BioSystems 94 (2008) 18–27

Contents lists available at ScienceDirect

BioSystems

journal homepage: www.elsevier.com/locate/biosystems

A real-time spiking cerebellum model for learning robot control

Richard R. Carrillo[a], Eduardo Ros[a], Christian Boucheny[b,1], Olivier J.-M.D. Coenen[b,*]

[a] Department of Computer Architecture and Technology, ETSI Informática y de Telecomunicación, University of Granada, Spain
[b] Sony Computer Science Laboratory, Paris, France

ARTICLE INFO

Article history:
Received 23 April 2007
Received in revised form 30 October 2007
Accepted 23 May 2008

Keywords:
Spiking
Neuron
Cerebellum
Adaptive
Simulation
Learning
Inferior olive
Probabilistic
Robot
Real time

ABSTRACT

We describe a neural network model of the cerebellum based on integrate-and-fire spiking neurons with conductance-based synapses. The neuron characteristics are derived from our earlier detailed models of the different cerebellar neurons. We tested the cerebellum model in a real-time control application with a robotic platform. Delays were introduced in the different sensorimotor pathways according to the biological system. The main plasticity in the cerebellar model is a spike-timing dependent plasticity (STDP) at the parallel fiber to Purkinje cell connections. This STDP is driven by the inferior olive (IO) activity, which encodes an error signal using a novel probabilistic low frequency model. We demonstrate the cerebellar model in a robot control system using a target-reaching task. We test whether the system learns to reach different target positions in a non-destructive way, therefore abstracting a general dynamics model. To test the system's ability to self-adapt to different dynamical situations, we present results obtained after changing the dynamics of the robotic platform significantly (its friction and load). The experimental results show that the cerebellar-based system is able to adapt dynamically to different contexts.

1. Introduction

Although the cerebellum architecture has been studied for more than 100 years (Ramón y Cajal, 1995; Golgi, 1967), its functional role is still an open topic. The cerebellum plays a major role in coordinated and accurate movements (Bastian et al., 2000; Ito, 2001). It is thought to be an essential computing tissue for our daily manipulation tasks. Its regular topology has inspired many artificial neural network models in the past decades (Kettner et al., 1997; Medina and Mauk, 1999; Schweighofer et al., 1998a, b; Spoelstra et al., 2000; Arbib et al., 1995; Eskiizmirliler et al., 2002). Furthermore, there are many research groups modelling in detail its cells (D'Angelo et al., 2001; Bezzi et al., 2004; Steuber et al., 2004) in order to elucidate the specific computations that take place at each part of the cerebellum architecture.

There have been great advances in robotics, mainly in industrial applications. Yet most of the industrial robots use stiff joints and high-gain closed-loop control. The movement of stiff joints facilitates control since it reduces (or even avoids) the necessity of dynamics models. Industrial robots are able to perform accurate trajectory-following adopting online closed-loop error-correction schemes. This strategy became possible due to the outstanding processing speed of current circuits that calculate errors and deliver feedback correction signals on a microsecond time scale. Nevertheless, stiff-joint control does not take advantage of the robot dynamics, which results in unnatural control, wasted energy and reduced robot autonomy.

In contrast, biological limbs have joints of variable stiffness and use low-gain control schemes where the dynamics cannot be ignored. Indeed, the dynamics, for instance, of an arm–hand system, is likely to be significantly modified when manipulating objects of different weights. Moreover, biological systems have delays in sensorimotor pathways up to several hundreds of milliseconds. This makes it impossible to apply on-line closed-loop error-correction strategies without having predictor modules able to abstract the kinematics and dynamics models of the platform.

There are plenty of challenges in robotics such as the development of accurate low-gain control schemes for robotic platforms of several degrees of freedom (DOF) and compliant joints (non-stiff). Compliant joints and low-gain control are of particular interest for robots interacting with humans. For security reasons, a robot becomes a safer platform if it is not able to apply dangerous forces and can absorb energy.

* Corresponding author.
E-mail addresses: rcarrillo@atc.ugr.es (R.R. Carrillo), eduardo@atc.ugr.es (E. Ros), christian.boucheny@wanadoo.fr (C. Boucheny), olivier@oliviercoenen.com (O.J.-M.D. Coenen).
URLs: http://atc.ugr.es/ eduardo (E. Ros), http://www.oliviercoenen.com (O.J.-M.D. Coenen).
[1] Current address: Evasion-Laboratoire GRAVIR INRIA, Rhône-Alpes, France.

0303-2647/$ – see front matter © 2008 Published by Elsevier Ireland Ltd.
doi:10.1016/j.biosystems.2008.05.008

In this paper, we emulated the learning strategy followed by biological systems to control low-gain compliant robotic platforms in the presence of sensorimotor pathways with delays of hundreds of millisecond. To do this, we studied how a cerebellum model can abstract dynamics models of the robotic platform to facilitate control by predicting and correcting errors in the motor space.

2. Cerebellum Model

Previous modelling of the cerebellum contribution in movement learning includes the modeling of smooth pursuit eye movements (Kettner et al., 1997). In this work, the cerebellar nuclei cells were not implemented in their model, and analog units, not spiking neurons were used. Schweighofer et al. (1998a) proposed a cerebellum model learning the inverse dynamics of a two-link six-muscle arm system. The parallel fiber-Purkinje cell (PF-PC) long-term depression (LTD) was biologically inspired, but not the long-term potentiation (LTP), which was implemented as a weights normalization process. Moreover, learning was performed over short trials only (less than 500 ms) and not continuously as in our contribution.

A few cerebellar models for eyelid conditioning have used spiking neurons (e.g. Medina and Mauk, 1999; Hofstötter et al., 2002). Learning was based on spikes coincidences between neurons, but

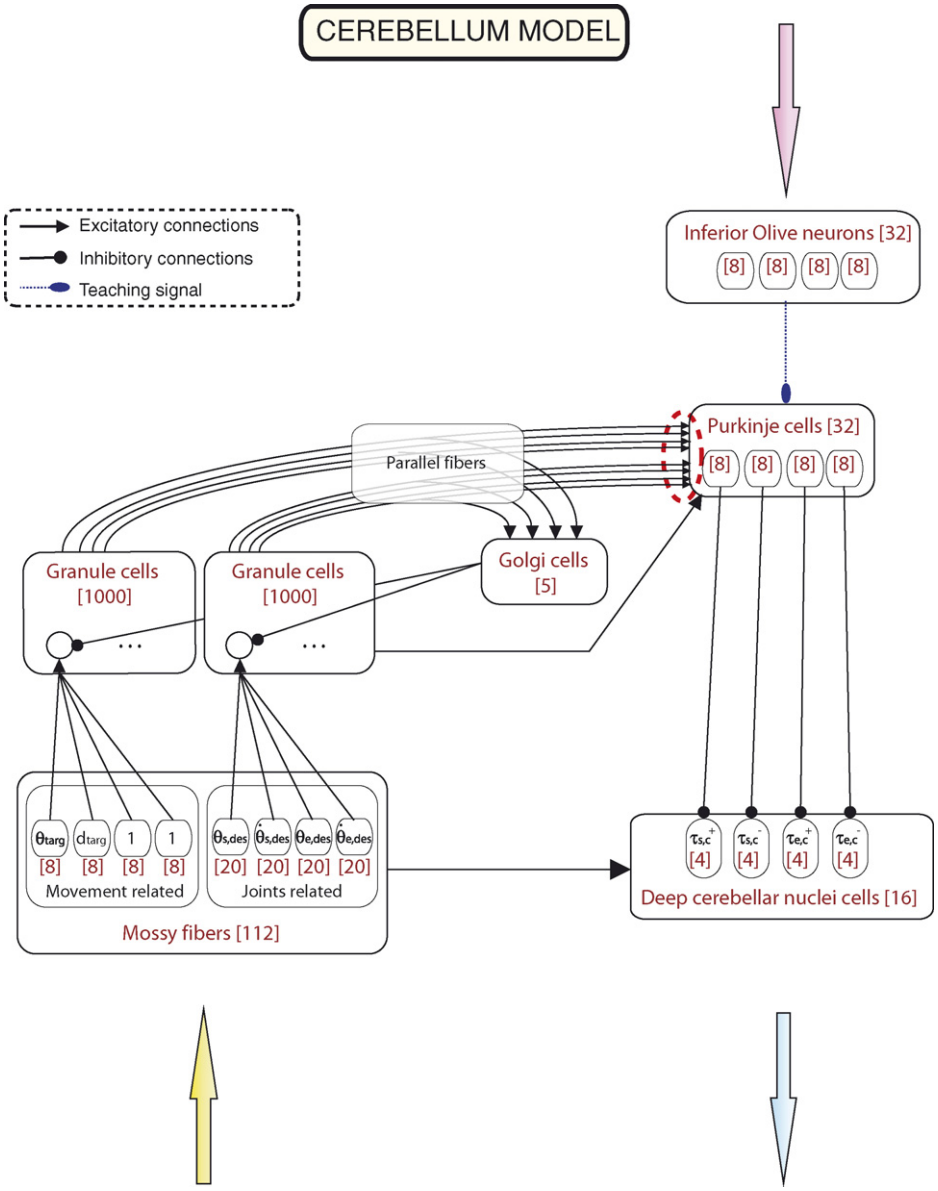

Fig. 1. Cerebellum model diagram. Inputs about the movement (desired arm state and target information) were sent (upward arrow) to the two layers of mossy fibers (MF): distance to the target and its absolute position in the experimental field ($d_{\text{targ}}(t)$ and θ_{targ}) as well as desired positions (θ) and speeds ($\dot{\theta}$) of the shoulder (s) and elbow (e) joints along the trajectory. These desired states were obtained from a crude inverse kinematic model (see Fig. 9), representing motor cortex and other motor areas. The mossy fibers projected to two layers of granule cells (GR, 1000 neurons per layer) and to 16 deep cerebellar nuclei (DCN) cells. The 32 Purkinje cells (PC), 16 DCN cells and 32 inferior olive (IO) neurons were divided along 4 functional zones (inspired from cerebellar microzone organisation), one for each of the actuators, agonist or antagonist, of the elbow and shoulder joints. The PC received excitatory inputs from all the desired joint state-related GR (ascending axons that maintain the cells in a state of excitability) and from all parallel fibers PF with a connection probability of $p_{\text{PC-PF}}$= 0.8. They also received an afferent from the IO in a one-to-one scheme. In turn, the DCN cells received inhibitory connections from two PCs of the same microzone. The teaching signal was processed by the IO cells (downward arrow, top) (see Fig. 9). At the output of the cerebellum (downward arrow, bottom), the DCN firing rates were interpreted as predictive positive (+) and negative (−) torque corrections (τ) for the shoulder (s) and elbow (e) joints. The numbers in brackets indicate the number of cells per layer and per zone.

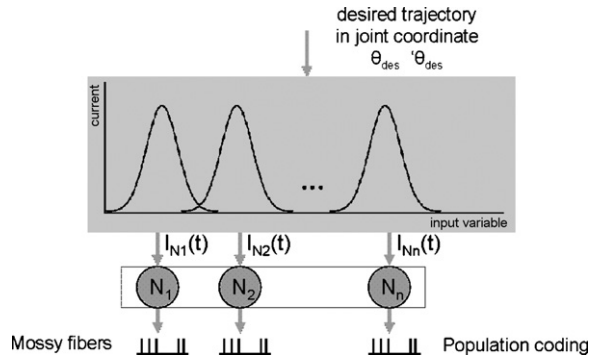

Fig. 2. Encoding of mossy fibers. The analog to spikes transformation for driving the mossy fibers used overlapping radial basis functions (RBF). The example here makes reference to encoded joint variables (see Figs. 1 and 9).

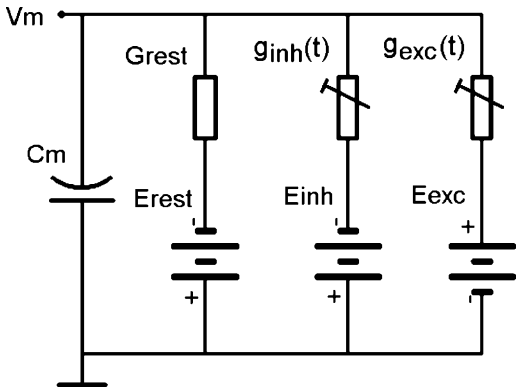

Fig. 3. Cell model with conductance-based synapses. C_m: membrane capacitance, V_m: membrane voltage, E_{rest}: resting reversal potential, G_{rest}: leakage conductance at rest, E_{exc}: excitatory reversal potential, $g_{exc}(t)$: conductance of excitatory synapse, E_{inh}: inhibitory reversal potential, $g_{inh}(t)$: conductance of inhibitory synapse.

none used the same probabilistic low-frequency firing of the inferior olive in their learning rules.

We simulated in real time a cerebellum spiking neural model (Boucheny et al., 2005; Arnold, 2001; Ros et al., 2006; Huang et al., 1998) made of approximately 2100 units: 112 mossy fibers (MF), 2000 granule cells (GR), whose axons form the parallel fibers (PF), 5 Golgi cells (GC), 32 inferior olive (IO) cells, whose axons form the climbing fibers (CF), 32 Purkinje cells (PC) and 16 deep cerebellar nuclei (DCN) cells (Fig. 1).

The cerebellum spiking neural model was simulated with the EDLUT simulator (Event-Driven simulator based on LookUp Tables) (Ros et al., 2006). The EDLUT simulator first compiles the neuron models offline to avoid heavy numerical calculation during the network simulation. The EDLUT environment facilitates a direct interface to real robotic platforms by permitting real-time simulation of large-scale spiking neural networks. The software is particularly suited for a cerebellar model in which sparse activity is expected (Coenen et al., 2001; Schweighofer et al., 2001) in the numerous neurons of the granular layer (there are approximately 10^{11} granule cells in the cerebellum (Kandel et al., 2000)).

Mossy fibers were implemented as leaky integrate-and-fire neurons. Their input current was determined by a radial basis function (RBF), which received one of the sensory variables (e.g. target position or velocity) or one of the desired joint states (position and speed of elbow and shoulder joints) (Figs. 1 and 2). The RBF centers were evenly distributed across the sensory dimensions, and their variance were chosen to ensure small responses overlap from consecutive mossy fibers.

Lasting functional changes at the synaptic level can be driven by the coincidence of multiple signals at individual synaptic sites (Brown et al., 1990). Long-term depression of the parallel fiber input to cerebellar Purkinje cells is a form of synaptic plasticity that can last from hours to days (Ito and Kano, 1982) and is thought to underlie several forms of associative motor learning (Mauk et al., 1998). In the cerebellar model that we present, long-term depression (LTD) is induced by correlating activation of parallel fiber (PF) and climbing fiber (CF) synaptic inputs (see Section 2.3).

The inferior olive (IO) neurons synapse onto the Purkinje cells and contribute to direct the plasticity of PF-PC synapses. These neurons, however, fire at very low rates (less than 10 Hz), which appears problematic to capture the high-frequency information of the error signal of the task being learned. This apparent difficulty may be solved by their irregular or chaotic firing (Keating and Thach, 1995; Kuroda et al., 2001; Schweighofer et al., 2004). We suggest that this is a very important property, which has the beneficial consequence to sample statistically the entire range of the

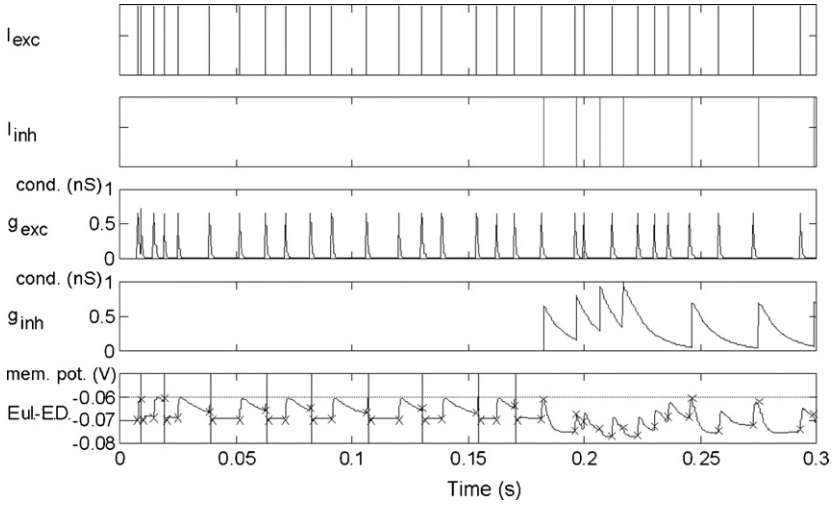

Fig. 4. Neuron simulation. (From top to bottom) The neuron receives excitatory (exc) and inhibitory (inh) input spikes (I), which provoke abrupt increases in the synaptic conductances (g), followed by their exponential decay. (Bottom) The neuron membrane potential (mem. pot) follows Eq. (1). A comparison is made between the continuous integration of Eq. (1) (Eul) (continuous line) and the EDLUT event-driven (ED) computations, which occur only at the time of the spikes indicated by the 'X's. The match is excellent.

error signal over multiple trials (see below). Here, we implemented this irregular firing using a Poisson model for spike generation.

Error correction was accomplished by changes in the activity of Purkinje cells that in turn influenced the activity of the deep cerebellar nuclei cells (Purves et al., 2001). These were translated into analog torque correction signals for the robot.

2.1. Neuron Models

Different neuron types (granule cell, Purkinje cell and Golgi cell) were included in the simulated network. An integrate-and-fire neuron is represented by the circuit in Fig. 3 and defined by Eq. (1). Its response behaviour using two different simulation strategies is shown in Fig. 4. This neuron model is a modified version of the Spike-Response-Model (SRM) (Gerstner and Kistler, 2002) widely used in the literature (Eckhorn et al., 1990; Schoenauer et al., 2002; Shaefer et al., 2002) to study, for example, temporal coding issues (Eckhorn et al., 2004).

$$C_m \frac{dV_m}{dt} = g_{exc}(t)\ (E_{exc} - V_m(t)) + g_{inh}(t)\ (E_{inh} - V_m(t)) + G_{rest}\ (E_{rest} - V_m(t))$$

$$g_{exc}(t) = \begin{cases} 0 & t < t_0 \\ G_{exc}\ e^{\frac{-(t-t_0)}{\tau_{exc}}} & t \geq t_0 \end{cases} \qquad (1)$$

$$g_{inh}(t) = \begin{cases} 0 & t < t_0 \\ G_{inh}\ e^{\frac{-(t-t_0)}{\tau_{inh}}} & t \geq t_0 \end{cases}$$

The synapses were modelled as input-driven conductances (with a positive abrupt change triggered by the post-synaptic event and an exponential decay); AMPA mediated excitatory synapses had a shorter time constant (τ_{exc}= 0.5 ms) than GABA inhibitory synapses (τ_{inh} = 10 ms). Each neuron type is defined by different characteristics (parameters) according to neurophysiological characterization studies (D'Angelo et al., 1995, 2001; Maex and De Schutter, 1998; Barbour, 1993; Solinas et al., 2003).

Table 1
Connectivity table of the cerebellum model

Cell type	Number	Afferents from	Efferents to
Granule	2000	4 mossy fibers	5 Golgi 32 Purkinje
Golgi	5	1000 granule	2000 granule
Purkinje	32	1500 granule 1 climbing fiber	2 deep cerebellar neurons

In our model, the inferior olive cells transmitted the error signal using probabilistic low rate spikes. Mossy fibers carried sensorimotor signals encoded into rate coded spike trains (activity 0–100 Hz). Finally, deep cerebellar nuclei cells provided spike trains which encoded corrective motor torque signals.

2.2. Cerebellum Topology

The model reproduced the cerebellum's different functional and topological features (Andersen et al., 1992; Kandel et al., 2000): sparse coding at the parallel fibers (Coenen et al., 2001; Schweighofer et al., 2001), converging topology into Purkinje cells, Purkinje cell receiving a dedicated "teaching climbing fiber" from the inferior olive, inhibition to the granule cells from collector Golgi cells, etc. (Table 1 and Fig. 1)

2.3. Cerebellar Learning Rules

We implemented learning at the parallel fibers to the Purkinje Cells connections (indicated by a dashed ellipse in Fig. 1) (Ito, 2001). The parallel fibers brought in the sensorimotor information and the Purkinje cells drove the cerebellum output through the deep cerebellar nuclei cells. The weight adaptation was driven by the activity generated by the inferior olive (IO), which encoded an error signal into a low frequency probabilistic spike train (from 0

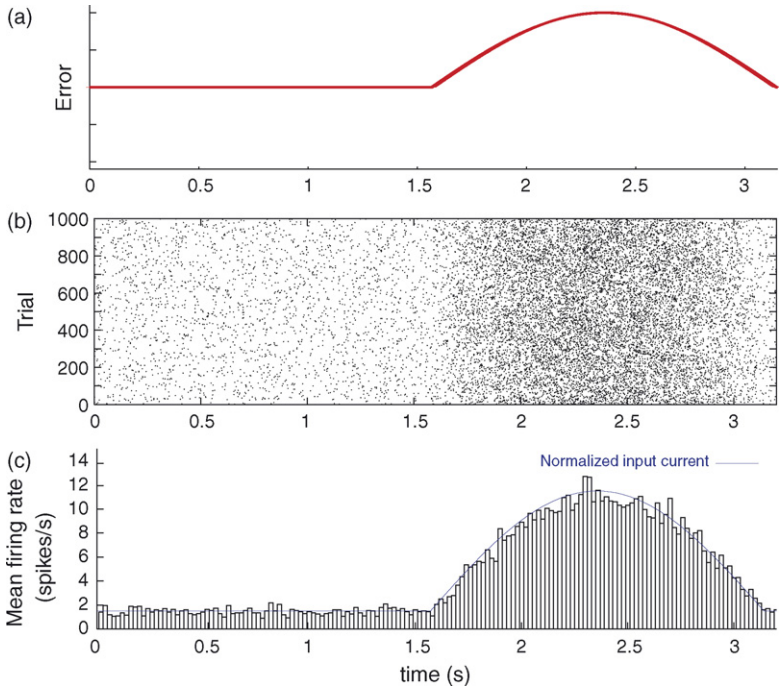

Fig. 5. Inferior olive probabilistic encoding of the error. (a) Example of the error to be encoded. (b) Probabilistic firing of an inferior olive cell to the error in "(a)" (see text). (c) Mean firing rate of the cell averaged over all trials in "(b)". Notice that the maximum firing rate is close to 10 Hz. The smooth curve shows the normalized input current to the cell related to the error amplitude. Notice how the cell never fires quite at the same moment relative to the error, but encodes it nevertheless.

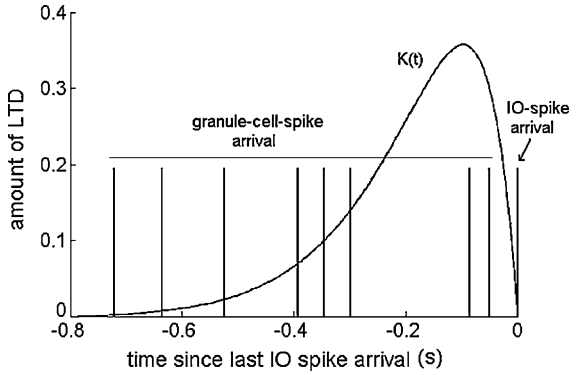

Fig. 6. Spike-Timing Dependent Plasticity (SDTP). Kernel used for granule cell (GR) and Purkinje cell (PC) synaptic long-term depression, corresponding to the solution of a second order differential system. The kernel is convolved with the spike train of the afferent PF (all spikes emitted for $t < 0$). This provides a measure of past PF activity setting the eligibility of the synapse to depression when the inferior olive (IO) neuron afferent to the PC emits a spike ($t = 0$).

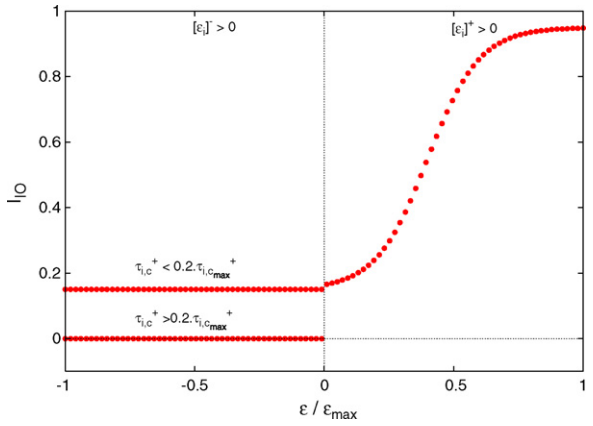

Fig. 7. Input current to inferior olivary cells. Each olivary cell is related to the agonist muscle i and its firing is dependent on the error signal for this muscle (see text). This reflects the influence of the deep cerebellar nuclei (DCN) feedback on the inferior olive (IO) together with an effector arm system made of agonist and antagonist muscle pairs. The left side of the vertical line is for an error on the antagonist muscle, whereas the right side is for the agonist muscle. The rule states that for an error on the antagonist muscle (left part), if the torque $\tau_{i,c}^+ > 0.2\tau_{i,c_{\max}}$, then the IO current $I_i^+ = 0$ (bottom line), otherwise $I_i^+ = 0.15$ (top line).

to 10 Hz, average 1 Hz) (Kuroda et al., 2001; Schweighofer et al., 1998b).

We modeled the inferior olive cell responses with a probabilistic Poisson process: given the normalized error signal $\epsilon(t)$ and a random number $\eta(t)$ between 0 and 1, the cell fired a spike if $\epsilon(t) > \eta(t)$, otherwise it remained silent (Boucheny et al., 2005). In this way, on one hand, a single spike reported an accurately timed information regarding the instantaneous error; and on the other hand, the probabilistic spike sampling of the error ensured that the whole error region was accurately represented over trials with the cell firing at most 10 spikes per second. Hence, the error evolution is accurately sampled even at low frequency. The histogram of the inferior olive output spikes reproduces the error signal temporal trace; see Fig. 5 for an example. This firing behavior is similar to the ones obtained in physiologial recordings (Kuroda et al., 2001).

The long-term potentiation (LTP) implemented at the parallel fiber to Purkinje cell synapses was a non-associative weight increase triggered by each granule cell spike (Eq. (2)) (Lev-Ram et al., 2003). The long-term depression (LTD) was an associative weight decrease triggered by spikes from the inferior olive (Eq. (3)) (Ito and Kano, 1982; Ito, 2001). This model of LTD uses a temporal kernel (Fig. 6), which correlates each spike from the inferior olive with the past activity of a granule cell and shows a peak at 100 ms (Kettner et al., 1997; Spoelstra et al., 2000; Raymond and Lisberger, 1998).

The network maximizes learning (LTD) at synaptic sites in which the input parallel fiber delayed activity is highly correlated with the error signal from the inferior olive. Hence, this kernel produces a predictive corrective output in the network that helps the control task in the presence of significance transmission delays.

$$\text{LTP:} \quad \Delta w(t_o) = \alpha\, \delta_{GR}(t_o) \tag{2}$$

$$\text{LTD:} \quad \Delta w(t_{IO}) = -\int_{-\infty}^{t_{IO}} K(t - t_{IO})\, \delta_{GR}(t)\, dt \tag{3}$$

The teaching signal relied on the motor error, namely the discrepancy between the desired state of the joints at time t and the actual one. The error for each joint, respectively, ϵ_s and ϵ_e, for shoulder and elbow, was computed as the sum of the position and velocity errors, weighted by coefficients $K_p = 10$ and $K_d = 23$ (same for each joint). The signals were delayed in order to align them in time. The desired command at time t was applied at time $t + \delta_1$ and the joint state at time $t + \delta_1$ was sensed by the system at time $t + \delta_1 + \delta_2$. Hence, the error signal for joint i at time t was given by: $\epsilon_i(t) = K_p(\theta_{i,des}(t - \delta_1 - \delta_2) - \theta_i(t - \delta_2)) + K_v(\dot{\theta}_{i,des}(t - \delta_1 - \delta_2) - \dot{\theta}_i(t - \delta_2))$.

Physiologically, the time-matching of the desired and actual joint states can be understood by the fact that the trajectory error would be detected at the level of the spinal cord, through a direct drive from the gamma motoneurons to the spinal cord (Contreras-Vidal et al., 1997; Spoelstra et al., 2000).

The error signal ϵ was used to compute the value of the input current to each IO cell. Smoothing was performed using a sigmoid, and inhibition of IO cells by DCN neurons was taken into account within a formal scheme. The positive part of the

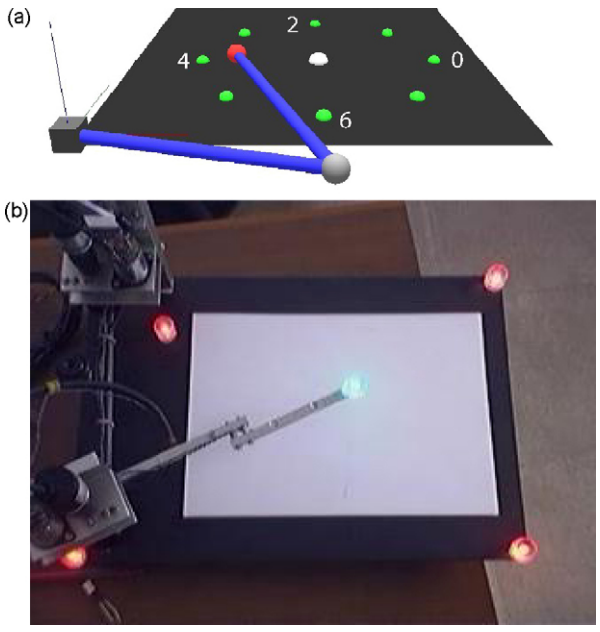

Fig. 8. Experimental robotic platform. (a) Representation of the arm in simulation. Each green point (grey in b&w) represents a target position (0–7) along a circle. (b) Two degrees of freedom (DOF) robotic arm used in the experiments. The motors have no gears and therefore are non-stiff (compliant) low-torque motors with non-linearities difficult to control.

error signal for joint i, $[\epsilon_i]^+$ was related to an error in the corresponding agonist muscle, and the negative part $[\epsilon_i]^-$ to an error in the antagonist muscle. If we denote $\tau^+_{i,c}$ the corrective torque command computed by the cerebellum for agonist muscle i at time $t-\delta_1-\delta_2$, then the input current I^+_i to IO cells within the microzone i^+ was given by: if$[\epsilon_i]^+ > 0$, then $I^+_i = 0.15 + (0.8)/(1+\exp(-10([\epsilon_i]^+)/(\tau^+_{i,c_{\max}}) + 4))$, if$([\epsilon_i]^- > 0 \,\&\, \tau^+_{i,c} > 0.2\tau_{i,c_{\max}})$, then $I^+_i = 0$, otherwise, $I^+_i = 0.15$ (see Fig. 7).

The three equations above correspond, respectively, to the cases when the cerebellar output undershot, overshot or equalled the output torque required for proper motor correction. The second equation modeled mathematically the efferent inhibition of the DCN to the IO and is interpreted as follows: if a non-negligible correction was output to agonist muscle i (DCN neuron output) whereas the movement required a positive correction for the antagonist muscle (error signal), then the unwilled correction was reduced (via IO inhibition by DCN neurons depending on the opposite error signal). The error currents were normalized using $\epsilon_{s_{\max}} = 1000$ and $\epsilon_{e_{\max}} = 600$ for the shoulder and the elbow, respectively.

3. Robotic Platform

The robotic platform was a two-DOF arm (Fig. 8b). The two joints were not stiff (compliant) and the motors applied low forces. The platform allowed continuous measurements of the position of each joint and power consumption. A pen or a weight could be attached to the arm's ending to change its dynamics.

The control system was simulated on computer. To relieve the computer from interface computation and permit real-time communication with the robot, an FPGA-based board contained position acquisition modules and motor-driver controller circuits. The controller modules translated the motor-torque commands from the computer into continuous signals using pulse-width mod-

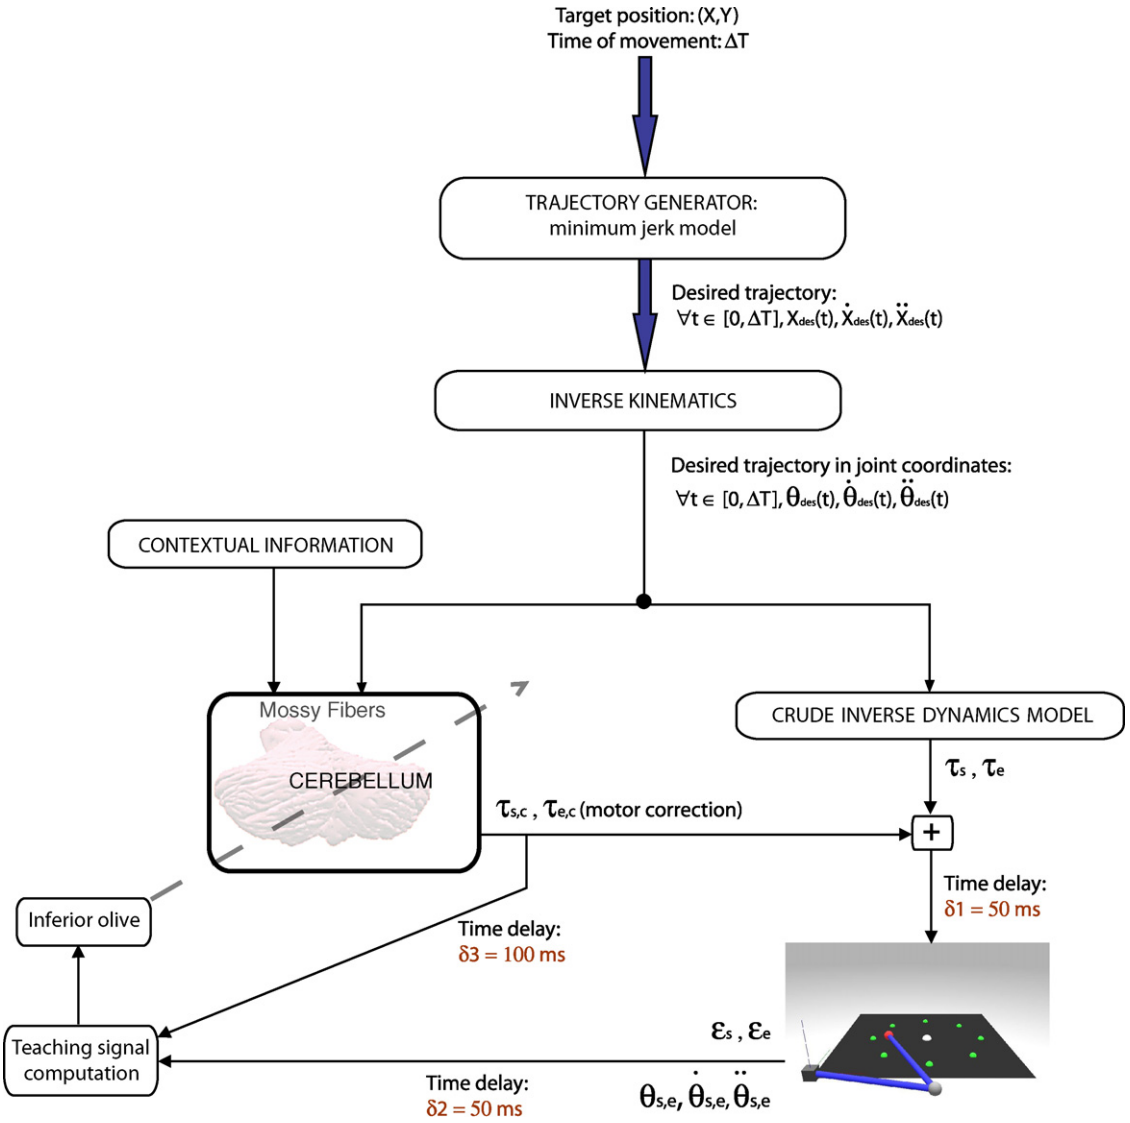

Fig. 9. Diagram of the control system for arm movement generation. The cerebellum acts as a predictive corrective module in the control loop. A desired smooth trajectory toward the target was computed in cartesian coordinates and transformed into joint coordinates. These desired arm states were used at each time step to compute a crude torque command and to update the predictive corrective command of cerebellum. The cerebellum command included information about the context of the movement. The two torques, crude and corrective torques, were summed to control the arm movement with a delay of δ_1 = 50 ms. In turn, the error of the resulting trajectory was sensed at the level of the limb and sent back to the system with a delay of δ_2 = 50 ms. This error was transformed to compute the cerebellum training signal by inferior olive neurons.

ulation (PWM). The PWM signal was supplied to the motors by a current-driver circuit.

4. Experimental Results

The spiking neurons of the cerebellar network (Section 2) were simulated using a computationally efficient table-based event-driven simulator, the EDLUT simulator (Ros et al., 2006) (see Fig. 4). EDLUT is particularly suited for a cerebellar model where sparse activity is expected in the numerous neurons of the granular layer (Coenen et al., 2001; Schweighofer et al., 2001). Plasticity for this model was also developed to allow online real-time learning.

The control system was first tested in simulations, then run on the experimental robotic setup (Fig. 8). Starting from a central position, the robotic arm performed straight movements to reach one of the different targets equally set on a circle (radius of 20 cm). The movements were performed at high speed (T = 0.5 s for each complete movement) to check the ability of the cerebellum to abstract the robotic platform dynamics.

To interact in real time, the robotic platform communicated with the EDLUT simulator every millisecond. At every time step the sensory data (robot joints position) was translated into spike trains transmitted through the mossy fibers. The cerebellar output spike trains were translated into torque correction signals (outputs of the deep cerebellar nuclei cells) and the error signal was transformed into a probabilistic low frequency spike train (inferior olive cell probabilistic model).

The simulations were run on a Pentium IV 2.8 GHz. There were 2100 neurons in the network for approximately 52,000 synaptic connections. During 1 s of simulation, the cerebellar network received an average of 395 spikes, delivered 405 output spikes, and processed 935,801 events. Under these conditions the simulator ran in real time the full network and the input-output transformations.

Considering the duration of motor execution (T = 0.5 s) relative to the time delays in cortico-spinal loops (up to 300 ms), we made the assumption that each reaching movement was performed in open-loop (no high-level motor correction were applied while reaching the target). Corrective commands to compensate for dynamics perturbations were computed only by the cerebellar model.

A movement was separated in two phases:

- **Open-loop movement phase:** A movement lasted T_{move} = 500 ms. The torque command applied to each articulation i was the sum of the cerebellar correction ($\tau_{i,c}$) and the i th torque (τ_i), computed by a basic inverse dynamic model according to the desired kinematic trajectory (Fig. 9). These two commands were sent to the limbs with a delay of δ_1 = 50 ms.
- **Post-movement phase:** It was set to a duration of T_{post} = 0.2 s. Its goal was to stop the movement of the arm, independently of its position relative to the target. The torque applied to each joint corresponds to the non delayed output of a derivative controller with a null-desired velocity: $\tau_i = K_{\text{vstop}}\theta_i$ with K_{vstop}=10. The lack of delay in such a command in a human arm control model can be explained by a different motor strategy, consisting, for example, in a high level co-contraction command of the antagonist muscles controlling an articulation.

The architecture of the model for the generation of accurate fast arm reaching movements is illustrated in Fig. 9. A minimum jerk model (Flash and Hogans, 1985) was used to compute the desired smooth trajectory of the arm end-point towards the target at (O_X, O_Y). The desired trajectory was expressed in Cartesian coordinates and transformed into joint coordinates by the inverse kinematic module. To solve the redundancy problem in the coordinates transformation, the robotic arm position was set to always be in a biological plausible posture, e.g. that the angle between the two links of the limb were to remain positive.

During the open-loop period of the movement, the torque commands sent to the joints were the sum of the output of a crude inverse dynamic controller and of the anticipative corrective cerebellar output. These torques were sent to the limb with a time delay δ_1 = 50 ms.

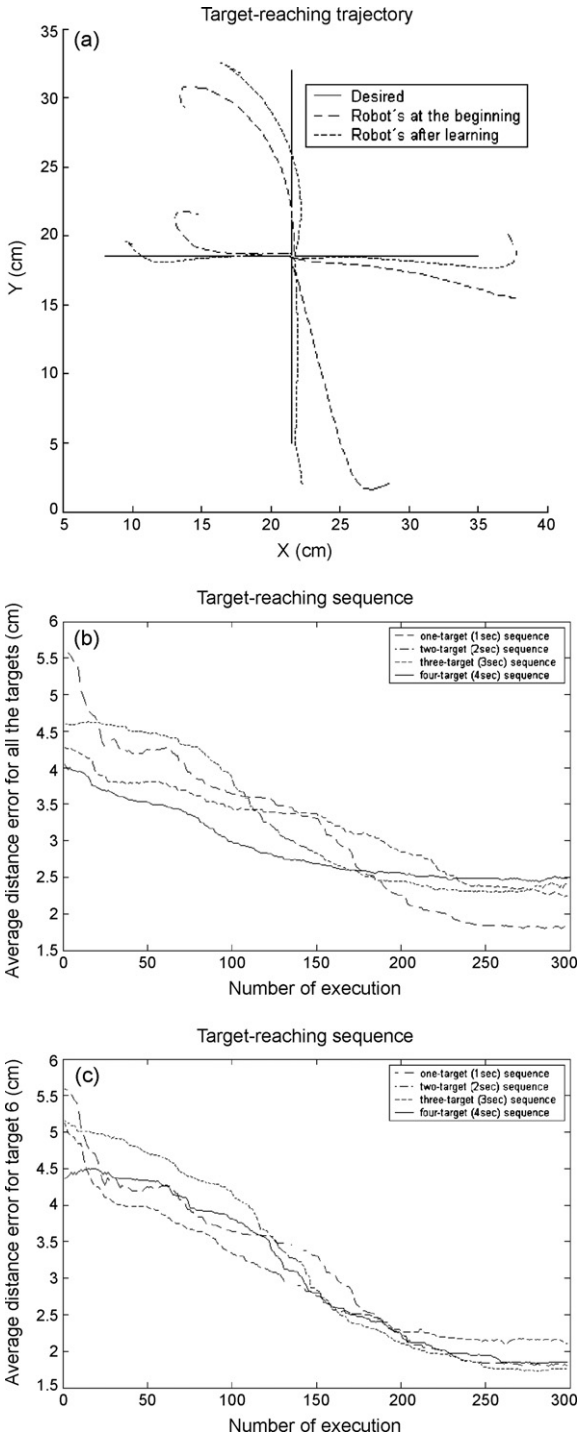

Fig. 10. Target-reaching experiments. (a) Trajectory followed by the arm's ending. (b) Average distance error computed over all trajectories when learning 1, 2, 3 or 4 different trajectories. (c) Distance error of the trajectory of target No. 6 when learnt conjointly with 1, 2, 3 or 4 different trajectories.

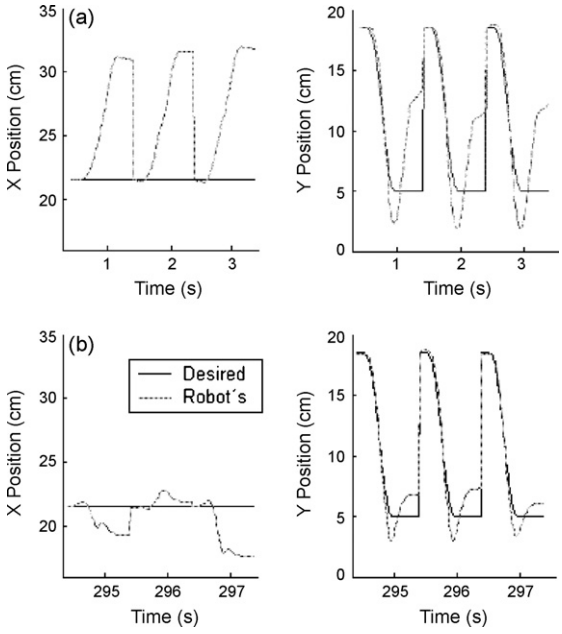

Fig. 11. Target-reaching example. Desired and actual arm ending position along the *x*- and *y*- axes (a) before learning and (b) after learning. 3 trials (3 s) are shown. The curve part of the trajectory shows the open-loop movement. The movements to reset the trials are not shown; this explains the abrupt vertical lines.

The error in the execution of movement was computed at the level of the arm, and sent back to the system with a delay of δ_2 = 50 ms. It was mainly used to determine the teaching signal conveyed by the inferior olive to the cerebellum to produce anticipative motor corrections. The error signal was composed of an angular position error and an angular velocity error for each articulation.

Finally, the cerebellar neural network received non delayed desired trajectory and movement context, modeling inputs originating from motor cortex and other areas, and its output participated to the construction of the teaching signal with a delay of δ_3 = 100 ms.

The inverse dynamic module was based on simplistic assumptions, such as mass homogeneity along the limbs and friction factor to compensate roughly for friction torques that reached 17 N cm for the shoulder motor and 3 N cm for the elbow. Other sources of dynamical perturbations, such as the forces exerted by the wires on the arm, were negligible compared to friction.

After defining an acceptable crude controller, we verified the repeatability of the movements and therefore of the errors of the crude controller. Indeed, the role of the cerebellum was to learn the anticipative corrections required across repeated trials of the same task. If the dynamics perturbations moving the arm to the desired paths varied too much across different trials under the same context (manipulating the same object) then no improvements could have been expected for the proposed control/correction scheme.

The model learned effectively and concurrently different target trajectories (Fig. 10). An example shows the movement in *x*–*y* coordinates before and after learning (Fig. 11). The cerebellum corrections build up over trials to compensate for the movement errors (Fig. 12).

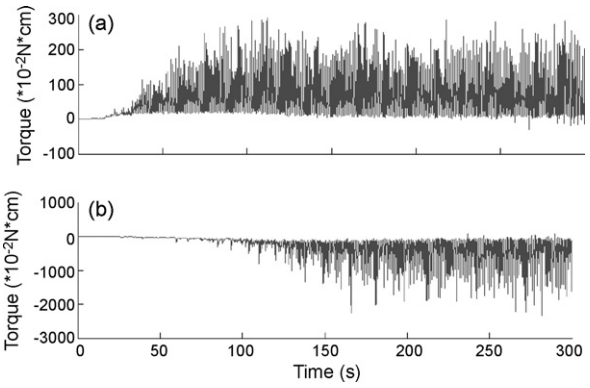

Fig. 12. Cerebellar torque contributions to target reaching experiments over the first 300 trials. Cerebellar torque increases as the system learns (a) at the elbow and (b) at the shoulder. Each trial lasted 1 s.

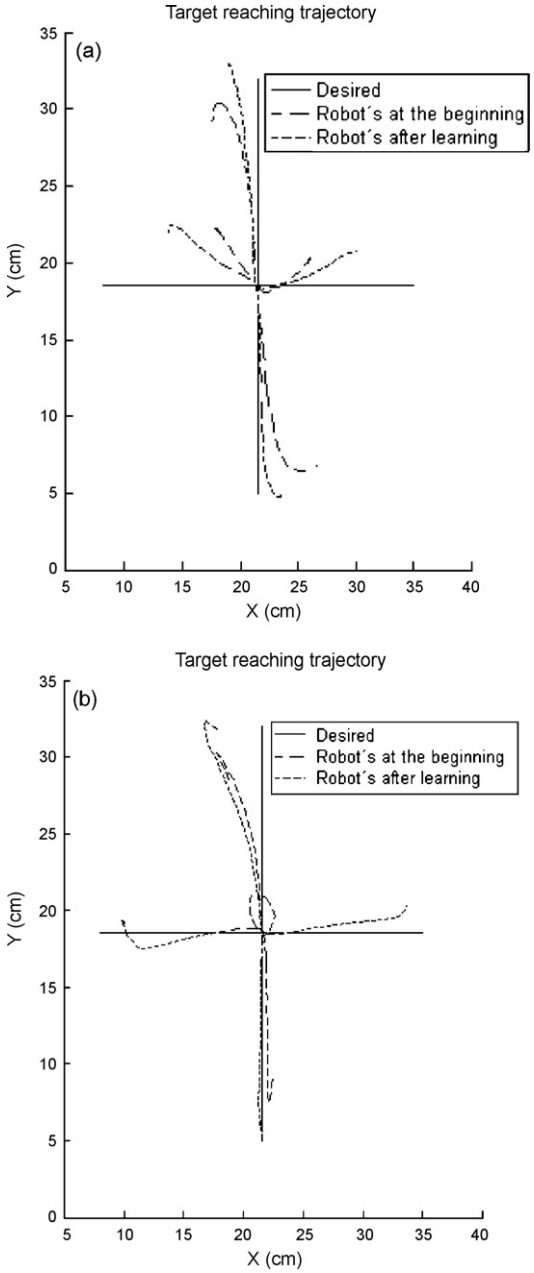

Fig. 13. Learning to compensate for the dynamics changes of the arm. (a) A 0.5 kg was added at the end of the robotic arm. (b) Friction was increased by inserting the end of the robotic arm into a sand pool. Notice how, before learning, the robot moved nearly vertically, whereas the intended movement was horizontal along the *x*-axis. It was indeed correct after learning.

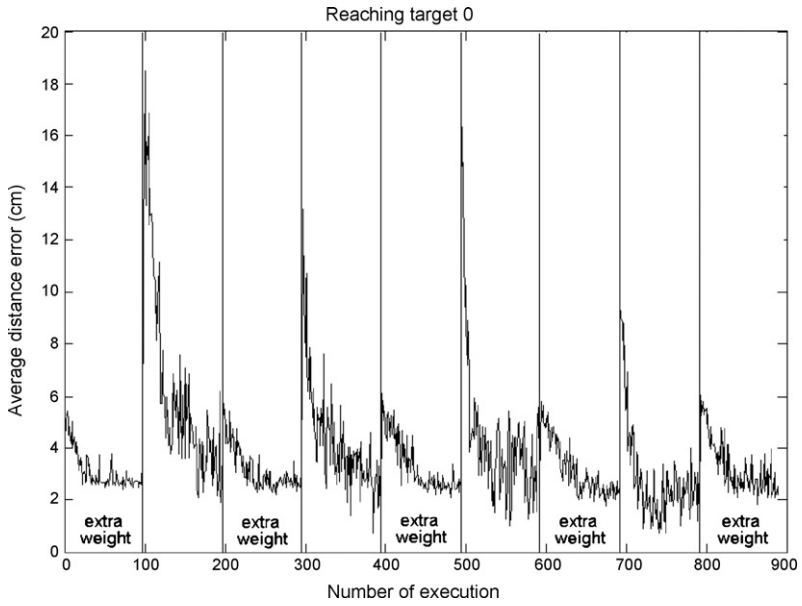

Fig. 14. Temporal adaptation. Error evolution as the task was changed from manipulating a 0.5 kg load to manipulating no load.

We also performed experiments where the dynamics of the arm was change either by a load of 500 g added to the end of the two-joint arm or by modifying the friction of the arm by inserting the end of the arm into a sand pool. The results of the cerebellum-driven improved trajectories are shown in Fig. 13.

The evolution of the error as the object/context was changed is shown in Fig. 14. The cerebellum network learned the new context every time it was changed. It also appeared to adapt more rapidly to the no-load condition over time, although a more detailed analysis is needed to confirm this. Note that the load, no-load condition was not explicitely encoded here, hence the system could not switch immediately to the right conditions without an adaptation period first.

These experimental results show that the control system with the cerebellum model can learn to compensate for dynamics perturbations caused by different contexts: friction or load changes that significantly alter the robot arm inertial moments.

We have shown how the spike-timing dependent plasticity (STDP) rule works as a temporal kernel filter relating the activity from the inferior olive (error dependent) with the sensorimotor inputs received through the granule cells. This scheme was able to construct predictive dynamic corrections for fast reaching movements.

A residual average distance error can be noticed even after the learning has stabilized. This error could be attributed to some system limitations but also to the fact that we dealt with a real robot which responded differently over time. For instance, over trials, the robot's motors increased considerably their temperature. This prevented the cerebellar model from adapting completely to the robot response, unless a richer and more complete sensorimotor context were made available.

However the goal of the present contribution was not to focus on designing a high performance control scheme but rather to evaluate an adaptive and robust working hypothesis based on a specific physiologically-relevant cerebellar network that could run and learn in real time. The performance obtained fulfilled this requirement although deeper studies on complementary mechanisms will be studied in the future to evaluate how the control strategy can take full advantage of further biologically plausible features of the system.

5. Discussion

In this work we have simulated a complete physiologically-relevant spiking cerebellar model in real time, and evaluated its potential role in generating predictive corrective actions towards accurate control in fast robotic reaching movements.

Whereas with previous simulators many computing hours would have been required to simulate a spiking cerebellar model learning to correct trajectories, with the current simulator, learning took place in less than a real hour to achieve acceptable performance levels allowing real-time control of a robot.

This performance was achieved even with the physiologically realistic firing of the inferior olive restricted to less than 10 Hz. To the best of our knowledge, this is the first time that such performance is obtained with the present biological characteristics in a complete action-perception loop using a real robot. This indeed suggests that one of the tasks of the inferior olive is to sample non-deterministically the input signals it receives to provide over time a complete representation of that signal to plasticity mechanisms at the Purkinje cells. Moreover, the results showed that little destructive interference occurred in learning the same task in different sensorimotor contexts, namely different overall target trajectories.

The robot arm we have used for the experiments had two non-stiff (compliant) joints controlled with low torque motors. In order to accurately control this platform it was necessary to build a predictive dynamics model of the arm. The cerebellum network essentially fulfilled this purpose.

Moreover, we implemented the delays in the sensorimotor pathways to evaluate the predictive strategy tested in this work. We implemented a STDP kernel filter that correlates the activity from the inferior olive (encoding the error using a probabilistic model) with the sensorimotor activity received through the parallel fibers. The correlation was done at the parallel fibers to Purkinje cells synaptic connections.

The experimental results showed how the cerebellum-based system was able to adapt dynamically to different contexts. Future work will test sensorimotor encoding strategies to learn multiple models and context switching mechanisms to choose optimal control action with minimal delays and relearning.

Acknowledgements

This work was supported by the EU project SENSOPAC (IST-028056) and the National Spanish Grant DEPROVI (DPI 2004-07032).

References

Andersen, B.B., Korbo, L., Pakkenberg, B., 1992. A quantitative study of the human cerebellum with unbiased stereological techniques. J. Comp. Neurol. 326 (4), 549–560.

Arbib, M.A., Schweighofer, N., Thach, W.T., Glencross, D.J., Piek, J.P., 1995. Modeling the cerebellum: from adaptation to coordination. Motor Control and Sensory-Motor Integration: Issue and Directions. Elsevier, Amsterdam, pp. 11–36.

Arnold, M., 2001. Feedback learning in the olivary–cerebellar system. Ph.D. Thesis. The University of Sydney.

Barbour, B., 1993. Synaptic currents evoked in Purkinje cells by stimulating individual granule cells. Neuron 11, 759–769.

Bastian, A.J., Zackowski, K.M., Thach, W.T., 2000. Cerebellar ataxia: torque deficiency or torque mismatch between joints? J. Neurophysiol. 83, 3019–3030.

Bezzi, M., Nieus, T., Coenen, O.J.M., D'Angelo, E., 2004. An integrate-and-fire model of a cerebellar granule cell. Neurocomputing 58–60, 593–598.

Boucheny, C., Carrillo, R., Ros, E., Coenen, O.J.-M.D., 2005. Real-time spiking neural network: an adaptive cerebellar model. Lecture Notes Comput. Sci. 3512, 136–144.

Brown, T.H., Kairiss, E.W., Keenan, C.L., 1990. Hebbian synapses: biophysical mechanisms and algorithms. Annu. Rev. Neurosci., 475–511.

Coenen, O.J.M.D., Arnold, M.P., Sejnowski, T.J., Jabri, M.A., 2001. Parallel fiber coding in the cerebellum for life-long learning. Auton. Robots 11 (3), 291–297.

Contreras-Vidal, J.L., Grossberg, S., Bullock, D., 1997. A neural model of cerebellar learning for arm movement control: cortico-spino-cerebellar dynamics. Learn. Memory 3 (6), 475–502.

D'Angelo, E., De Filippi, G., Rossi, P., Taglietti, V., 1995. Synaptic excitation of individual rat cerebellar granule cells in situ: evidence for the role of NMDA receptors. J. Physiol. (Lond.) 482, 397–413.

D'Angelo, E., Nieus, T., Maffei, A., Armano, S., Rossi, P., Taglietti, V., Fontana, A., Naldi, G., 2001. Theta-frequency bursting and resonance in cerebellar granule cells: experimental evidence and modeling of a slow K+-dependent mechanism. J. Neurosci. 21 (3), 759–770.

Eckhorn, R., Gail, A.M., Bruns, A., Gabriel, A., Al-Shaikhli, B., Saam, M., 2004. Different types of signal coupling in the visual cortex related to neural mechanisms of associative processing and perception. IEEE Trans. Neural Networks 15 (5), 1039–1052.

Eckhorn, R., Reitboeck, H.J., Arndt, M., Dicke, P., 1990. Feature linking via synchronization among distributed assemblies: simulations of results from cat visual cortex. Neural Comput. 2, 293–307.

Eskiizmirliler, S., Forestier, N., Tondu, B., Darlot, C., 2002. A model of the cerebellar pathways applied to the control of a single-joint robot arm actuated by McKibben artificial muscles. Biol. Cybernet. 86, 379–394.

Flash, T., Hogans, N., 1985. The coordination of arm movements: an experimentally confirmed mathematical model. J. Neurosci. 5 (7), 1688–1703.

Gerstner, W., Kistler, W.M., 2002. Spiking Neuron Models. Cambridge University Press.

Golgi, C., 1967. The neuron doctrine: theory and facts. In: Nobel Lectures: Physiology or Medicine, [1906], 1901–1921. Elsevier, Amsterdam, pp. 189–217.

Hofstötter, C., Mintz, M., Verschure, P.F.M.J., 2002. The cerebellum in action: a simulation and robotics study. Eur. J. Neurosci. 16 (7), 1361–1376.

Huang, J., Jabri, M.A., Coenen, O.J.-M.D., October 1998. Models of basal ganglia and cerebellum for sensorimotor integration and predictive control in real-time robot navigation. Laboratory Report. Sydney University.

Ito, M., Kano, M., 1982. Long-lasting depression of parallel fiber-Purkinje cell transmission induced by conjunctive stimulation of parallel fibers and climbing fibers in the cerebellar cortex. Neurosci. Lett. 33, 253–258.

Ito, M., 2001. Cerebellar long-term depression: characterization, signal transduction, and functional roles. Physiol. Rev. 81 (3), 1143–1195.

Kandel, E.R., Schwartz, J.H., Jessell, T.M., 2000. Principles of Neural Science. McGraw-Hill Professional Publishing, New York.

Keating, J.G., Thach, W.T., 1995. Nonclock behavior of inferior olive neurons: interspike interval of purkinje cell complex spike discharge in the awake behaving monkey is random. J. Neurophysiol. 73 (4), 1329–1340.

Kettner, R.E., Mahamud, S., Leung, H., Sittkoff, N., Houk, J.C., Peterson, B.W., Barto, A.G., 1997. Prediction of complex two-dimensional trajectories by a cerebellar model of smooth pursuit eye movement. J. Neurophysiol. 77 (4), 2115–2130.

Kuroda, S., Yamamoto, K., Miyamoto, H., Doya, K., Kawato, M., 2001. Statistical characteristics of climbing fiber spikes necessary for efficient cerebellar learning. Biol. Cybern. 84 (3), 183–192.

Lev-Ram, V., Mehta, S.B., Kleinfeld, D., Tsien, R.Y., 2003. Reversing cerebellar long-term depression. Proc. Natl. Acad. Sci. U.S.A. 100 (26), 15989–15993.

Maex, R., De Schutter, E., 1998. Synchronization of Golgi and granule cell firing in a detailed network model of the cerebellar granule cell layer. J. Neurophysiol. 80 (5), 2521–2537.

Mauk, M.D., Garcia, K.S., Medina, J.F., Steele, P.M., 1998. Does cerebellar LTD mediate motor learning? Toward a resolution without a smoking gun. Neuron 20, 359–362.

Medina, J.F., Mauk, M.D., 1999. Simulations of cerebellar motor learning: computational analysis of plasticity at the mossy fiber to deep nucleus synapse. J. Neurosci. 19 (16), 7140–7151.

Purves, D., Augustine, G.J., Fitzpatrick, D., Katz, L.C., Lamantia, A., McNamara, J.O., Williams, S.M., 2001. Neuroscience, 2nd ed. Sinauer Associates, Inc.

Ramón y Cajal, S., 1995. Histology of the Nervous System of Man and Vertebrates, vol. 1/2. Swanson, N., Swanson, L.W. (Trans.) [1909]. Oxford University Press, New York.

Raymond, J.L., Lisberger, S.G., 1998. Neural learning rules for the vestibulo-ocular reflex. J. Neurosci. 18 (21), 9112–9129.

Ros, E., Carrillo, R., Ortigosa, E.M., Barbour, B., Agís, R., 2006. Event-driven simulation scheme for spiking neural networks using lookup tables to characterize neuronal dynamics. Neural Comput. 18 (12), 2959–2993.

Ros, E., Ortigosa, E.M., Agís, R., Arnold, M., Carrillo, R., 2006. Real-time computing platform for spiking neurons (RT-Spike). IEEE Trans. Neural Networks 17 (4), 1050–1063.

Schoenauer, T., Atasoy, S., Mehrtash, N., Klar, H., 2002. NeuroPipe-Chip: a digital neuro-processor for spiking neural networks. IEEE Trans. Neural Networks 13 (1), 205–213.

Schweighofer, N., Arbib, M.A., Kawato, M., 1998a. Role of the cerebellum in reaching movements in humans. I. Distributed inverse dynamics control. Eur. J. Neurosci. 10, 86–94.

Schweighofer, N., Doya, K., Fukai, H., Chiron, J.V., Furukawa, T., Kawato, M., 2004. Chaos may enhance information transmission in the inferior olive. Proc. Natl. Acad. Sci. U.S.A. 101, 4655–4660.

Schweighofer, N., Spoelstra, J., Arbib, M.A., Kawato, M., 1998b. Role of the cerebellum in reaching movements in humans. II. A neural model of the intermediate cerebellum. Eur. J. Neurosci. 10 (1), 95–105.

Schweighofer, N., Doya, K., Lay, F., 2001. Unsupervised learning of granule cell sparse codes enhances cerebellar adaptive control. Neuroscience 103 (1), 35–50.

Shaefer, M., Schoenauer, T., Wolff, C., Hartmann, G., Klar, H., Rueckert, U., 2002. Simulation of spiking neural networks—architectures and implementations. Neurocomputing 48, 647–679.

Solinas, S., Maex, R., De Schutter, E., 2003. Synchronization of Purkinje cell pairs along the parallel fiber axis: a model. Neurocomputing 52–54, 97–102.

Spoelstra, J., Schweighofer, N., Arbib, M.A., 2000. Cerebellar learning of accurate predictive control for fast-reaching movements. Biol. Cybernet. 82, 321–333.

Steuber, V., De Schutter, E., Jaeger, D., 2004. Passive models of neurons in the deep cerebellar nuclei: the effect of reconstruction errors. Neurocomputing 58–60, 563–568.