j^2), \quad (5)$$

where g is the gating term defined above, $\alpha_2 = 5 \times 10^{-7}$ is a learning rate, $\alpha_3 = 0.01$ is a decay rate, GO the firing rate of the Golgi cell (although the model could be extended, there is only one GO in this present model; see below). Note that Eq. (5) is an anti-Hebbian rule because the strength of each GO–GC inhibitory synapse is increased in proportion to pre- and postsynaptic activity. Whereas anti-Hebbian rules alone are inherently stable (i.e. no decay term is normally needed), the addition of the threshold adaptation in the present model makes it run away; the decay term is thus necessary. Moreover, we impose the constraint $w_j^{GO-GC} \geq 0$ so that the GO–GC synapses are always inhibitory [as shown by the minus sign in Eq. (B2)].

Simulations

Motor task: learning fast and accurate reaching movements. We simulated fast planar reaching movements (duration 300 ms) to eight targets situated on a circle of radius 20 cm around the start position. The target to the right of the start position was defined as the 0° direction and targets at successive counterclockwise positions were given in 45° increments. At $t < 0$, the hand was on the 90° target. At $t = 0$ s, the target was first moved to the central position. At $t = 1$ s (i.e. 0.7 s after the end of the desired movement), the target was displaced to a new position, and so on. The straight, desired trajectories were generated by a minimum-jerk trajectory generator in extra-personal space and then converted to joint angles by an inverse kinematics transformation (see Schweighofer *et al.*).⁴⁴

Table 2. Comparison of mean square error for the different models

	Fix	GT	GTH	GTA	GTAH
MSE (cm ²)	2.9 ± 0.4	1.5 ± 0.3	1.1 ± 0.05	0.9 ± 0.2	0.6 ± 0.1

Fix, fixed model. Because we are interested in the fast learning capabilities of the model, the results shown here are not asymptotical errors, but errors after 100 (supervised) learning trials.

We ran several specific simulations to selectively test the different adaptation rules. First, as a benchmark, we examined the performance of a “fixed” model, i.e. a model with static GC thresholds and synaptic efficacies that is simply trained by supervised learning in the PCs [Eq. (B7) in Appendix B]. Second, we included the gated threshold adapting (GT) rule [Eq. (3)] for both the GCs and the GO to allow each cell to be responsive to its inputs. Third, to study the effect of MF–GC long-term potentiation on the one hand and GO–GC synaptic plasticity on the other, we looked at the additional effect of the gated Hebbian (GTH) rule [Eqs. (3) and (4)] and the gated anti-Hebbian learning (GTA) rule [Eqs. (3) and (5)]. Fourth, we assessed the computational power of the fully learned model, i.e. a model comprising the gated threshold adaptation rule, the gated anti-Hebbian learning rule and the gated Hebbian learning rule (GTAH model). Finally, we analyzed the role of putative neuromodulators in gating the adaptation rules by running two non-gated models: a model with only threshold adaptation (T model), and a model with non-gated threshold adaptation, anti-Hebbian and Hebbian learning (TAH model).

For every set of learning rules, simulations were run 10 times, each time with different initial conditions for the random number generator. The time step was 3 ms.

Data analysis. To assess performance, we computed the mean square error (MSE) in trajectory during a test trial in which every target was visited once, sequentially. Thus, there were 16 movements in a test trial: eight movements to the targets and eight movements back to the initial position. The MSE is defined by:

$$\text{MSE} = \frac{1}{M} \sum_{t=1}^M [x^d(t) - x(t)^r]^2 + [y^d(t) - y(t)^r]^2, \quad (6)$$

where M is the number of time steps, the superscript “d” refers to the desired trajectory and the superscript “r” to the real trajectory.

To compute the entropies of individual cells and the redundancy [see Eq. (7)], we binarized the GC firing rates (note that this binarization was only performed for computational tractability of the statistical analysis, but it was not used for the on-line cerebellar model; see Appendix B). Preliminary simulations showed that when the GCs fired, they often did it close to GC^{\max} (100 spikes/s). As a consequence, we set the threshold for binarization at 90 spikes/s. The normalized GC redundancy R is computed from the mutual information between GCs [Eq. (2)], and is defined by:

$$R = \frac{1}{C} I(GC_1, \dots, GC_N) = \frac{1}{C} \left(\sum_{k=1}^N H(GC_k) - H(\mathbf{GC}) \right), \quad (7)$$

where N is the number of GCs, and $C = N \log(2)$ the capacity of the binarized GC network. The overall entropy and the sum of each cell’s entropy are, respectively:

$$H(\mathbf{GC}) = - \sum_{j=1}^M p(\mathbf{GC}^j) \log(p(\mathbf{GC}^j)) \quad (8)$$

and

$$\sum_{k=1}^N H(GC_k) = - \sum_{k=1}^N [p_k \log p_k - (1 - p_k) \log(1 - p_k)], \quad (9)$$

where $p(\mathbf{GC}^j)$ is the probability of pattern \mathbf{GC}^j and p_k the firing probability of cell GC_k .

To assess the information content of the GC discharges, we computed the mutual information between individual GCs and movement variables, which comprised both kinematic input variables (position, velocity, acceleration) and motor output variables (inertial torques, interaction inertial torques, centripetal torques and, for the shoulder only, Coriolis torques). For each computation of the mutual information, both the GC firing rate and the input variable (one of 13 variables: six kinematic and seven dynamic) were binned in 100 bins, thus creating 100 levels s_i and n_j . The mutual information is:

$$I = \sum_{i,j=1}^{100} P(s_i, n_j) \log_2 \left(\frac{P(s_i, n_j)}{P(s_i)P(n_j)} \right). \quad (10)$$

We then computed the normalized mutual information I/S , where S is the entropy of the input (for no information loss the ratio is 1).

RESULTS

Simulations showed that learning performance for the fixed model, as measured by the mean MSE, was quite sensitive to the average GC firing value; best learning occurred for $\overline{GC} \sim 5$ spikes/s. However, although we carefully chose the mean GC and GO threshold values [see Eqs. (B2) and (B3) in Appendix B] so that the smallest possible MSE was obtained, the mean MSE was still large (see Table 2) after 100 training movements: the MSE was only reduced to an average of 2.9 cm², i.e. about 20% of the initial error (15.7 cm²). Moreover, 16.2 ± 3.8% of the GCs did not fire during the 16 test movements.

In agreement with the results obtained in the fixed model, in the GT model [Eq. (3)] the desired mean GC firing rate was set to $GC^d = 5$ spikes/s and the desired mean GO firing rate was set at $GO^d = 30$ spikes/s. We ran the learning experiment in two parts. First, we ran the unsupervised learning experiment (100 reaching movements to targets randomly chosen among the set of eight; same movement specifications as above) to adjust the individual GC and GO thresholds. Second, we ran the supervised learning experiment (100 movements to randomly chosen targets), i.e. the GC–PC synaptic efficacies were varied [Eq. (B7) in Appendix B]. Because adjusting the thresholds created a crude sparse code, learning could be speeded up by increasing the learning rate in the GC–PC synapses from $\eta = 5 \times 10^{-7}$ for the fixed case to $\eta = 5 \times 10^{-6}$ (see Appendix A). The MSE after 100 (supervised) learning trials was about twice as small as that of the fixed model (see Table 2). Finally, we verified that all the GCs fired at some time during the series of movements.

To train the GTA and GTH models, we first let the model perform 300 movements to randomly chosen targets (same movement specifications as above) to adjust the thresholds and the synaptic efficacies of the GC layer (during the first 100 movements only the thresholds were adjusted so that all the GCs were in their working ranges). Second, as above, we ran the supervised learning experiment (100 movements to random targets). As seen in Table 2, inclusion of the GTH rules resulted in a decrease of the mean MSE (compared to the gated threshold model) and a large decrease in inter-trial variability ($\eta = 5 \times 10^{-6}$). The GTA had an even stronger effect on the mean MSE.

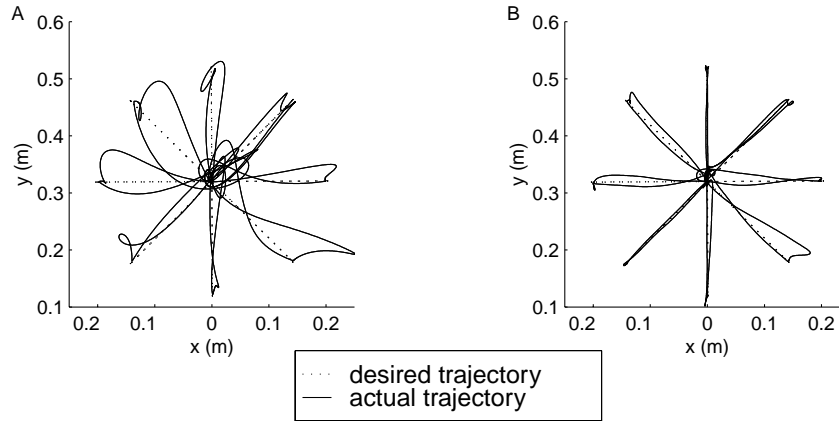


Fig. 2. Performance of the network in 300-ms reaching movements to targets at a distance 20 cm from the rest position. The dotted lines represent the desired trajectories and the plain lines the actual trajectories. (A) Trajectories before learning. The error is $MSE = 15.7 \text{ cm}^2$. (B) Trajectories generated by the trained GTAHA model, after 100 supervised learning movements. The error in this case is $MSE = 0.57 \text{ cm}^2$ and the actual movements closely follow the desired movements.

Finally, the GTAHA model gave a mean MSE that was less than a fifth of the mean MSE produced by the fixed model and about a third of the MSE produced by the GT model. Moreover, there was a large decrease in inter-trial variability.

The trajectories generated by the cerebral cortex–spinal cord controller before cerebellar learning and the trajectories generated by the GTAHA model are shown in Fig. 2. Before learning, large interaction torques occurring in fast 300-ms movements caused the actual trajectories to largely deviate from the desired trajectories (Fig. 2A). After learning (same conditions as above), the GTAHA model resulted in greatly improved motor performance and the actual trajectories precisely followed the desired trajectory (Fig. 2B).

Comparing the learning curves for a GTAHA model and for a fixed model (Fig. 3) reveals that the Hebbian and anti-Hebbian rules were not so much necessary for the

first rapid and large decrease in MSE, but for the smaller refinements in movement performance crucial for accurate reaching. Because the Hebbian rule leads to a better recoding of the inputs (see below) and the anti-Hebbian rule to a decrease in redundancy between GCs (see below), the PCs can better learn to compensate for the non-linear component of the interaction torques.

Cerebellar neuron activities during three movements for the fully trained GTAHA model are shown in Fig. 4. In Fig. 4A, we show the firing rates of three MFs responding to a combination of position, velocity and acceleration. In Fig. 4B, we show the GC raster plot. The cells fired mostly during the movements. This is consistent with experiments showing that the GCs have no intrinsic firing at rest.¹⁶ Note, further, that no cell fired for more than a fraction of a movement time; moreover, only a small percentage of cells discharged at any time. In contrast, the GO was continuously activated during each




Fig. 3. Example of learning curves for the fixed model and the GTAHA model. Note how the gated Hebbian and anti-Hebbian rules play a non-negligible role in the refinement motor performance.

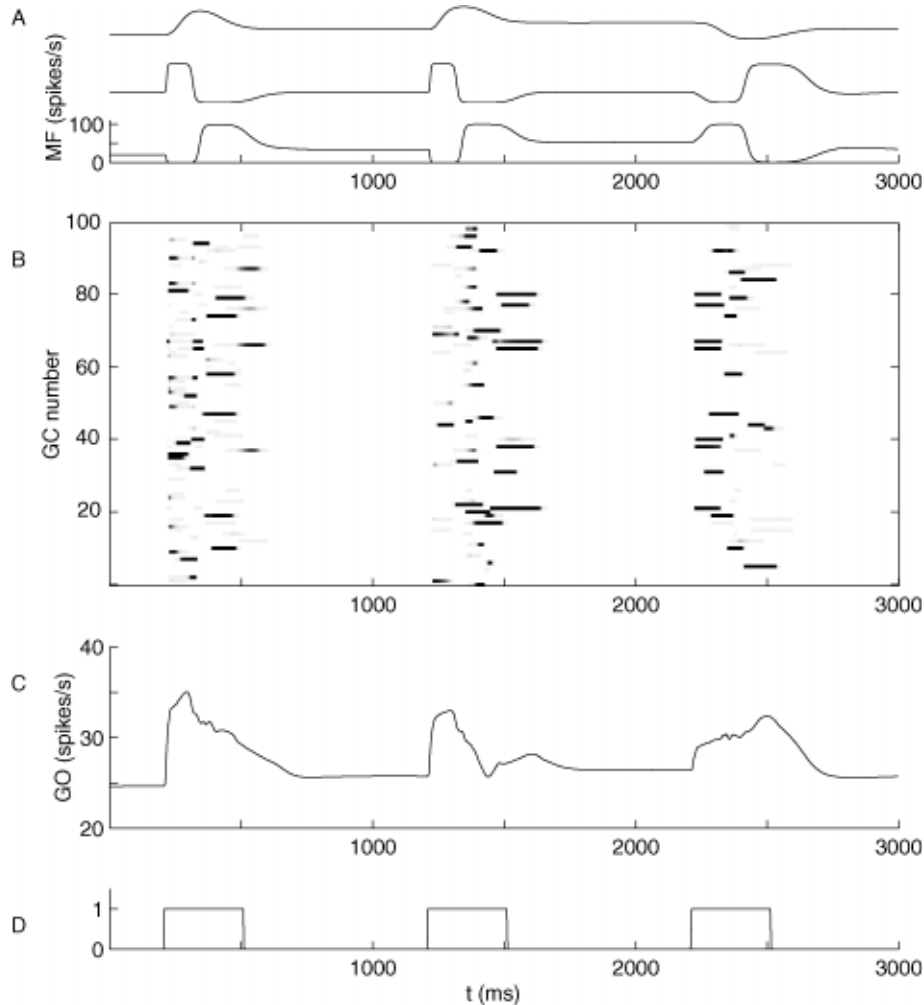


Fig. 4. Activities of the cerebellar neurons during the first three movements after learning for the GTA model for 300 trials. (A) Firing rates of three MFs that respond to a combination of position, velocity and acceleration. (B) Raster plot of the GCs shows that the cells fire mostly during the movements, and that only a small percentage fire at any instant. The black traces correspond to 100 spikes/s. (C) Response of the GO. (D) The sequence of action (1) and inaction (0). Movements occur when the steps are at 1.

movement, as shown in Fig. 4C. Note that a histogram of the firing rates of all the cells recorded during the 16 movements showed a sharp peak at the low desired frequency, 5 spikes/s. However, a histogram of the firing rates recorded during the whole experiments (i.e. during movements + pause between movements) shows a peak closer to 0 spikes/s (results not shown).

A comparison of the activities of five GCs for the GT and GTA models (same initial conditions) in the course of one movement (from the target at -45° to the central position) after learning is shown in Fig. 5. This figure illustrates the role of the Hebbian and anti-Hebbian learning rules in information transfer maximization and redundancy reduction. Although the cells showing activity during this movement did not all fire simultaneously, cells 4 and 5 both discharged around 50 ms in the GT model. Because too many cells fired at this instant (i.e. the neural code was redundant), the anti-Hebbian learning at the GO–GC synapses suppressed GC activity of cell 4 (also seen for cell 1 around 100 ms). In contrast, if too few cells fired at some instant (i.e. if the information transfer was poor), their activities were

reinforced by Hebbian learning at the MF–GC synapses (cells 2 and 3).

We then examined a possible causal relationship between GC redundancy and performance error by plotting the MSE as a function of the redundancy [Eq. (7)] for the different models discussed above, for different initial conditions (same set of initial conditions for all models). First note that the redundancies were largely reduced in the GT model compared to the fixed model. Second, the mean redundancy between cells of the GTA model was about half that of the redundancy of the fixed model, indicating that even a single GO could nicely decorrelate the GC activities. Third, the redundancy for the GTH model was lower than that of the GT model: because Hebbian learning increased the input/output mutual information (see below) by increasing $H(\mathbf{GC})$ [see Eq. (1)], it also decreased the redundancies [see Eq. (2)]. A strong linear trend between the two variables ($y = 19.1x - 2.67$; $P < 0.0001$ and $R^2 = 0.84$) was observed, as shown in Fig. 6. Thus, GC redundancy reduction was crucial in order to achieve adequate cerebellar motor performance.




Fig. 5. Comparison of five GC activities during one movement for the GT model after 100 movements (plain lines) and with the GTA model (dashed lines) after 300 movements. The results are comparable because the GT model has converged after 100 movements. The bottom panel shows the time of movement.




Fig. 6. Linear regression between the redundancy and the MSE shows a strong linear trend. The data points are obtained with different initial conditions and different models. Fix, fixed model. The redundancy is a dimensionless quantity.




Fig. 7. Mean mutual information between GC firing and movement variables recorded after learning for the GT model and the GTA model. There are six kinematic variables (position, velocity and acceleration for each joint) and seven dynamic variables (for each joint: inertial torques, interaction inertial torques, centripetal torques and, for the shoulder only, Coriolis torques). The thin bars show the standard deviations. Note that the mutual information is larger for all variables for the GTA model than for the GT model.

What is the information contained in the GC discharges? Is this information transformed qualitatively or quantitatively by the unsupervised learning rules? To answer these questions, we computed the mean mutual information [normalized by the entropy of the input; see Eq. (10)] between each GC and six kinematic variables and seven motor command variables (see Experimental Procedures) during a test trial of 16 movements, after learning for both the GT and the GTA models. The results, shown in Fig. 7, allow some general remarks. First, the mutual information contained by each cell was small for all the variables, as expected with a sparse coding scheme (when the cells did not fire, they carried no information). Second, although the mean firing rate was identical, there was an overall increase in mutual information for all the kinematic variables in the GTA model compared to the GT model. In particular, the GCs of the GTA model have a much higher information content in acceleration (more than 60% for both joints) compared to the GT model. Third, overall, the dynamic variables were better encoded than the kinematic variables; moreover, the GCs of the GTA model encode the centripetal and Coriolis torques best of all variables. Thus, interestingly, the gated Hebbian and anti-Hebbian learning rules not only increased the information transfer, they also helped the GCs to compute the non-linear interaction torques before the PC layer.

The role of the gating process in maintaining stable learning can be well appreciated when long periods of inactivity separated each movement. The gate was

crucial when the interval between movements was more than 1.0 s (Fig. 8). The gated models (bottom panel) gave good, stable performance even when the periods of inaction were prolonged. In contrast, when gating was not included (top panel), the unsupervised learning algorithms diverged even for small intervals between movements (note the difference in scale in the ordinate of both plots): for intervals greater than 1 s, the MSE was about the same as before any cerebellar learning (see Fig. 2A). Also note that the GTA model surpassed the GT model for all intervals.

DISCUSSION

We proposed a novel model of unsupervised learning in the cerebellum compatible with cerebellar anatomy and physiology that incorporates new findings in cerebellar plasticity. We showed that gated unsupervised learning rules construct stable, goal-directed sparse GC codes that are informative of the input and have minimal redundancies between neurons. Moreover, we demonstrated that these sparse codes allow larger learning rates at the GC-PC synapses, which speed up cerebellar learning considerably. Hence, we suggest that, for accurate and fast cerebellar learning, the GCs may be much more plastic than envisioned by Marr and Albus.

Comparison with other cerebellar models

To show the capabilities of the GC sparse code in cerebellar learning, Tyrrell and Willshaw⁵² simulated

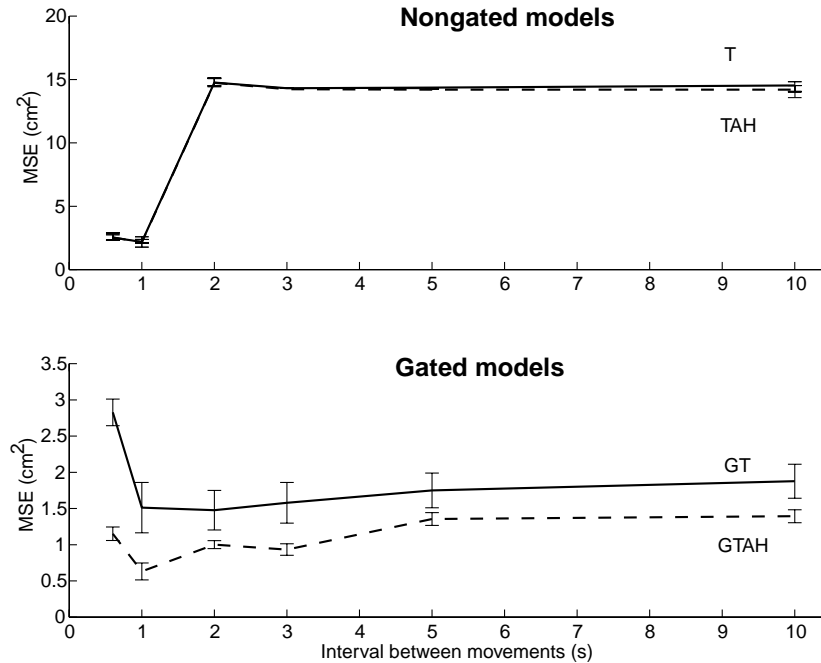


Fig. 8. Variation of the relative error on the interval between movements. Top: the T model (plain line) and TAH model (dashed line). Note that the two curves almost overlap; moreover, for intervals greater than 1 s, the MSE is about the same as before any cerebellar learning. Bottom: GT and (plain line) and GTAHA (dashed line). Note the different scales between the ordinates. The gate becomes crucial as soon as the interval between movements is more than 1.0 s; the non-gated models diverge for higher values. On the contrary, the gated models produce stable and accurate movements even when the periods of inaction are prolonged. Also, the GTAHA largely surpasses the GT.

Marr's model. Because of the binary GCs and synaptic efficacies, these authors found a low capacity, as only about 100 different patterns could be learned by a single Purkinje cell. The analog Albus² model is bound to have a much higher capacity; however, although Albus claimed to have ran sparse coded GC simulations, no actual results were presented.

The Kettner *et al.*,³² Schweighofer *et al.*⁴⁵ and Barto *et al.*⁸ models indicated that GC sparse codes were effective in learning to control dynamical systems. In the Kettner *et al.* and Barto *et al.* models, however, the GCs were binary and the sparse representation was artificially implemented. In the Schweighofer *et al.* model, the GC parameters (thresholds and synaptic efficacies) had to be carefully hand-tuned. Moreover, because these parameters were static, the GC firing ranges were not equally used and many GCs often fired simultaneously or not at all; these limitations resulted in poor and slow learning (presumably because of both poor information transfer and high redundancies). In contrast, in the present model, a GC sparse code that leads to good motor performance is automatically learned.

The Buonomano and Mauk model¹² showed that the GO negative feedback could create time-varying, non-periodic, subsets of active GCs in response to periodic inputs. On one hand, because their model uses spiking GCs, it goes farther than the present model; on the other hand, the task to be learned (a series of time intervals) was much simpler than the motor task described here.

Note finally that the present cerebellar unsupervised learning model does not critically depend upon the choice of the control model; it could therefore also be

tested in other models, such as the model of Barto *et al.*,⁸ which does not make use of desired trajectories.

Unsupervised learning in the granule cell layer

Our results show that, if appropriately gated, the MF–GC Hebbian and the GO–GC anti-Hebbian plastic processes can improve motor performance dramatically. Because the MF–GC Hebbian synapses create event detectors that are decorrelated by the GO–GC anti-Hebbian synapses, the dense and correlated MF inputs are projected onto highly dissimilar sparse GC patterns easily recognizable by the PCs. Furthermore, Hebbian learning allows a good representation of the input space, independently of the initial conditions, leading to a small inter-trial variability. Note that, in spite of the cellular “specialization” that is created by the Hebbian and anti-Hebbian rules, the responses to new inputs are not degraded and the model can still generalize gracefully. We tested the generalization properties of the GTAHA model compared to those of the GT model by training the network as in Fig. 2B, but tested with targets rotated 22.5°. In spite of the neurons' specialization, we observed no relative degradation of the generalization performance of the GTAHA model compared to the GT model (results not shown).

In the GTAHA model, computing the mutual information between movement variables and the GC activity revealed that the centripetal and Coriolis torques were represented best of all variables in the GC code. These are precisely the variables (with the inertial interaction torques) that need to be computed by the cerebellum to

achieve good performance in the Schweighofer *et al.* model. Thus, our results suggest that, if the GCs are equipped with gated Hebbian and anti-Hebbian learning rules, the adaptive, non-linear, summation of four MF and one GO inputs is an efficient pre-processing computational step for supervised learning in the PCs.

Diffuse gating signals

In the present model, the diffuse gating signal that turns off the unsupervised learning rules during inaction is crucial for cerebellar adaptive control: when no gate is included, learning is severely degraded if the arm is at rest during long periods (see Fig. 8, top). In this case, the GCs eventually learn how to respond best to position inputs (since during periods of inaction only position is represented in the MF input; see van Kan *et al.*).⁵³ Then, when a movement is generated, the GCs have “forgotten” how to code for velocity and acceleration. Because acquisition of the inverse dynamic models requires position, velocity and acceleration inputs, supervised learning in the PC is extremely poor when movements are alternated with long periods of inaction.

Because the diffuse gating signal that regulates the unsupervised learning rules is a function of the error in performance, our gated unsupervised learning algorithm may seem to be a supervised learning algorithm. However, two important features distinguish it from a supervised learning algorithm. First, it is only the magnitude of the overall performance error that is taken into account; the direction of the error, which is crucial in supervised learning, is not used. Second, while a vector carrying joint-specific error signals is needed for supervised learning, the gating signal is a function of the overall error in performance. Thus, our gating signal is reminiscent of the reward signal in reinforcement learning;⁴⁸ unlike the reward signal in reinforcement learning, however, our gating signal is only positive or null.

Biological plausibility

Although the neurons in the model are very simplified, they share similar behavior with real neurons. The current frequency curve of real GC neurons exhibits linearity until about 100 Hz.¹⁶ In the model, the firing rate GC_i is obtained by passing the membrane potential through a sigmoid that approximates the current–frequency curve of a realistic GC model.²⁴ The input gain and the large maximum firing value $GC^{\max} = 100$ spikes/s have been chosen to reproduce the linearity of the granule cells for a large range of input currents. Furthermore, as shown in Figs 4 and 5, the GCs fired mostly during the movements. This is consistent with experiments showing that the GCs have no intrinsic firing at rest.¹⁶ Furthermore, the threshold of the GO is adjusted so that the cell has a background firing rate of about 20 spikes/s, in accordance with experimental results.⁵³

The learning algorithms proposed are plausible and well supported by existing data. Our threshold adaptation rule is compatible with the potentiation of intrinsic excitability recently shown in cerebellar granule neurons.⁵

Armano *et al.* showed that high-frequency MF stimulation decreased GC spike threshold and that this change in excitability depends on the membrane potential. In the present model, we propose that the modulation of intrinsic excitability is a regulatory process, i.e. a GC can be either potentiated or depressed, depending on the mean firing rate of the cell compared to a “desired” firing rate.

The Hebbian rule is compatible with the experimental induction of long-term potentiation in the real granule cells.¹⁷ In these experiments, like in the model, long-term potentiation was induced by MF stimulation [MF_i in Eq. (4)] paired with GC depolarization [gc_j in Eq. (4)].

Unfortunately, no direct proof of the existence of GO–GC activity-dependent plasticity exists to date; however, the proposed learning rule is plausible and can be tested by GO stimulation paired with GC depolarization.

There are two neuromodulator candidates for the gating signal. First, serotonin levels have been shown to be low during sleep compared to waking in both the raphe nucleus⁴² and in the cerebellum¹³ (however, evidence linking serotonin levels to cerebellar deficits are scant, but see Trouillas).⁵⁰ Second, cerebellar norepinephrine has been shown to be involved in the acquisition of new motor tasks.⁵⁷

Although the fully plastic GTAH model could accurately learn to control a two-dimensional arm with as little as 100 neurons, there are of the order of 10^{11} GCs in the human cerebellum. The following facts allow us to put our model in perspective with the anatomy of the real cerebellum. First, it should be noted that because neurons are noisy devices, each one of our idealized neurons corresponds in fact to a small GC population. Second, given the large number of degrees of freedom and the pervasive non-linearities of the seven degrees of freedom human arm, an internal model of the arm’s dynamics is an extremely complex mapping between kinematic and dynamic variables and thus requires a large number of encoding GCs. Third, many more GCs would be required to learn to control the arm over the whole movement space. To actually test if a few GOs can actually decorrelate $>100,000$ GC noisy firing patterns, a detailed cerebellar model composed of a large number of spiking neurons,³⁵ based on the actual physiology of the GOs,^{19,54} could be embedded in a realistic noisy sensorimotor system.

Proposed experimental studies

We now propose four experiments to test the validity of our model.

Low mean granule cell firing rates and granule cell information content. Although individual GCs are too small to be recorded with electrodes *in vivo*, the recent two-photon microscopy technique could be used to measure single GC activity *in vivo*.⁴⁹ The very high temporal and spatial resolution of two-photon microscopy would be adequate to measure calcium dynamics in GCs, which appears to correlate well with neural activity.²⁴ Note, however, that two-photon microscopy may not be powerful enough to measure into the granular

layer, which lies more than 300 μm below the cerebellar surface.

Anti-Hebbian learning. *In vitro* experiments probing the plasticity of the GO–GC synapses by stimulating the pre- and postsynaptic sites simultaneously could be designed to test our prediction of anti-Hebbian GO–GC synapses.

Gated unsupervised learning. Our prediction of the neuromodulatory gating process could be tested *in vitro* by probing the speed of induction of long-term potentiation at the MF–GC synapses under application of neuromodulators or their agonists and antagonists at different concentrations.

Granule cell redundancies affect motor performance. This prediction could be tested experimentally by monitoring both the GC activity and the performance error in the behaving animal in both intact and mutant mice lacking GOs.⁵⁶ Because many GCs need to be simultaneously and independently recorded to compute the redundancies, optical recording would be the appropriate

experimental tool. Although optical recording of GC activity has so far only been performed in anesthetized animals (for a review, see Ebner and Chen),²⁰ recent advances that allow recording in behaving animals²⁸ might confirm the relationship between redundancies among GCs and motor performance.

CONCLUSION

Because the present model spans from the system to the cellular level, we could demonstrate by simulation the effect of the GC neural code onto adaptive motor control. Although such a direct, causal, effect is difficult to show experimentally, a series of experiments such as those proposed above could further support the hypothesis that learning goal-directed sparse GC codes that are informative of the input and have minimal redundancies between neuron codes is crucial for fast, precise and stable cerebellar learning.

Acknowledgements—We are grateful to the members of the Kawato Dynamic Brain Project at ATR for their comments on earlier drafts of this paper.

REFERENCES

1. Aizenmann C. and Linden D. J. (2000) Rapid, synaptically driven increases in the intrinsic excitability of cerebellar nuclear neurons. *Nat. Neurosci.* **3**, 109–111.
2. Albus J. S. (1971) The theory of cerebellar function. *Math. Biosci.* **10**, 25–61.
3. Albus J. S. (1981) *Brains, Behaviours and Robotics*, Lawrence Erlbaum, Mahwah, New Jersey.
4. Amari S. and Takeuchi A. (1978) Mathematical theory on formation of category detecting nerve cells. *Biol. Cybern.* **29**, 127–136.
5. Armano S., Rossi P., Taglietti V. and D'Angelo E. (2000) Long-term potentiation of intrinsic excitability at the mossy fiber–granule cell synapse of rat cerebellum. *J. Neurosci.* **20**, 5208–5216.
6. Armstrong D. L., Hay M. and Terrian D. M. (1987) Modulation of cerebellar granule cell activity by iontophoretic application of serotonergic agents. *Brain Res. Bull.* **19**, 699–704.
7. Attkick J. J. (1992) Could information theory provide an ecological theory of sensory processing? *Network* **3**, 213–251.
8. Barto A. G., Fagg A. H., Sitkoff N. and Houk J. C. (1999) A cerebellar model of timing and prediction in the control of reaching. *Neural Comput.* **11**, 565–594.
9. Bastian A. J., Martin T. A., Keating J. G. and Thach W. T. (1996) Cerebellar ataxia: abnormal control of interaction torques across multiple joints. *J. Neurophysiol.* **76**, 492–509.
10. Bishop C. (1995) *Neural Networks for Pattern Recognition*, Oxford University Press, Oxford.
11. Brocher S., Artola A. and Singer W. (1992) Agonists of cholinergic and noradrenergic receptors facilitate synergistically the induction of long-term potentiation in slices of rat visual cortex. *Brain Res.* **573**, 27–36.
12. Buonomano D. V. and Mauk M. D. (1994) Neural network model of the cerebellum: temporal discrimination and the timing of motor responses. *Neural Comput.* **6**, 38–55.
13. Cespuglio R., Faradji H. and Jouvet M. (1983) Voltammetric detection of extracellular 5-hydroxyindole compounds at the level of cell bodies and the terminals of the raphe system: variations during the wake–sleep cycle in the rat in chronic experiments. *C. r. hebdomadaire Séances Acad. Sci., Paris* **296**, 611–616.
14. Chauvet G. (1986) Habituation rules for a theory of the cerebellar cortex. *Biol. Cybern.* **55**, 201–209.
15. Console-Bram L. M., Baird D. H., Fitzpatrick-McElligott S. G. and McElligott J. G. (1998) Modulation of GAP-43 mRNA by GABA and glutamate in cultured cerebellar granule cells. *Brain Res.* **783**, 316–325.
16. D'Angelo E., De Fillipi G., Rossi P. and Taglietti V. (1998) Ionic mechanism of electroresponsiveness in cerebellar granule cells implicates the action of a persistent sodium current. *J. Neurophysiol.* **80**, 493–503.
17. D'Angelo E., Rossi P., Armano S. and Taglietti V. (1999) Evidence for NMDA and mGlu receptor-dependent long-term potentiation of mossy fiber–granule cell transmission in rat cerebellum. *J. Neurophysiol.* **81**, 277–287.
18. Desai N. S., Rutherford L. C. and Turrigiano G. G. (1999) Plasticity in the intrinsic excitability of cortical pyramidal neurons. *Nat. Neurosci.* **2**, 515–520.
19. Dieudonne S. (1998) Submillisecond kinetics and low efficacy of parallel fibre–Golgi cell synaptic currents in the rat cerebellum. *J. Physiol.* **510**, 845–866.
20. Ebner T. J. and Chen G. (1995) Use of voltage-sensitive dyes and optical recordings in the central nervous system. *Prog. Neurobiol.* **46**, 463–506.
21. Eldracher M., Staller A. and Pompl R. (1997) Adaptive encoding strongly improves function approximation with CMAC. *Neural Comput.* **9**, 403–417.
22. Field D. J. (1994) What is the goal of sensory coding? *Neural Comput.* **6**, 559–601.
23. Foldiak P. (1990) Forming sparse representations by local anti-Hebbian learning. *Biol. Cybern.* **64**, 165–170.
24. Gabbiani F., Midgaard J. and Knopfel T. (1994) Synaptic integration in a model of cerebellar granule cells. *J. Neurophysiol.* **72**, 999–1009.
25. Grossberg S. (1976) On the development of feature detectors in the visual cortex with applications to learning and reaction–diffusion systems. *Biol. Cybern.* **21**, 145–159.
26. Harpur G. F. (1997) *Low Entropy Coding with Unsupervised Neural Networks*, Queen's College, University of Cambridge.

27. Hertz J., Krogh A. and Palmer R. (1991) *Introduction to the Theory of Neural Computation*, Addison-Wesley, Redwood City, CA.
28. Inase M., Iijima T., Takahashi I., Takahashi T., Shinoda M., Hirose H., Niisato L. and Tsukada K. (1998) Optical recording of the motor cortical activity during reaching movements. *Soc. Neurosci. Abstr.* **24**.
29. Ito M., Sakurai M. and Tongroach P. (1982) Climbing fiber induced long term depression of both mossy fiber responsiveness and glutamate sensitivity of cerebellar Purkinje cells. *J. Physiol.* **324**, 113–134.
30. Jonker H. J., Coolen A. C. and Denier van der Gon J. J. (1998) Autonomous development of decorrelation filters in neural networks with recurrent inhibition. *Network* **9**, 345–362.
31. Kawato M. and Gomi H. (1991) A computational model of four regions of the cerebellum based on feedback-error-learning. *Biol. Cybern.* **68**, 95–103.
32. Kettner R. E., Mahamud S., Leung H. C., Sitkoff N., Houk J. C., Peterson B. W. and Barto A. G. (1997) Prediction of complex two-dimensional trajectories by a cerebellar model of smooth pursuit eye movement. *J. Neurophysiol.* **77**, 2115–2130.
33. Kitazawa S., Kimura T. and Yin P. B. (1998) Cerebellar complex spikes encode both destinations and errors in arm movements. *Nature* **392**, 494–497.
34. Linsker R. (1992) Local synaptic rules suffice to maximize mutual information in a linear network. *Neural Comput.* **4**, 691–702.
35. Maex R. and Schutter E. D. (1998) Synchronization of Golgi and granule cell firing in a detailed network model of the cerebellar granule cell layer. *J. Neurophysiol.* **80**, 2521–2537.
36. Malsburg C. v. d. (1973) Self-organization of orientation sensitive cell in the striate cortex. *Kybernetik* **14**, 85–100.
37. Marr D. (1969) A theory of cerebellar cortex. *J. Physiol.* **202**, 437–470.
38. Mitoma H. and Konishi S. (1999) Monoaminergic long-term facilitation of GABA-mediated inhibitory transmission at cerebellar synapses. *Neuroscience* **88**, 871–883.
39. Oja E. (1982) A simplified neuron model as a principal component analyzer. *J. math. Biol.* **15**, 267–273.
40. Olshausen B. A. and Field D. J. (1996) Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* **381**, 607–609.
41. Plumbley M. D. (1993) Efficient information transfer in anti-Hebbian neural networks. *Neural Net* **6**, 823–833.
42. Portas C. M., Bjorvatn B., Fagerland S., Gronli J., Mundal V., Sorensen E. and Ursin R. (1998) On-line detection of extracellular levels of serotonin in dorsal raphe nucleus and frontal cortex over the sleep/wake cycle in the freely moving rat. *Neuroscience* **83**, 807–814.
43. Schweighofer N. and Arbib M. A. (1998) A model of cerebellar metaplasticity. *Learn. Mem.* **4**, 421–428.
44. Schweighofer N., Arbib M. A. and Kawato M. (1998) Role of the cerebellum in reaching movements in humans. I. Distributed inverse dynamics control. *Eur. J. Neurosci.* **10**, 86–94.
45. Schweighofer N., Spoelstra J., Arbib M. A. and Kawato M. (1998) Role of the cerebellum in reaching movements in humans. II. A neural model of the intermediate cerebellum. *Eur. J. Neurosci.* **10**, 95–105.
46. Simpson J. I., Wylie D. R. and De Zeeuw C. I. (1996) On climbing fibers and their consequences. *Behav. Brain Sci.* **19**, 384–398.
47. Spoelstra J., Schweighofer N. and Arbib M. A. (1999) Cerebellar learning of accurate predictive control for fast reaching movements. *Biol. Cybern.* **82**, 321–333.
48. Sutton R. S. and Barto A. G. (1998) *Reinforcement Learning*, MIT, Cambridge, MA.
49. Svoboda K., Denk W., Kleinfeld D. and Tank D. W. (1997) *In vivo* dendritic calcium dynamics in neocortical pyramidal neurons. *Nature* **385**, 161–165.
50. Trouillas P. (1993) The cerebellar serotonergic system and its possible involvement in cerebellar ataxia. *Can. J. Neurol. Sci.* **20**, (Suppl. 3) S78–S82.
51. Turrigiano G. G. (1999) Homeostatic plasticity in neuronal networks: the more things change, the more they stay the same. *Trends Neurosci.* **22**, 221–227.
52. Tyrrell T. and Willshaw D. (1992) Cerebellar cortex: its simulation and the relevance of Marr's theory. *Phil. Trans. R. Soc. Lond. B, Biol. Sci.* **336**, 239–257.
53. van Kan P. L., Gibson A. R. and Houk J. C. (1993) Movement-related inputs to intermediate cerebellum of the monkey. *J. Neurophysiol.* **69**, 74–94.
54. Vos B. P., Volny-Luraghi A. and De Shutter E. (1999) Cerebellar Golgi cells in the rat: receptive fields and timing of responses to facial stimulation. *Eur. J. Neurosci.* **112**, 621–634.
55. Wang S. J., Cheng L. L. and Gean P. W. (1999) Cross-modulation of synaptic plasticity by beta-adrenergic and 5-HT_{1A} receptors in the rat basolateral amygdala. *J. Neurosci.* **19**, 570–577.
56. Watanabe D., Inokawa H., Hashimoto K., Suzuki N., Kano M., Shigemoto R., Hirano T., Toyama K., Kaneko S., Yokoi M., Moriyoshi K., Suzuki M., Kobayashi K., Nagatsu T., Kreitman R. J., Pastan I. and Nakanishi S. (1998) Ablation of cerebellar Golgi cells disrupts synaptic integration involving GABA inhibition and NMDA receptor activation in motor coordination. *Cell* **95**, 17–27.
57. Watson M. and McElligott J. G. (1984) Cerebellar norepinephrine depletion and impaired acquisition of specific locomotor tasks in rats. *Brain Res.* **296**, 129–138.
58. Weiss M. and Pellet J. (1982) Raphe–cerebellum interactions. II. Effects of midbrain raphe stimulation and harmaline administration on single unit activity of cerebellar cortical cells in the rat. *Expl Brain Res.* **48**, 171–176.
59. Zhu W. J., Vicini S., Harris B. T. and Grayson D. R. (1995) NMDA-mediated modulation of gamma-aminobutyric acid type A receptor function in cerebellar granule neurons. *J. Neurosci.* **15**, 7692–7701.

(Accepted 17 November 2000)

APPENDIX A. SPARSE CODING ALLOWS FAST LEARNING

Here we show how a sparse input code increases learning speed for supervised learning in a linear neuron whose activity is given by:

$$z = \sum_{i=1}^n w_i y_i, \quad (\text{A1})$$

where y_i can be considered as GC firing rates, w_i are the efficacies (weights) of the LTD synapses and z the firing rate of a linear PC. The desired association is $z^* = z$, i.e. z^* is the target pattern ($z^* - z$ can be thought as the IO firing rate). We define an error function for the set of




Fig. A1. The shape of the contour plot of the error landscape in weight space illustrates the effect of sparse coding on learning speed. For the purpose of illustration, we consider a hypothetical neuron receiving only two inputs, weighted by synaptic efficacies w_1 and w_2 . If the input code is not sparse (left), the error function is elongated and the gradient descent algorithm requires many iterations. If the input code is sparse, the error function is spherical; the learning rate can then be increased and the gradient descent rapidly converges to the minimum \mathbf{w}^* of the error function.

input patterns and target patterns indexed by μ

$$E(\mathbf{w}) = \frac{1}{2} \sum_{\mu} (z^{*\mu} - z^{\mu})^2 = \frac{1}{2} \sum_{\mu} \left(z^{*\mu} - \sum_{i=1}^n w_i y_i^{\mu} \right)^2. \quad (\text{A2})$$

We can minimize $E(\mathbf{w})$ by sliding down the surface E defined in weight space. The gradient descent algorithm suggests changing each weight by, at time step τ , an amount $\Delta \mathbf{w}^{(\tau)}$ proportional to the gradient of E at the present location $\mathbf{w}^{(\tau)}$. This leads to the least mean square (or delta) learning rule,

$$\Delta \mathbf{w}^{(\tau)} = -\eta \nabla E \Big|_{\mathbf{w}^{(\tau)}} = \eta \sum_{\mu} \left(z^{*\mu} - \sum_{i=1}^n w_i y_i^{\mu} \right) y_i^{\mu}. \quad (\text{A3})$$

We now consider a local quadratic Taylor approximation of the error function around a minimum \mathbf{w}^* of the error function

$$E(\mathbf{w}) \cong E(\mathbf{w}^*) + \frac{1}{2} (\mathbf{w} - \mathbf{w}^*)^T \mathbf{H} (\mathbf{w} - \mathbf{w}^*), \quad (\text{A4})$$

where

$$(\mathbf{H})_{ij} \equiv \frac{\partial^2 E}{\partial w_i \partial w_j} \Big|_{\bar{\mathbf{w}}}$$

is the matrix of second derivatives of the error function evaluated at \mathbf{w}^* and is called the Hessian matrix of the error.

In the neighborhood of \mathbf{w}^* , contours of constant errors are ellipses (see Fig. A1) whose surface axes are aligned with the eigenvectors \mathbf{u}_i of \mathbf{H} , with lengths inversely proportional to the square root of the corresponding eigenvalues λ_i .

It can be shown^{10,27} that the speed of learning along each eigenvector \mathbf{u}_i of the Hessian matrix is determined by the factor $(1 - \eta \lambda_i)$. To speed up learning, we can increase the learning rate η , but to ensure convergence we must have $|1 - \eta \lambda_i| < 1, \forall i$. This limits the value of η to $\eta < 2/\lambda_{\max}$. The convergence rate along the eigenvector corresponding to the smallest eigenvalue is governed by:

$$\left(1 - \frac{2 \lambda_{\min}}{\lambda_{\max}} \right).$$

Consequently, if there is a large difference between the largest and the smallest eigenvalues of the Hessian, learning will be slow (this case corresponds to the narrow valleys in the landscape of the error function E). In contrast, if all the eigenvalues λ_i are equal, the learning can be sped up by increasing the learning rate (this corresponds to a spherical shape of the error function around the minimum, as shown in Fig. A1).

For the linear PC [Eq. (A1)], we have

$$(\mathbf{H})_{ij} \equiv \frac{\partial^2 E}{\partial w_i \partial w_j} \Big|_{\bar{\mathbf{w}}} = y_i y_j. \quad (\text{A5})$$

If the GCs implement a sparse code, all GCs fire independently at the same low probability p . Thus, the Hessian matrix is such that:

$$(\mathbf{H})_{ij} = p^2 \text{ if } i \neq j \text{ and } (\mathbf{H})_{ii} = p. \quad (\text{A6})$$

Because the mean firing rate of the sparse code is low, $p \ll 1$, thus $p^2 \ll p$, and the Hessian can be approximated by a diagonal matrix with identical eigenvalues p . The landscape of E around the minimum is thus spherical and a high learning rate speeds up convergence (see Fig. A1).

APPENDIX B. CEREBELLAR NEURAL MODEL

We describe here the cerebellar neural circuit, based on the Schweighofer *et al.* model,⁴⁵ simulated in the present paper. Learning in the original Schweighofer *et al.* model was arduous because of both poor GC recoding of the MF inputs and inclusion of realistically large delays in the afferent and efferent pathways. We concentrate here only on the GC coding problem by setting all the delays to zero; we refer the reader to Spaelstra *et al.*⁴⁷ for a treatment of the delay problem.

Mossy fiber inputs

MFs in the intermediate cerebellar cortex have clear preference for movements around a specific joint; moreover, MF activities are correlated with joint position, velocity and what appears to be acceleration, but often carries more than one modality.⁵³ MFs have two origins: central and sensory.⁵³ As in the Schweighofer *et al.* model, we assume that the MFs of central origin carry information about the desired trajectory (see Barto *et al.*⁸ for a different approach). Because there are no delays and no noise in the model, desired and sensed trajectories are very similar after learning (also, thanks to the action of the cerebral feedback controller, they are not so dissimilar during learning). Thus, for simplicity, we assume in the present model that all the inputs to the MFs are of central origin. We model the inputs to the MFs by linear combinations of kinetic variables specific to one joint (10 MFs for each joint). For each cell, the membrane potential mf_i and the mean firing rate MF_i are given by:

$$\begin{aligned} \tau \frac{dmf_i}{dt} &= -mf_i + \lambda_{i1} \theta^d + \lambda_{i2} \dot{\theta}^d + \lambda_{i3} \ddot{\theta}^d \\ MF_i &= \frac{MF_i^{\max}}{1 + \exp[-G^{mf}(mf_i - B^{mf})]}, \end{aligned} \quad (B1)$$

where the coefficients λ are generated randomly with equiprobability to $\{-1; 0; 1\}$, the kinematic variables θ^d , $\dot{\theta}^d$ and $\ddot{\theta}^d$ are desired normalized angular position, velocity and acceleration (i.e. for each of these variables, the mean was subtracted and the resulting value was divided by the standard deviation), MF_i^{\max} the maximum firing rate, G^{mf} the input gain to the transfer function and B^{mf} the potential value for which the firing rate is half of its maximum (parameter values are given in Table 1).

Golgi-granule cell complex

To allow a rich combination of the 20 MF inputs and to reproduce the divergence between the MFs and the GCs, we model a layer of $N=100$ GCs. The GC membrane potential and the firing rate are given by:

$$\begin{aligned} \tau_{gc} \frac{dgc_i}{dt} &= -gc_i + \sum_{k=1}^4 w_{ki}^{MF-GC} MF_k - w_i^{GO-GC} \\ GC_i &= \frac{GC_i^{\max}}{1 + \exp[-G^{gc}(gc_i - B_i^{gc})]}, \end{aligned} \quad (B2)$$

where w_{ki}^{MF-GC} reflects the strength of the excitatory MF input and w_i^{GO-GC} that of the inhibitory GO input, GC_i^{\max} is the maximum firing rate, G^{gc} the input gain to the transfer function and B_i^{gc} the potential at which the firing rate is half of its maximum for cell i (that we call ‘‘threshold’’ for convenience). The synaptic efficacies are initially set according to a uniform distribution and the cell parameters are given in Table 1.

Because the Golgi cells receive both direct MF inputs and GC inputs, they regulate the level of GC activity via both feedforward and feedback pathways. The membrane potential and the firing rate of a single GO included in the model is given by:

$$\begin{aligned} \tau_{go} \frac{dgo}{dt} &= -go + \sum_i w_i^{MF-GC} MF_i + \sum_j w_j^{GO-GC} GC_j \\ GO &= \frac{GO^{\max}}{1 + \exp[-G^{go}(go - B^{go})]}, \end{aligned} \quad (B3)$$

where the synaptic efficacies w_j^{GO-GC} reflect the strength of granule input (same notation as above). The cell parameters are given in Table 1.

Cerebellar output and error in performance and supervised learning

Because this study focuses on GC coding, we do not model the PCs or the IO cells according to biological plausibility, but merely retain the essence of the Marr and Albus theory: the GC-PC synapses are modified by IO inputs carrying error information. We make the following simplifications. First, there are two PCs, each controlling a joint, and the PCs are the cerebellar output neurons (i.e. there are no deep nuclear cells). Second, PC and IO cells are linear. Third, the cerebellar LTD is approximated by a least mean square, or delta, rule with momentum (Schweighofer *et al.*⁴³ and Schweighofer and Arbib⁴³ proposed biologically more plausible LTD models).

The membrane potential of each PC is given by:

$$\tau_{pc} \frac{dpc_i}{dt} = -pc_i + \sum_j w_{ji}^{ltd} GC_j, \quad (B4)$$

where w_{ji}^{ltd} are LTD synaptic efficacies. The activity is simply the membrane potential.

As in Kawato and Gomi,³¹ IO cells transmit torque-like errors in performance; sensory representations of actual trajectories are excitatory and efference copies of desired movements are inhibitory. Thus, for the shoulder, for example, the input to the inferior olive is simply:

$$E_1 = K_p(\theta^{fs} - \theta^{ds}) + K_v(\dot{\theta}^{rs} - \dot{\theta}^{ds}) + K_a(\ddot{\theta}^{rs} - \ddot{\theta}^{ds}), \quad (B5)$$

where the coefficients K_p , K_v and K_a are feedback gains, the superscript ‘‘d’’ refers to the desired trajectory of central origin, the superscript ‘‘r’’ refers to the actual trajectory and the superscript ‘‘s’’ refers to the shoulder. Here, the IO activity is simply taken equal to its input; thus for the shoulder, $IO_1 = E_1$.

The global (as opposed to local to each neuron) ‘‘gating’’ term g , which enables threshold adaptation when performance errors occur, is a function of the performance error:

$$\begin{aligned} g &= \gamma \frac{|E_1 \cdot E_2|}{|E_1^{\max} \cdot E_2^{\max}|} \quad \text{if} \quad \frac{|E_1 \cdot E_2|}{|E_1^{\max} \cdot E_2^{\max}|} < 1 \\ g &= 1 \quad \text{otherwise,} \end{aligned} \quad (B6)$$

where E_1 and E_2 are the errors for, respectively, the shoulder and the elbow joint, E_1^{\max} and E_2^{\max} are the maximal joint errors (recorded before learning), and $\gamma = 500$ is a gain parameter. Note that, in the real cerebellum, g does not have to be related to the error movement *per se*, but could also be computed by comparing the desired movement to an efferent copy of the trajectory.

During the supervised learning experiments, at each time step, the synapses w_{ji}^{ld} are updated according to:

$$\Delta w_{ji}^{ld}(t+1) = -\eta IO_j(t) GC_j(t) + \delta \Delta w_{ji}^{ld}(t), \quad (\text{B7})$$

where η is a learning rate and δ a momentum parameter set at 0.99.