



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/171926/>

Version: Accepted Version

Article:

Wilson, E.D., Assaf, T., Rossiter, J.M. et al. (2021) A multizone cerebellar chip for bioinspired adaptive robot control and sensorimotor processing. *Journal of The Royal Society Interface*, 18 (174). 20200750. ISSN: 1742-5689

<https://doi.org/10.1098/rsif.2020.0750>

© 2021 The Author(s). Published by the Royal Society. This is an author-produced version of a paper subsequently published in *Journal of the Royal Society, Interface*. Uploaded in accordance with the publisher's self-archiving policy.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

A Multizone Cerebellar Chip for Bioinspired Adaptive Robot Control and Sensorimotor Processing

Emma D. Wilson, Tareq Assaf, Jonathan M. Rossiter, Paul Dean, John Porrill,
Sean R. Anderson and Martin J. Pearson

Abstract

The cerebellum is a neural structure essential for learning, which is connected via multiple zones to many different regions of the brain, and is thought to improve human performance in a large range of sensory, motor and even cognitive processing tasks. An intriguing possibility for the control of complex robotic systems would be to develop an artificial cerebellar chip with multiple zones that could be similarly connected to a variety of subsystems to optimize performance. The novel aim of this paper, therefore, is to propose and investigate a multizone cerebellar chip applied to a range of tasks in robot adaptive control and sensorimotor processing. The multizone cerebellar chip was evaluated using a custom robotic platform consisting of an array of tactile sensors driven by dielectric electroactive polymers mounted upon a standard industrial robot arm. The results demonstrate that the performance in each task was improved by the concurrent, stable learning in each cerebellar zone. This paper, therefore, provides the first empirical demonstration that a synthetic, multizone, cerebellar chip could be embodied within existing robotic systems to improve performance in a diverse range of tasks, much like the cerebellum in a biological system.

1 Introduction

The importance of cerebellar function can be inferred from the estimate that the cerebellum contains 80% of neurons in the human brain [1]. Interestingly, a single cerebellar microcircuit, with virtually identical neuronal circuitry, is repeated throughout the entire cerebellar cortex [2]. Despite this uniformity the cerebellum is implicated in a multitude of diverse tasks. It is traditionally regarded as a motor control structure, capable of adaptively modulating motor commands by providing corrections required for accurate movements [2, 3, 4]. There is also mounting evidence indicating that the cerebellum is involved in sensory processing, sensory perception and cognitive functions (such as emotion and language) [5, 6, 7, 8].

The microcircuit regularity of cerebellar cortex implies that the same basic signal processing algorithm is implemented by each region of the cerebellum (a microzone), whether used for control of reflexive or voluntary movements, sensory noise cancellation, or higher functions such as language. The combination of general applicability with uniformity suggests that functional differences between cerebellar microzones emerge from differences in the input and output connectivity [9, 10, 2]. This has brought about the ‘chip’ metaphor of cerebellar organisation, where the function of each region depends on both the uniform internal algorithm implemented by all chips, and on the architecture in which the chip is embedded (e.g. external connections, which differ dependent on function). The observation that cerebellar lesions impair but do not abolish function indicates that the cerebellum is not the sole pathway subserving each particular function, suggesting that the cerebellum modulates a range of behaviours and generally optimises performance [11, 5]. Prior research has also shown that serial, or tandem learning enables efficient learning in the cerebellum [12].

Engineering control solutions are often designed on a case-by-case basis and optimised to a specific task. When controlling nonlinear, multi-degree of freedom, compliant, soft robots this is a non-trivial problem. Neural control strategies, implemented by structures including the cerebellum, have evolved

over long periods of time to deal with compliant, nonlinear materials such as muscle and biological systems are able to achieve remarkable levels of control performance despite using a relatively flat and homogenous structure. A synthetic, uniform cerebellar chip algorithm that could be plugged into existing systems to fine tune and improve the performance in a range of tasks, much like the cerebellum in the biological system, has great potential for robotic applications. Such an algorithm could greatly reduce initial control design effort and reduce the need for extensive offline system identification. This is because initial control structures would need only to provide an approximate solution (e.g. the brainstem in motor plant compensation only approximates the dynamics of the plant) and the adaptive cerebellar element could fine tune this control. Such a cerebellar chip would be especially useful in lightweight, multi-dimensional, anthropomorphic robots, where considerable control efforts are needed.

Therefore, an intriguing possibility for the control of complex robotic systems with multiple degrees of freedom is an artificial ‘cerebellar chip’ that could be plugged into existing control systems to fine-tune and improve performance in a wide range of tasks. Previous work in robotics has used cerebellar-inspired algorithms to provide adaptive solutions for robot control in single applications including variable stiffness, lightweight actuators with varying dynamics [13, 14, 15] especially in the context of robot arm control [16, 17, 18, 19, 20, 21, 22]. Cerebellar inspired models have also been applied in robotics to locomotion [23], collision, or obstacle avoidance, tasks [19, 24, 25, 26, 27], to gaze stabilisation tasks [19, 28, 29], to the adaptive cancellation of self-generated sensory signals [30], and to provide anticipatory control [31]. However, cerebellar-inspired control algorithms have not yet been tested through simultaneous application of the same microzone algorithm to a range of different tasks within a single robotic system. Even in cases when the same model has been applied to different sensorimotor tasks in a robotic system [30, 28, 19], this has not been done simultaneously.

The novel aim of this paper is to propose and investigate simultaneous learning in multiple microzones of cerebellar microcircuitry applied to different tasks in motor control and sensory processing within a robotic system. The impact of interactions between algorithms applied to each distinct task is investigated for the first time here. This paper, therefore, represents an essential step towards developing a cerebellar chip for robotic systems.

The cerebellar microzones were each based on the adaptive filter model of the cerebellum [32, 33], itself derived from the original Marr-Albus models [3, 4]. This algorithm is computationally powerful and is widely used in the analysis of cerebellar function [33, 34]. It can represent both forward and inverse models [9], and has been evaluated in a range of robotic tasks, including image stabilisation [28], reafferent noise cancellation [30] and both linear and nonlinear control of artificial muscle [35, 36].

The following three robot tasks were chosen to investigate: (i) control of an array of artificial whiskers using artificial muscle (trajectory control using motor plant compensation); (ii) removal of self motion or re-afferent noise from the sensory response of an array of active tactile whisker sensors (sensory noise cancellation); (iii) the calibration of a head centred map of sensory space to improve accuracy of directed motor commands toward points of interest in that map (sensorimotor map calibration). Our design principle, derived from the principle that the same cerebellar algorithm is used in a wide variety biological tasks, is that there were to be no ad-hoc changes to the algorithm internal circuitry used in each task.

A custom built robot platform (*‘Bellabot’*) [37] was used to evaluate the real world performance of the cerebellar chips. This platform consists of an array of tactile whisker-like sensors driven by artificial muscle actuators (Dielectric ElectroActive Polymers - DEAPs) mounted as the end-effector of a standard industrial robot manipulator. DEAPs have inconsistent, time varying, non-linear dynamics which present a number of control challenges.

The remainder of this paper is organised as follows. The following sections describe the robotic tasks, platform, and the adaptive filter model of cerebellar function including its application to these tasks. The results section provides data on the performance of the cerebellar adaptive filter applied to each task concurrently, as well as results showing how concurrent learning affects performance. The paper concludes with a discussion of the results.

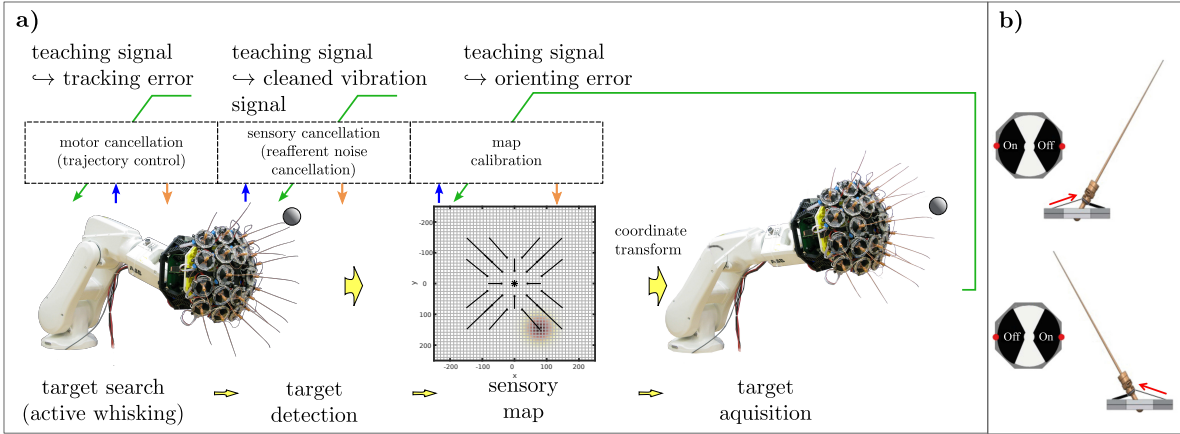


Figure 1. a) Summary of the basic robotic task. Each of the three areas (motor compensation, sensory cancellation and map calibration) where cerebellar control is applied are highlighted. The robotic platform (*Bellabot*) used to implement the algorithms is pictured, with a picture of it oriented to look at a target shown on the right of the figure. b) Whisker module activated by DEAPs. The figures on the left give a top down view of the whisker module and those on the right a side on view. Each whisker is activated by two separate actuators (shown in black in top down view). Activation of each side (indicated by on/off label) sequentially drives the whiskers back and forward. Arrows (shown in the side on view) indicate the direction of movement resulting from the activation of each actuator.

2 Methods

2.1 Robotic tasks

The three tasks (trajectory control using motor plant compensation, sensory noise cancellation, sensorimotor map calibration) were carried out using the robotic platform ‘*Bellabot*’ (details given in the subsequent section). These tasks were chosen as they are distinct, difficult tasks. They are modeled on the behaviour of whiskered rodents attending to points of salient contact made by their whiskers [38]. Challenges include, non-linearities and time variations in the motor plant and reafferent response, and mapping of points of whisker contacts to the true head centred topographic location.

The *Bellabot* platform used an array of artificial whiskers, mounted upon a manipulator, to detect novel tactile targets. The location of the detected targets was represented in a topographic map of whisker sensory space and then used to drive orienting movements towards the target. Accuracy of the orienting movements was assessed using a camera located at the centre of the whisker array. Orienting errors were used to update learning in a cerebellar module calibrating the topographic map. Additional cerebellar modules were used to control the trajectory of each whisker (modeled on rodent whisking behaviour [39]) and to remove re-afferent, self-induced noise signals from each whisker during target detection. A picture of the robot platform and summary of the three basic tasks in the context of detecting and orienting towards novel objects is given in Fig. 1a and the individual whisker module in Fig. 1b. To evaluate the performance of each adaptive filter over time data were obtained while the *Bellabot* cyclically performed four sequential behaviours: **explore**, **recoil**, **orient** and **reset**. The full protocol duration was 40 minutes, during which these four cyclic behaviours were repeated.

Before each trial a small ball (target), mounted upon a clamp stand, was placed in front of the platform (see left hand picture in Fig. 1a). This was placed approximately 200-300mm out, within a radius of 100mm from the centre of the camera. During the **explore** behaviour (average duration 4.5 secs/ cycle), the array of whiskers were actively driven with the noise cancelling and trajectory tracking chips active and learning.

The manipulator was moved forward until a contact was made by one of the whiskers (if no

contact was made the *Bellabot* was reset and the target re-placed). During the contact (i.e. when the deflection signal was above threshold) learning was gated, and the adaptive filter weights were not updated. Detected targets were read from the 2D topographic map of whisker sensory space.

After contact the **recoil** behaviour (average duration 8.3 secs/ cycle) was initiated, where the robot moved backwards a safe distance from the target. During the recoil behaviour active whisking continued, with the noise cancelling and trajectory tracking chips active and learning (to increase the time period for learning C_m and C_s weights - Sections 2.4, 2.5).

The head was then moved to **orient** (average duration 5.9 secs/ cycle) such that the centre mounted camera was directed towards the estimated target point in space as determined from the head centred topographic map of the whisker sensory space.

At the end of the orienting movement an image of the ball was captured from the camera. A coloured blue ball (that differed from the colours in the background) was used as a target. The centre of the target was then calculated from the image as the centre of mass of the blue area of pixels. This estimated target location was used to obtain the target acquisition error and update the weights of the map-calibration adaptive filters. During image capture the whiskers were held stationary and the noise cancelling and trajectory tracking adaptive filters were not active, and filter weights not updated. The *Bellabot* was then **reset** and the platform returned to its original start configuration before beginning another trial.

Between trials the contact ball was relocated pseudo-randomly such that different whiskers were contacted in subsequent cycles, with a similar number of total contacts being made on each whisker. To simplify the experiments, we only considered cases where the inner circle of eight whiskers were contacted (contacting just the inner eight whiskers was decided as a trade-off between testing the map calibration algorithms works and limiting the length of experiments to keep manageable). The behavioural cycle was repeated 85 times to obtain a rich set of data to calibrate the topographic map.

2.2 Robotic Platform

A custom built robotic platform was used to evaluate the real world performance of each task. The *Bellabot* (Fig. 1a) platform consisted of an array of 20 DEAP actuated whiskers surrounding a central camera, mounted on a 5 degree-of-freedom industrial manipulator (ABB - IRB120) (for further details see [40, 37]).

Each DEAP driven whisker was made of a single conical membrane with two carbon grease electrodes printed onto the membrane. The whisker actuation was constrained to 1-dof, and whisker deflections and rotations were measured using Tri-axis Hall effect sensors. Signals from these sensors were sampled at 500Hz and passed via USB to an external computer. Motor commands to move each whisker were relayed from this computer and converted to the high voltages required for DEAP actuation.

A standard USB camera mounted at the centre of the whisker array was used to capture a series of images at the end of an orienting movement to calculate errors in the x -, y - directions, i.e. the difference between the centre of the image (desired position) and the centre of the object actually oriented to (actual position).

The coordination of motor command and sensory response to and from the platform was maintained using the BRAHMS modular execution framework [41]. BRAHMS was developed specifically to integrate heterogeneous models of neural components by providing standardised interfaces and maintaining synchronous communication and execution across these models suitable for application in hard-real-time constrained robot control. The core of BRAHMS was written and compiled using C++, however, it has numerous language wrappers for users to adopt including, python, C# and MATLAB, and can be deployed using MPI for parallel compute environments. In this study the cerebellar micro-circuits were coded using C++ and the interface to robot control modules in C#, offline data analysis was performed using MATLAB. Raw sensory data gathered from the whisker array was sampled and marshaled at 500Hz before being downsampled to 25Hz for presentation to the modularized algorithms under test, the coordination for which is one of the native operations of BRAHMS.

2.3 Cerebellar Microcircuit

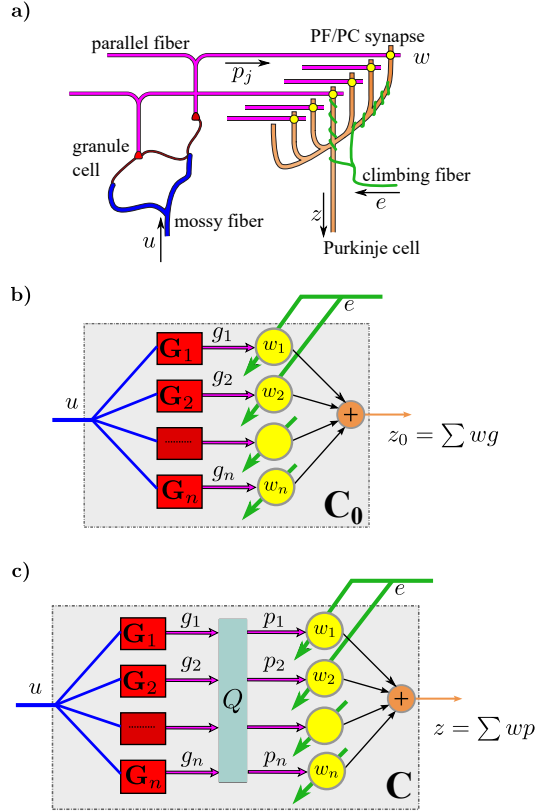


Figure 2. The cerebellar microcircuit as an adaptive filter. (a) Simplified cerebellar microcircuit. (b) Interpretation of the cerebellar microcircuit as an adaptive filter. Mossy fibre inputs (u) are analysed into component parallel fibre signals (g_i) in the granule layer by a bank of fixed filters. We use a bank of fixed alpha basis filters (see Eq. 6). These parallel fibre signals are weighted (w_i) and recombined to give the Purkinje cell output (z). (c) The cerebellar microcircuit as an adaptive filter with additional Q matrix applied to filter outputs (g_i) to approximately orthonormalise into parallel fibre signals (p_i).

In the basic cerebellar microcircuit (Fig. 2a) there are two input pathways (climbing fibres and mossy fibres) and a single output (Purkinje cells). In the adaptive filter model of cerebellar function (Fig. 2b), mossy fibre inputs are analysed into component parallel fibre signals which are weighted (parallel fibre - Purkinje cell synapses) and recombined to form the Purkinje cell output. The climbing fibre inputs are a teaching or error signal, used to train the weights of parallel fibre-Purkinje cell synapses [32, 33]. An extension to the adaptive filter model (Fig. 2c) used here, includes a fixed matrix Q to orthonormalise¹ signals and speed up learning.

In interpreting the cerebellar microcircuit as an adaptive filter, the time-varying signals carried by mossy fibre inputs to the cerebellum are represented at sample time T as $u(T)$ (for clarity only a single input is considered here, i.e. the number of mossy fibre inputs $n_u = 1$). These inputs are passed through a basis of n fixed filters G_i to produce signals $g_1(T), g_2(T), \dots, g_{n_p}(T)$ (where n_p is the number of parallel fibre signals, in this single mossy fibre input case $n_p = n$, for cases where there are more mossy fibre inputs $n_p = n \times n_u$). A fixed matrix Q is used to orthonormalise signals to produce

¹This orthonormalisation step is described in further detail in [36]. Although necessary to speed up learning, the use of a fixed matrix is not a biologically plausible mechanism and future work is required to establish how biology solves the problem of fast learning

parallel fibre signals $p_1(T), p_2(T), \dots, p_{n_p}(T)$. These parallel fibre signals synapse on Purkinje cells to produce the Purkinje cell output

$$z(T) = \sum_{i=1, j=1}^{n_p} w_i Q_{i,j} G_j u(T) = \sum_{i=1}^{n_p} w_i p_i(T) = \mathbf{w}^T \mathbf{p} \quad (1)$$

where w_i gives the weight of the i^{th} parallel fibre-Purkinje cell synapse and $Q_{i,j}$ is the (i, j) element of matrix Q , additionally \mathbf{w} and \mathbf{p} denote a vector of weights and parallel fibre signals respectively, where $\mathbf{w} = [w_1, w_2, \dots, w_{n_p}]$ and $\mathbf{p} = [p_1(T), p_2(T), \dots, p_{n_p}(T)]$.

The climbing fibre input is interpreted as a teaching, or error signal (e) that adapts these synaptic weights using the decorrelation learning rule. This is a computationally powerful rule for supervised learning [42] that minimises the mean squared error and is guaranteed to converge to the optimal solution for an appropriate error signal [43]. It is equivalent to the least mean squares (LMS) rule from adaptive control theory [44]. Filter weights are adjusted by,

$$\delta w_i = -\beta e p_i \quad (2)$$

where β is a positive learning rate parameter.

This basic adaptive filter interpretation of the cerebellar microcircuit is used for each of the three robotic tasks (trajectory control, noise cancellation, map calibration), with each particular task specified by the external connections of the adaptive filter circuit (section 2.7).

2.4 Trajectory control

The biohybrid scheme for controlling the whisker rotation angle (ψ_a) uses a form of motor plant compensation based on that implemented by the vestibular ocular reflex (VOR), and shown in Fig. 3a. It extends the basic VOR circuit by adding a reference model (\mathbf{M}) to specify the desired response of the controlled system, which enables the plant-compensation algorithm to function independently of the order of the plant [36]. The reference signal input to the trajectory control scheme is given as r (Fig. 3a) and the aim is to track the signal d_a , which is given as the signal r filtered through \mathbf{M} .

An individual whisker was driven by two separate DEAP actuators to give a one degree of freedom whisker angle response, meaning that each whisker was an over-actuated, redundant system. To simplify the redundancy problem we coupled the inputs to the two actuators, so they were specified by a single variable. A fixed tonic input, c_0 ($c_0 = 0.5$) was sent to each actuator, and the driving signal input to each actuator set to $c_0 - u$, and $c_0 + u$ where u is the motor command (Fig. 3a).

The mean of the whisker angle signals was tracked and removed to give the controlled whisking angle (ψ_a), so in effect just the amplitude and frequency of whisking were controlled. This was done to compensate for asymmetry and initial offsets in the whisker responses.

As indicated in Fig. 3a, plant compensation was provided by both a fixed element \mathbf{B} (corresponding to the brainstem in oculomotor plant compensation), and a recurrently connected adaptive filter element \mathbf{C}_m (corresponding to the cerebellum). \mathbf{B} was designed to provide approximate, under-gained compensation of the average whisker angle plant (\mathbf{P}_a) and \mathbf{C}_m used to fine tune this control (further details in [36]). The fixed filter \mathbf{B} was defined in discrete time as,

$$\mathbf{B}(q, \gamma) = \frac{b_0 - b_1 q^{-1} - b_2 q^{-2}}{1 - a_1 q^{-1} - a_2 q^{-2}} \quad (3)$$

where q^{-1} is the backward shift operator and γ a vector of filter parameters, here $\gamma = [b_0 \ b_1 \ b_2 \ a_1 \ a_2] = [0.012 \ 0.015 \ -0.0033 \ 1.64 \ -0.65]$. These parameters were estimated by performing a least squares fit to input (u), output (ψ_a) data collected from all whiskers. Here, the whisker input signals were coloured noise, generated by low-pass filtering pseudo-random white noise with range 0-1 using a 1Hz filter cut-off. This fit gave an approximate estimate of the plant \mathbf{P}_a . The brainstem parameters were then obtained as the inverse of \mathbf{P}_a , scaled by 0.65 to be under-gained.

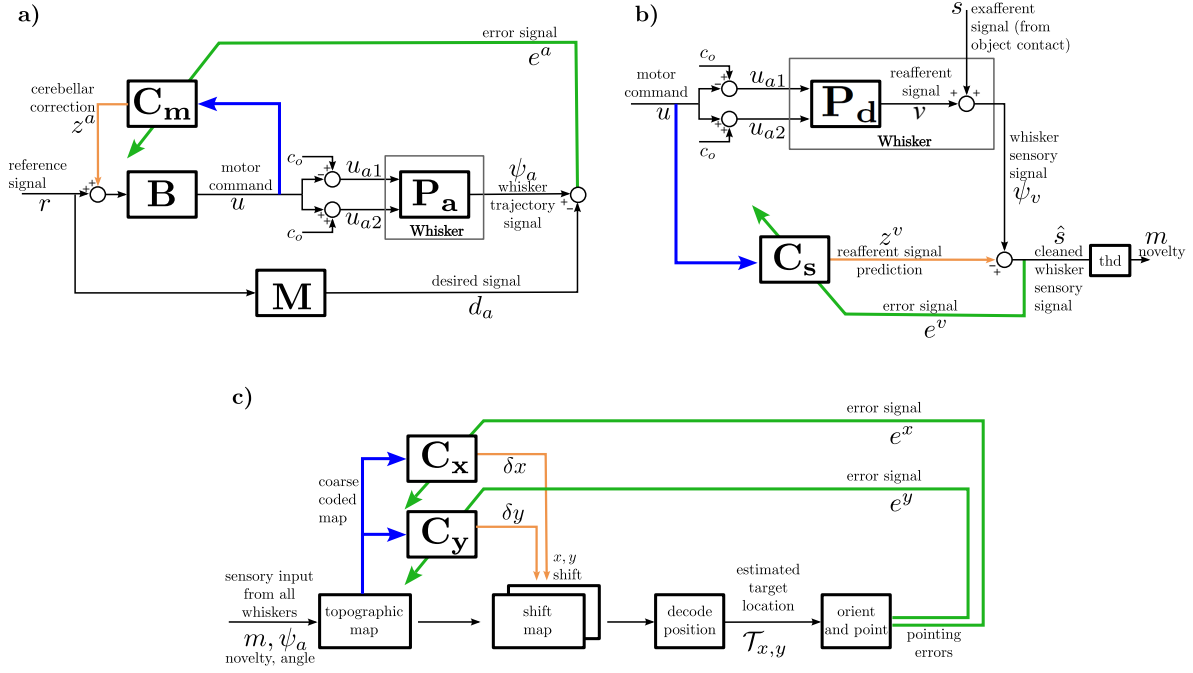


Figure 3. Architectures for implementing different tasks. To split the motor command (u) into signals to drive each actuator a fixed tonic input c_0 ($c_0 = 0.5$) was used and the input to actuator 1 given as $u_{a1} = c_0 - u$ and to actuator 2 as $u_{a2} = c_0 + u$. (a) Trajectory control scheme (based on the VOR). Control of the whisker angle (where P_a is a plant that describes the whisker angle response to motor inputs) is provided by a combination of a fixed feedforward brainstem element (B), and recurrently connected cerebellar element (C_m). A reference model (M) is used to specify the desired behaviour of the controlled plant. (b) Novelty detection scheme with noise canceling adaptive filter. The adaptive cerebellar filter (C_s) learns to remove components of the whisker deflection signals (ψ_v) that are correlated with motor commands (refferent noise cancellation). P_d is a plant that describes the whisker deflection response to inputs. (c) Map calibration control scheme. The accuracy of a 2D topographic, tactile map of the environment is improved using two adaptive cerebellar filters. The adaptive filters C_x and C_y learn to shift the map in the x - and y - directions respectively to improve the accuracy of orients to novel events in the map.

The reference model \mathbf{M} , used to specify the behaviour of the controlled plant was

$$\mathbf{M}(q, \gamma) = \frac{\rho_0}{1 - \alpha_0} \quad (4)$$

where $\gamma = [\rho_0 \ \alpha_0] = [0.33 \ 0.67]$ (for further details on the use of a reference model to extend the VOR control scheme see [36]).

The cerebellar element \mathbf{C}_m used to fine tune control was,

$$\mathbf{C}_m(q, T_i) = \sum_{i=1, j=1}^{n_p} w_i^a Q_{i,j}^a G(q, T_i) \quad (5)$$

where the superscript a is used to indicate the general signals introduced in Section 2.3 here correspond to trajectory tracking task signals. Within \mathbf{C}_m the fixed filters G (Fig. 2b,c) were implemented by two alpha functions (critically damped 2nd order filters) with time constants of 0.05, and 0.5 sec. These alpha functions are defined in discrete time as,

$$G_i(q, T_i) = \frac{\frac{dt^2}{T_i^2}}{1 + \left(2\frac{dt}{T_i} - 2\right) q^{-1} + \left(\frac{dt^2}{T_i^2} - 2\frac{dt}{T_i} + 1\right) q^{-2}} \quad (6)$$

where G_i is the i^{th} basis filter, T_i the time constant, and dt the sampling period (0.04s). The fixed matrix Q^a is used to speed up learning (where the vector of optimised outputs $\mathbf{p}^a = Q^a \mathbf{g}^a$). It is given as

$$Q^a = \begin{pmatrix} -0.1036 & 0.0056 \\ 0.0652 & 1.2019 \end{pmatrix} \quad (7)$$

and was designed off-line to exactly orthonormalise the filtered brainstem output when there is no cerebellar contribution (for further details on the design of this matrix see [36]).

A major advantage of the recurrent architecture shown in Fig 3b is that it allows sensory errors to be used to drive adaptation of the adaptive filter weights directly, rather than requiring an estimate of the unobservable motor errors [45]. The learning rule to estimate adaptive filter weights (from initially zero) is

$$\delta \mathbf{w}^a = -\beta e^a \bar{\mathbf{p}}^a \quad (8)$$

where $\bar{\mathbf{p}}^a$ is the vector of optimised filter outputs (optimised parallel fibre signals) filtered through reference model \mathbf{M} , \mathbf{w}^a a vector of corresponding weights (i.e. parallel fibre-purkinje cell synapses), for the trajectory control task the number of parallel fibre signals per whisker was $n_p = 2$. In equation (8) β is the learning rate ($\beta = 5$) and e^a is the sensed tracking error, and given as the difference between measured whisker angle (ψ_a) and reference model output (d_a),

$$e^a = d_a - \psi_a \quad (9)$$

2.5 Self noise cancellation

Movement of the robot whiskers generated sensory signals in the whiskers that could mask external signals (the reafference problem). A noise cancelling scheme (similar to that described in [30]) used the architecture shown in Fig. 3b to predict the sensory consequences of whisker movement and remove them from whisker deflection measurements to improve the accuracy of novelty detection.

Whisker deflection signals were measured in 2-dimensions orthogonal to the whisker shaft. These vector signals were combined to give the overall magnitude of whisker deflection (ψ_v) for each whisker. An adaptive cerebellar filter was used to predict the whisker deflections resulting from self-movement

(i.e. \mathbf{C}_s adapted to approximate \mathbf{P}_d) and this prediction was subtracted from the actual signal to provide an estimate of the whisker signals generated by the external world (e.g. contacts with objects).

In this scheme the cerebellar adaptive filter (\mathbf{C}_s) also contained a bank of two alpha basis filters with time constants 0.05 and 0.5 seconds. A matrix (Q^v) was again used to speed up the rate of learning, given as

$$Q^v = \begin{pmatrix} -0.1015 & -0.0169 \\ -0.0672 & 0.4049 \end{pmatrix} \quad (10)$$

this matrix was again estimated from the filter outputs to the motor command input u prior to any learning. The optimised parallel fibres in the noise cancelling adaptive filter are given as $\mathbf{p}^v = Q^v \mathbf{g}^v$.

The weights of the noise cancelling adaptive filter were learnt using the decorrelation learning rule to remove correlations between the estimated ex-afferent signal (\hat{s}) and the parallel fibre signals (\mathbf{p}^v).

$$\delta \mathbf{w}^v = -\beta e^v \mathbf{p}^v \quad (11)$$

where the superscript v corresponds to the noise cancellation task signals, \mathbf{w}^v is a vector of the weights of the noise cancelling adaptive filter, β the learning rate ($\beta = 5$), and e^v the error or teaching signal which is equal to the estimated ex-afferent signal ($e^v = \hat{s}$). For the self noise cancellation task, the number of parallel fibres per whisker was $np = 2$.

A simple threshold element (thd in Fig. 3b) was used to detect novelty. Signals were said to include novel contacts if the estimated ex-afferent signal (\hat{s}) was above threshold for a number of samples. Novel contacts were written into a topographic map and used to drive orienting movement to the estimated contact location.

2.6 Map calibration

In the third task an adaptive cerebellar filter was required to calibrate a 2D topographic map of the whisker sensory space that was employed to drive orienting movements to novel targets. The topographic map was constructed by using a 2D Gaussian function (with centre placed at the assumed tip of the contacted whisker) to provide a probabilistic representation of the location of detected targets (Fig. 4a). This location was used to drive the robot orienting movement. At the end of the movement the target acquisition error (obtained from the camera) was used to drive map calibration. Errors in the target position were only provided in two dimensions (in-plane with the camera), so the estimated perpendicular distance to a detected target was fixed at the whisker tip.

The sensory topographic map was initially miscalibrated by artificially adjusting whisker locations to give a distorted sensory map, and the cerebellar algorithm was required to compensate for these miscalibrations and reduce target acquisition errors. The map-calibration architecture is shown in Fig. 3c. Two cerebellar microzones (\mathbf{C}_x and \mathbf{C}_y) were used to shift the map in the x - and the y -directions, so for a given sensory map with a target centre at the 2D location (x, y) , the cerebellar bias $\mathbf{z} = (dx, dy)$ will shift the map so the target has centre $(x + dx, y + dy)$. The output from each cerebellar filter effectively slides the map activity across by an amount $\mathbf{z} = (dx, dy)$, where the outputs from each filter are dx and dy .

The inputs to \mathbf{C}_x and \mathbf{C}_y were a normalised, coarse coded representation of the topographic map (see Fig. 4), and the error signals (equivalent to climbing-fibre signals) carried information about movement inaccuracy. For this application the fixed filters G_i in \mathbf{C}_x and \mathbf{C}_y were represented by a single unitary gain, so their outputs (equivalent to parallel-fibre signals) were identical to their inputs (equivalent to mossy-fibre signals). Only a single unitary gain is required as the target was moved in a pseudo random way between contacts, meaning no predictive information is required. The significance of the difference in fixed filters (G_i) used here in comparison to in the trajectory control and noise cancellation schemes is addressed in the discussion.

Both \mathbf{C}_x and \mathbf{C}_y shift the position of the peak map activity, with the total bias calculated as,




Figure 4. Example of coarse coding for map calibration. a) Example of 2D topographic map providing a probabilistic representation of the location of a detected target. b) Coarse-coded, normalised version of the topographic map, here an 8x8 grid is used to represent the coarse coded image. c) Example of estimated x, y weights for the i^{th} parallel fibre.

$$\delta x = \sum_{i=1}^{n_p} w_i^x p_i^t \quad (12)$$

$$\delta y = \sum_{i=1}^{n_p} w_i^y p_i^t \quad (13)$$

where p_i^t is the intensity of the i^{th} pixel and $\mathbf{p}^t = [p_1^t, p_2^t, \dots, p_{n_p}^t]$ is a vector of the coarse-coded, normalised parallel fibre representation of the topographic map. Here, the number of parallel fibre signals are determined from the size of the coarse coded image (e.g. for Fig. 4 $n_p = 8^2$). For the map calibration task, the number of parallel fibre signals $n_p = 64$. The weights w_i^x and w_i^y are the weights for the i^{th} parallel fibre which are used to shift map in x and y directions respectively and learnt over trials (from initially zero). We assume the error signal inputs to the adaptive filter model carry information related to the x, y components of target acquisition error so the learning rule for estimating these weights can be written as

$$\delta \mathbf{w}^x = -0.1\beta e^x \mathbf{p}^t \quad (14)$$

$$\delta \mathbf{w}^y = -0.1\beta e^y \mathbf{p}^t \quad (15)$$

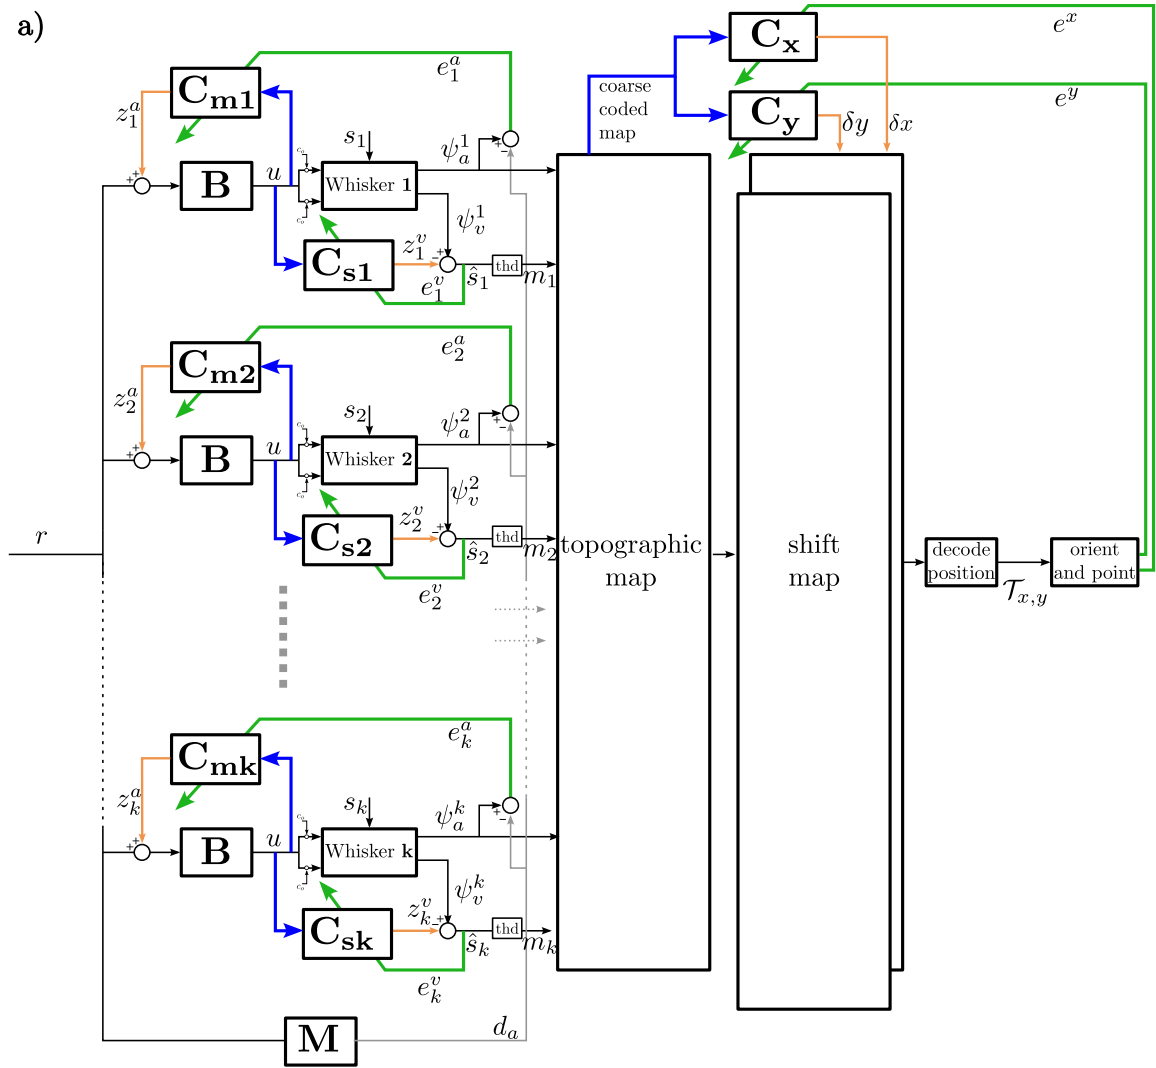
where \mathbf{w}^x and \mathbf{w}^y are the vectors of the weights used to shift the map in the x - and y - directions respectively, 0.1β is the learning rate (a factor of 0.1 is included in the learning rate, so that here the learning rate is 10 times smaller than the learning rate used for trajectory control and noise cancellation). The values e^x and e^y are the target acquisition errors in the x and y directions respectively.

2.7 Chip Connectivity and Interactions

The architecture for controlling the three tasks concurrently is shown in Fig. 5a. Each of the 20 whiskers on the robotic platform had two cerebellar chips, one for trajectory control and one for sensory noise cancellation. Two additional chips were used to shift the map in the x, y directions to correct for errors in orienting towards novel targets.

The output of the brainstem (corrected using the trajectory controller microzones \mathbf{C}_{mk}) feeds into the microzones for self noise cancellation (\mathbf{C}_{sk}), the cleaned deflection signal then feeds into the mapping algorithm. In other-words trajectory control is followed by self noise cancellation which is followed by map calibration (as shown in the simplified diagram of Fig. 5b).

A simulation was used to assess the impact of concurrent learning on the performance in each task. Contacts to just the inner eight whiskers were simulated to match experimental conditions. The location of each whisker was simulated using the base geometry of the whisker head, whisker geometry and the current whisking angle. The whisking angle was calculated directly from the motor command (whisker input) using the following,



b)




Figure 5. Schematic of the overall control architecture. a) The architecture for implementing the cerebellar algorithm simultaneously in three distinct areas is shown for k whiskers. For trajectory control and sensory noise cancellation, $2 \times k$ separate cerebellar microzones are used. The microzones for trajectory control and sensory noise cancellation corresponding to the k^{th} whisker are denoted C_{mk} , and C_{sk} respectively. The outputs write into a single topographic map. Two microzones (C_x , C_y) are used to shift the map in the x, y directions to correct for errors in orienting to novel targets. b) Pipeline of tasks that the cerebellar algorithm is applied to.

$$\psi_a^j(T) = b^j u(T) - a^j \psi_a^j(T-1) \quad (16)$$

Where ψ_a^j is the angle of whisker j at sample time T , u^j the motor command (or input) to whisker j and $a \sim N(300, 100)$, $b \sim N(10, 1.1)$ model parameters for each whisker.

The motor command was also used to calculate the whisker deflection signal. This was done using the following,

$$\psi_v^j(T) = v_s^j(T) + v_n^j(T) \quad (17)$$

$$v_s^j(T) = rs u(T) \quad (18)$$

$$v_n^j(T) = \begin{cases} 40 rn, & \text{if } rt > 0.98 \\ 0.01 rn, & \text{otherwise} \end{cases} \quad (19)$$

where w_v^j is the whisker deflection signal for each whisker j , and rs , rt , rn are random values generated at each step T and for each whisker j . rs is a pseudo random value drawn from the standard uniform distribution over the range (150,300), rt is a pseudo random value drawn from the range (0,1), and rn is drawn from the normal distribution $rn \sim N(0, 1)$.

In the simulation, up to 300 targets, each contacting the tip of one of the eight whiskers randomly, were simulated over a 30 minute period. The duration of each contact was 2 seconds, and the contact was simulated through changes to the whisker angle and whisker deflection signals. The angle of the contacted whisker (ψ_a^j) was fixed to that at the start of the contact time. The deflection signal of the contacted whisker (ψ_v^j) was set to a fixed value of 120.

3 Results

Experimental results from the *Bellabot* platform for each task are first given, followed by simulated results on the interactions between microzones.

3.1 Whisker trajectory control

The input reference signal, r , to each whisker controller was a sinusoid of frequency 1Hz and amplitude 5 degrees. Before learning, the response varied significantly between whiskers (Fig. 6a), likely as a result of both wide manufacturing tolerances leading to differences in response from each whisker, and of differences in the locations and orientations of the whiskers upon the robot head with respect to gravity. The cerebellar algorithm was able to compensate for these differences between the whiskers, so that at the end of learning the responses of each whisker were very similar and closely matched the desired response (Fig. 6b).

Initial and final responses of a representative individual whiskers are shown in Figs 6c and 6d, indicating how the algorithm compensated for an initially under-gained and slow response.

The evolution of adaptive filter weights over time for the representative whiskers are shown in Fig. 6f. For the first 500 sec the weights adapt relatively quickly to minimise the tracking error, then change more slowly presumably to track changes in the whisker dynamics caused by the time varying properties of the DEAP actuators ('creep'). Again, this result is typical of each whisker.

Root mean square (RMS) errors, where the error is the difference between the desired and actual whisker angle response (Eq. (9)), are shown for each whisker in Fig. 6f. The RMS errors, and the changes in errors over time, are similar for each whisker. On average, the RMS tracking errors are reduced by 87% when comparing errors during the first 60secs of the trial with errors during the final 60secs.

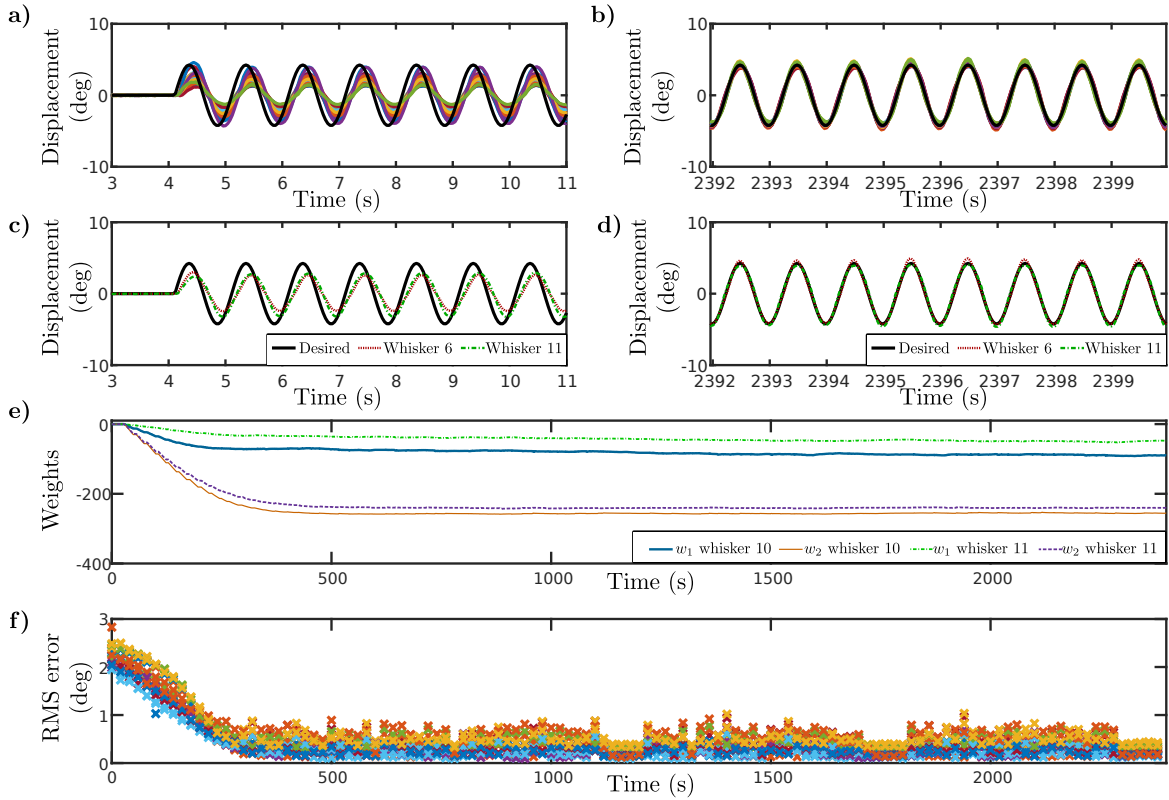


Figure 6. Whisker trajectory control using the cerebellar algorithm. The trajectory of each of the 20 whiskers was controlled to follow a 1Hz sine wave. No data are given for whisker 13 as this whisker was damaged so inactive during the tasks. a) Trajectory of each DEAP whisker before learning, compared with the desired trajectory (—). b) Trajectory of each DEAP whisker after learning, compared with the desired trajectory (—). c) The desired and actual whisker trajectories for two example whiskers (whisker 10 and 11) before learning. d) The desired and actual whisker trajectories after learning. e) Example of learnt adaptive filter weights for two whiskers (whisker 10 and 11). Here, the weights w_1 and w_2 correspond to the weighting applied to the parallel fibre signals p_1 and p_2 respectively. f) Windowed RMS error for trajectory tracking for every whisker. Errors are averaged over a sliding window with length 20secs.




Figure 7. Whisker re-afferent noise cancellation. a) The variance of the whisker deflection signal (ψ_v) during active whisking before (\cdot) and after ($\dot{\cdot}$) noise cancellation is plotted for every whisker. No data are plotted for whisker 13, as this whisker was damaged and not active during the tasks. b) Comparison between percentage change in signal variance after re-afferent noise cancellation and the maximum value of the cross correlation between the whisker driving signal and total whisker deflection signal for each whisker (indicated by no.). The reduction in signal variance after re-afferent noise cancellation is increased in cases where there are stronger correlations between the whisker drive signal (u) and whisker deflection signal (ψ_v)

3.2 Noise Cancellation

In parallel to trajectory control, an additional cerebellar adaptive filter for each whisker was used to cancel reafferent noise signals. Fig 7a shows the variance of the deflection signal from each whisker (ψ_v) during active whisking (and no external contacts) with and without noise cancellation. For the majority of whiskers the signal variance increased over time in the absence of noise cancellation, probably because the whisker itself moved more (Fig 6) as a result of adaptive trajectory control. A similar increase is seen right at the start for some whiskers when noise-cancellation is used, but is then often followed by a decrease.

On average, the noise cancelling adaptive filter gives a 32% reduction in noise variance in the last 60secs of the trial in comparison to the first 60secs across all whiskers. The noise cancelling adaptive filter makes a significant improvement and reduced the variance of the deflection signal, by up to 80% in the cases where the final variance without noise cancellation was large (e.g. whiskers 6, 11, 19) (Fig 7a, b). For cases when the final variance without noise cancellation was already small (e.g. whiskers 3, 4, 5, 19, 12) the noise cancelling adaptive filter did not significantly reduce the noise further (Fig 7a, b). This was because the noise cancelling adaptive filter only reduced noise that was correlated with the parallel fibre signals (i.e. filtered motor commands).

When the maximum cross-correlation between the motor command (u) and the whisker deflection signal (ψ_v) is plotted for each whisker against the change in variance caused by the noise cancelling adaptive filter (Fig 7b), it can be seen that when the motor command and deflection signal are well correlated the noise-cancelling adaptive-filter gives a significant reduction in signal variance, but when signals are not well correlated the change in variance is small.

3.3 Map Calibration

During the **explore** phase of testing (Methods), detected contacts with the target were written into a topographic map of the sensory space, and subsequently used to drive the robot’s orienting response. The topographic map was initially miscalibrated (Fig. 8a-c) by artificially adjusting the actual whisker positions (Fig. 8a). The results shown in Fig. 8d indicate how the weight vectors (x - horizontal, y - vertical) for each cerebellar chip changed as the whiskers made contact with the target in order to reduce subsequent orienting errors.

The vector combination of the x - and y - weights associated with each coarse coded parallel fibre at the end of learning is shown in Fig. 8e. The vectors show how the initial, erroneous map is distorted into a better representation of the actual whisker map. As a result orienting errors are on average reduced (Fig. 8f). Map calibration gives an 82% reduction in average orienting errors.

Fluctuations between the errors of individual contacts are caused by (i) random changes in the location of the ball between contacts in different areas of the sensory map, (ii) the fact that the ball has finite size and the actual location of the whisker upon the ball’s surface is unknown and varies, and (iii) the fact the location of the contact upon the whisker shaft is unknown and varies. For three of the contacts the target was incorrectly identified, meaning the camera estimate of the target location was incorrect, resulting in artificially large errors. These errors are omitted from the results shown in Fig. 8f.

The learning rate was set low enough such that the occasional artificially large error did not have much influence on the learned weights. Estimates of the contact locations over time, compared to the actual position of the whisker tips (Fig. 8g) show that the adaptive cerebellar filters for map calibration learnt to place the contacts near the actual whisker locations.

3.4 Interactions between microzones

All tasks were tested concurrently by implementing in the robot system. A simulation of the system was used to investigate the affect on learning and performance of coupling between microzones.

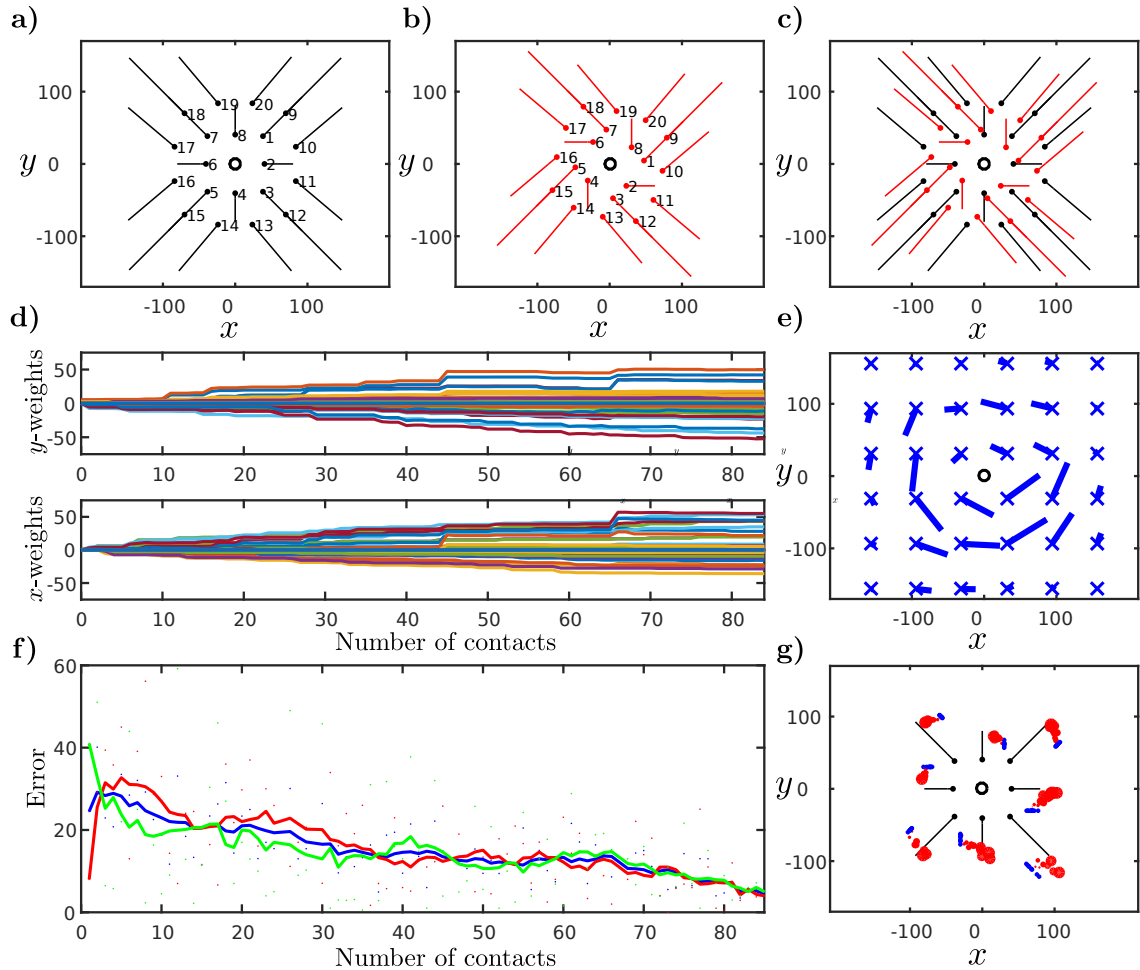


Figure 8. Map calibration using the cerebellar algorithm. The algorithm was used to correct for a miscalibrated map of sensory space. Whisker locations were artificially adjusted to give a distorted sensory map and the cerebellar algorithm used to compensate for these miscalibrations. a) Estimated x -, y - position of each actual whisker on the robot head. The centre of the head is plotted at the origin (o) and the location of the base of each whisker shown as a dot. b) Miscalibrated x - y - whisker positions. c) Comparison of the estimated actual (—) and miscalibrated (—) whisker locations. d) Learned weights in the horizontal (x -weights) and vertical (y -weights) directions for each contact. Weights were learned using the error in orienting to a contact, with a separate set of weights used to compensate for errors in the x - and y - directions. e) Learned weights in the sensory head space. The learnt weights (x - y - weights combined into given vector) at each coarse coded parallel fibre location (x) after learning are shown. f) Change in errors as the number of contacts increases. The x -errors (•), y -errors (•) and combined average errors (•) for each contact number are plotted. The solid lines show a 5-point moving average of the error. In all cases the error decreases over-time as the cerebellar algorithm learns to compensate for miscalibrations in the sensory map. g) Estimate of whisker contact locations over time. The estimated actual whisker locations (of the inner circle of eight whiskers – whiskers 1-8) are plotted in black, the estimated location of contacts when using the miscalibrated map given in b) are plotted in blue, and the learnt corrected location of these contacts plotted as red circles (where the size of the red circle increases with the number of contacts). The algorithm learns to place the contacts at the estimated actual whisker positions over time.

3.4.1 Effect on Map Calibration

In the overall architecture, self noise cancellation precedes map calibration (see Fig. 5). To demonstrate the effect of self noise cancellation on map calibration, the scenarios summarised in Table 1 were simulated. Results from these simulations are shown in Fig. 9. This shows how self noise cancellation affects the map calibration results when noise simulating false contacts is applied. When there is no noise simulating false contacts the map calibration results are good: RMS errors are small throughout (Fig 9b), learned parallel fibre weights shift the map to counteract miscalibrations (Fig. 9c), and as contacts increase weights change smoothly to reach a steady state value (Fig. 9f). When noise simulating false contacts is included, and map calibration uses the error signal directly, results are poor: RMS errors are significant (Fig. 9b), learned parallel fibre signals shift the map in a seemingly random way (Fig. 9d) and as the number of contacts increases the learned weights seem to change in a random way and do not reach a steady state as the number of contacts increase (Fig. 9g). The results, when noise simulating false contacts is included, are improved if the sign of the error signal is used for map calibration (this prevents large random errors skewing the results): RMS errors are generally small with occasional high errors caused by the included noise (Fig. 9b), learned parallel fibre weights shift the map to counteract miscalibrations (Fig. 9e), as contacts increase weights change to reach an approximate steady state (Fig. 9h – note that the noise in weight change is caused by the additional simulated noise)

Scenario	False contact noise	Self-noise cancellation	Map calibration
1	No	No learning	Learning using error signal
2	Yes	Learning using error signal	Learning using error signal
3	Yes	Learning using error signal	Learning using sign of error

Table 1: Summary of scenarios simulated to show the affect of self-noise cancellation on map calibration

If there are many false positive contacts (which can be caused by self-noise cancellation learning not being complete), then these can cause large errors and affect the learned weights for map calibration. This is demonstrated in the poor results shown in Fig. 9d,g). Using the sign of the error signal in the map calibration learning rule prevents these large errors from occurring, so gives better results (Fig. 9e,h). These results are close to the ideal case when there is no noise to simulate false contacts.

3.4.2 Effect on Self-Noise Cancellation

In the overall architecture, trajectory control precedes self noise cancellation (see Fig. 5). The motor command (that is modified by the trajectory control microzone) directly generates the self noise signal, as well as the parallel fibre signals for self noise cancellation. To demonstrate the effect of trajectory control on self-noise cancellation, the scenarios summarised in Table 2 were simulated. Results from these simulations are shown in Fig 10.

When the self-noise learning rate is set to zero, the level of self-noise still varies as the trajectory is learned (see Fig. 10b,c). This is because the motor command is increased during trajectory control, the motor command is also used to generate the self-noise, so the level of self noise is increased. The trajectory-controller affects the magnitude of the signals in the learning rule, so affects the speed of learning when self noise cancellation is included. This is shown in the increased fluctuations in weights (which converge around similar values) when trajectory control learning is included (see Fig. 10e). This is because, trajectory learning affects the magnitude of the parallel fibre signals that are used in the learning rule for self noise cancellation (increasing PF signals in this case). If the learning rate is normalised by the magnitude of the parallel fibre signals (see Fig. 10f) then this ensures there is no

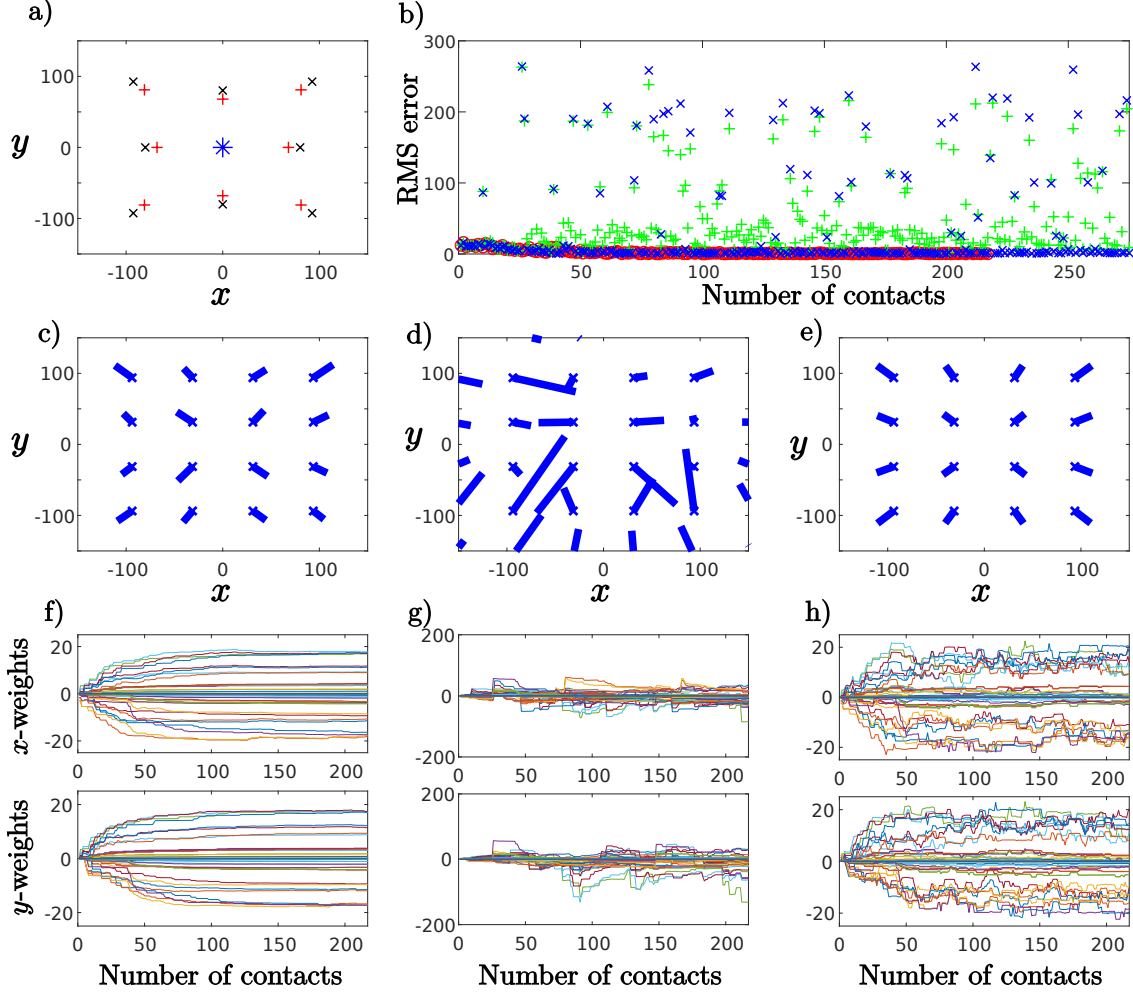


Figure 9. Simulated results from scenarios in Table 1 showing the affect of noise cancellation on map calibration. a) Actual (x) and misscalibrated ($+$) tip locations; b) RMS errors for scenario 1 (\circ - note errors small, so hard to see), scenario 2 ($+$), and scenario 3 (\times); c) learned weights for scenario 1; d) learned weights for scenario 2; e) learned weights for scenario 3; f) learned weights over time for scenario 1; g) learned weights over time for scenario 2; h) learned weights over time for scenario 3.

Scenario	Trajectory control	Self-noise cancellation
1	No learning	Learning using sign of error/ no learning
2	Learning using sign of error	Learning using sign of error/ no learning
3	Learning using sign of error	Learning using sign of error, rate normalised by magnitude of PF signals/ no learning

Table 2: Summary of scenarios simulated to show the affect of trajectory control on self-noise cancellation

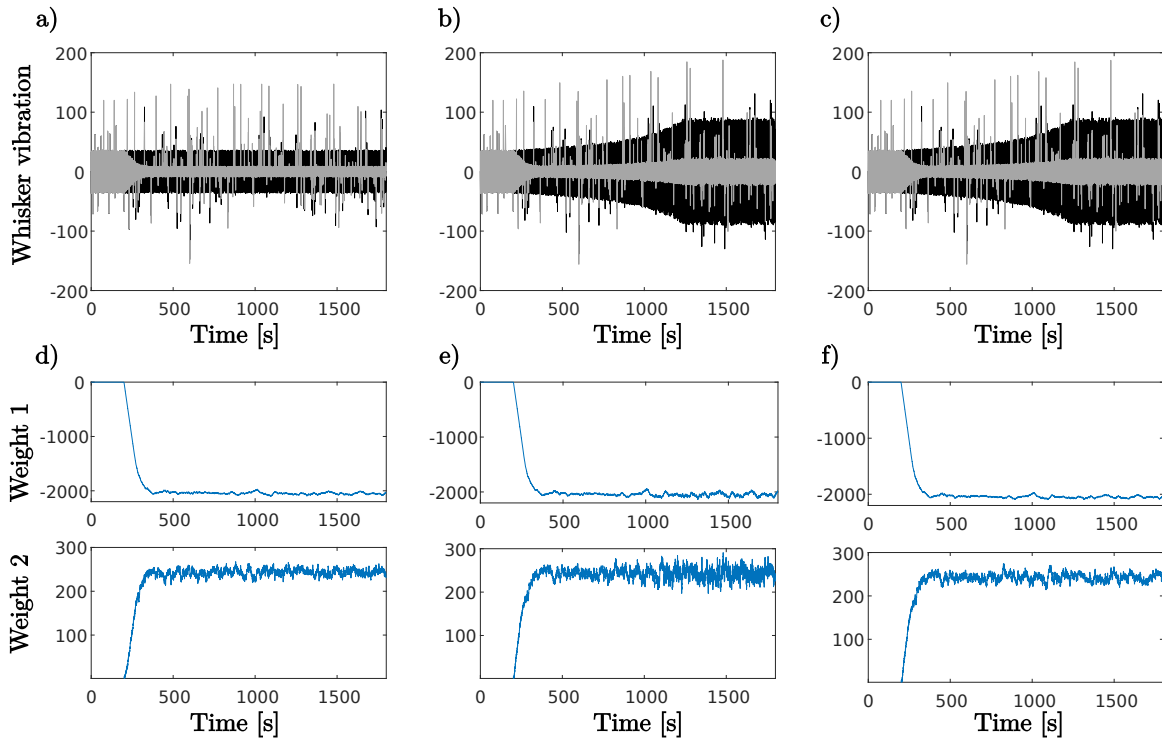


Figure 10. Simulated results from scenarios in Table 2 showing the affect of trajectory control on self-noise cancellation. a-c) Deflection signals over time with (—) and without (---) self noise cancellation for a) Scenario 1, b) Scenario 2, c) Scenario 3; d-f) Learned self-noise cancellation weights over time, where the signals weight 1 and weight 2 correspond to the weighting applied to the parallel fibre signals p_1 and p_2 respectively, for d) Scenario 1, e) Scenario 2, f) Scenario 3

relative increase in speed of learning as the parallel fibre signals increase in magnitude.

3.4.3 Effect on trajectory control

The trajectory control microzone is upstream, so unaffected by self noise cancellation, or map calibration.

4 Discussion

Multiple cerebellar microzones, 42 in total, based on the adaptive filter model of cerebellar function, have been applied simultaneously to improve sensor trajectory control, sensory noise cancellation and map calibration in real-time in a whiskered robot platform. The results show that it is possible to implement a large number of parallel adaptive cerebellar microzones in a real-time robotics task and in each case they improve the performance in the task. The average improvements, measured experimentally as the percentage reduction in error, for trajectory control, sensory noise cancellation and map calibration were 87%, 32% and 82% respectively.

Each adaptive controller used for trajectory control substantially improved on the fixed, under-gained brainstem controller. In cases where there were reafferent, or self generated, components present in the whisker deflection signals the adaptive controllers for sensory noise cancellation removed these contributions and reduced the variance of the whisker deflection signals during active whisking. The adaptive filters for map calibration also performed well. The results show that even with a relatively coarse encoding of the sensory space, here represented using just 64 parallel fibres, the algorithm is capable of significantly reducing the error introduced from a miscalibrated distorted initial map.

4.1 Internal Algorithm

The internal algorithm was based on the adaptive filter model of cerebellar function, the advantages of this computational approach in comparison to detailed spiking models is considered at length in [46]. Each task was implemented using the same adaptive filter algorithm but with different external connections. The only difference in this internal algorithm between tasks was in the different basis filters used. For map calibration a single filter implementing a unity gain was used, whereas for sensory cancellation and motor plant compensation a bank of two alpha basis filters were used.

Only a single, unity gain filter is required for map calibration as the target was moved randomly on a case by case basis, meaning that the cerebellum does not need to supply any predictive information about the future location of a target. Prior research has shown that a bank of leaky integrator filters enable accurate calibration during a predictive pursuit task [47]. The number of such filters, and their time constants is dependent on the characteristics of the moving target, and the accuracy required.

The properties of the decorrelation learning rule mean that any parallel fibres that signals carry signals uncorrelated with the error signals (i.e. irrelevant signals) will learn associated synaptic weights that are zero. This means that the bank of basis filters could have included all relevant, as well as any irrelevant basis filters or input signals.

The fact that in the biological system there are many (up to 180,000 in the human) parallel-fibre inputs to a given Purkinje cell [48] implies that the majority of the parallel fibre signals carry information that is weakly, or unrelated to the task. This is supported by experimental evidence that has shown that the vast majority of synapses between parallel fibres and Purkinje cells are silent [49, 50, 43].

All parameter values were determined by the experimental conditions, as detailed in the methods section. The learning rates were set such that fast stable learning could be achieved over the duration of the experiments and in all cases weights were learned from initially zero. The learning rates themselves were therefore determined by the experimental conditions, and not driven by biological considerations. For map calibration, a learning rate 10 times smaller than that for trajectory control and self noise

cancellation was used, this was to ensure the occasional large errors caused by inaccurately detecting the target in experiments did not have a significant effect on learning.

To speed learning here, the basis filter outputs were transformed by a fixed matrix, Q to approximately orthonormalise parallel fibre signals. This matrix was estimated from batch Singular Value Decomposition (SVD) of the expected motor command. This step was required to ensure that learning was fast enough to provide satisfactory performance over the time-span of the robotic experiments, as well as to compensate expediently for any actuator changes over time. For robotic applications, where learning needs to occur over short time frames, such a step to speed learning is essential. It is uncertain if and how optimisation for fast learning might occur in biology, this is discussed further in Section 4.4.

4.2 Microzone connectivity

For many regions of the cerebellum details of the external connections and the corresponding functions are poorly understood. The external connectivity of the microzones used here was therefore based in part on biological evidence and in part on the control requirements of the different robotic tasks.

The scheme for controlling the whisking angle is based on a simplified model of the VOR [51]. The use of a reference model (\mathbf{M}) is an extension of this model, which enables the control of strictly proper plants [36], but it is at present unclear how this problem is solved in biological systems. The scheme for noise cancellation is based on adaptive noise cancellation [30, 52, 53], a neural equivalent of which has been proposed that makes predictions about specific neural connections [52]. Further experiments are needed to check these predictions. The map calibration architecture used here is proposed as a way of calibrating a map within the adaptive filter framework. The model also makes firm predictions concerning external connectivity and the signals carried by component pathways. These predictions are discussed further in [47], however, further investigation is again required to check them.

The coupling between the architectures for the different tasks proposed here are not directly motivated by biology, but they raise important questions for biology of how adaptation of one cerebellar microcircuit can influence and effect performance of other circuits. In the experimental results, coupling between the noise cancellation and trajectory control schemes is evident in the noise cancellation results: in some cases the variance increased before decreasing due to the increase in input signal driven by the trajectory control circuit. Gating was needed to avoid learning in certain circumstances (e.g. a contact will modify the whisker angle trajectory and so result in tracking errors, but as these errors are caused by an external influence they should not affect or adapt the controller). A similar gating process has been observed in some motor behaviours [54].

4.3 Robot control

A synthetic, uniform cerebellar chip algorithm that could be plugged into existing systems to fine tune and improve performance in a range of tasks has great potential for robotic applications. Particularly in the context of complex, nonlinear, lightweight, multi-degree of freedom robots, where considerable control efforts are required. A cerebellar chip algorithm could help to reduce the initial control design effort as initial control structures would need only to provide an approximate solution which could be tuned using the cerebellar chip algorithm.

To successfully apply many cerebellar chip algorithms concurrently for the control of different tasks within multi-degree of freedom robots the computational requirements of the algorithm must be reduced to a minimum. Here, where memory was required for the task (in trajectory tracking and noise cancellation tasks) we used a bank of two alpha basis filters. These basis provide a compact way of representing a range of delays and reduce the number of computations required in comparison to using tapped delay lines. The computational efficiency of our cerebellar chip algorithm enabled us to apply forty-two separate chips simultaneously during a real time robotics task. This suggests that the algorithm presented here has promise for the control of complex robots.

4.4 Limitations and future directions

The biological plausibility of our scheme is in part limited by a lack of detailed knowledge of how the biological system operates in full. Overall, the biological evidence appears to be consistent in broad terms with the multizones algorithm proposed here. The next step is to consider more detailed evidence, that could be provided by testing specific findings generated by the presented results. Key areas of future work, where the biological solution remains unclear (and so the implications on our results) include: fast learning, functional difference between microzones, and coupling between learning in different tasks.

To speed learning, a matrix multiplication is used to approximately orthonormalise parallel fibre signals. It is uncertain if and how such an orthonormalisation procedure might occur in biology. If it did, it would likely occur in granular layer processing [55, 36]. It has previously been suggested that granular plasticity allows optimal basis to be learned overtime [55, 56], with the granular layer using an architecture that bears resemblance to machine learning algorithms that perform signal orthogonalisation [56, 57]. If and how an orthonormalisation procedure might occur in biology is an intriguing question.

There is some evidence to suggest functional differences between microzones might occur due to physiological differences within the microcircuit, as well as the input-output connections [58]. It is possible that differences in cerebellar function arise not only because of differences in input-output connections, but also because the cortical microcircuit itself varies from region to region of the cerebellum (e.g. see [59]). However, it appears at present that these regional variations are consistent with predictions from the adaptive filter model of cerebellar function [59] and it is possible that they serve to change the parameters of the basic algorithm (such as learning rates, or types of basis filters used) rather than the algorithm itself. However, the issue of actual functional differences between microzones is very unclear at present. As research into functional differences progresses, an avenue of future work would be to return to our results to assess how they are impacted by such functional differences.

Our results show that coupling between microcircuits means that learning for one task can impact performance in other tasks. The coupling between the architectures for different tasks proposed in the paper is not directly motivated by biology. Another area of future work, is to investigate methods of coupling between trajectory control and self-noise cancellation within the biological system from the neuroscientific literature.

5 Conclusions

The experimental results demonstrate it is possible to design and implement a large number of adaptive cerebellar microzones simultaneously, in real-time, and for each case they improve performance. This is the first embodied demonstration of the cerebellar chip metaphor. Our analysis demonstrates that learning in microzones downstream of a given zone impacts performance, therefore learning is coupled across microzones. Empirically, learning was demonstrated to be stable in each task. This was likely facilitated by using low learning rates, mitigating the adverse affects of coupling. Further investigation is required to theoretically analyse the stability of the learning system and also to investigate if this problem occurs in the biological system, as well as biologically plausible mechanisms for dealing with coupling.

Acknowledgment

The authors would like to thank the EPSRC for financial support provided by grant EP/I032533/1.

References

- [1] F. A. Azevedo, L. R. Carvalho, L. T. Grinberg, J. M. Farfel, R. E. Ferretti, R. E. Leite, R. Lent, S. Herculano-Houzel, *et al.*, “Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain,” *Journal of Comparative Neurology*, vol. 513, no. 5, pp. 532–541, 2009.
- [2] M. Ito, *The Cerebellum and Neural Control*. New York, Raven, 1984.
- [3] D. Marr, “A theory of cerebellar cortex,” *J. Physiol.*, vol. 202, no. 2, pp. 437–470, 1969.
- [4] J. S. Albus, “A theory of cerebellar function,” *Math. Biosci.*, vol. 10, no. 1, pp. 25–61, 1971.
- [5] J.-H. Gao, L. M. Parsons, J. M. Bower, J. Xiong, J. Li, and P. T. Fox, “Cerebellum implicated in sensory acquisition and discrimination rather than motor control,” *Science*, vol. 272, no. 5261, pp. 545–547, 1996.
- [6] S.-J. Blakemore, C. D. Frith, and D. M. Wolpert, “The cerebellum is involved in predicting the sensory consequences of action,” *Neuroreport*, vol. 12, no. 9, pp. 1879–1884, 2001.
- [7] B. E. Murdoch, “The cerebellum and language: historical perspective and review,” *Cortex*, vol. 46, no. 7, pp. 858–868, 2010.
- [8] O. Baumann, R. J. Borra, J. M. Bower, K. E. Cullen, C. Habas, R. B. Ivry, M. Leggio, J. B. Mattingley, M. Molinari, E. A. Moulton, *et al.*, “Consensus paper: the role of the cerebellum in perceptual processes,” *The Cerebellum*, vol. 14, no. 2, pp. 197–220, 2015.
- [9] J. Porrill, P. Dean, and S. R. Anderson, “Adaptive filters and internal models: multilevel description of cerebellar function,” *Neural Networks*, vol. 47, pp. 134–149, 2013.
- [10] R. Apps and M. Garwicz, “Anatomical and physiological foundations of cerebellar information processing,” *Nature Reviews Neuroscience*, vol. 6, no. 4, pp. 297–311, 2005.
- [11] M. Manto, “Mechanisms of human cerebellar dysmetria: experimental evidence and current conceptual bases,” *Journal of neuroengineering and rehabilitation*, vol. 6, no. 1, p. 10, 2009.
- [12] T. Honda, S. Nagao, Y. Hashimoto, K. Ishikawa, T. Yokota, H. Mizusawa, and M. Ito, “Tandem internal models execute motor learning in the cerebellum,” *Proceedings of the National Academy of Sciences*, vol. 115, no. 28, pp. 7428–7433, 2018.
- [13] J. Peters and P. van der Smagt, “Searching a scalable approach to cerebellar based control,” *Applied intelligence*, vol. 17, no. 1, pp. 11–33, 2002.
- [14] P. Van Der Smagt, “Cerebellar control of robot arms,” *Connection Science*, vol. 10, no. 3-4, pp. 301–320, 1998.
- [15] D. M. Wolpert and M. Kawato, “Multiple paired forward and inverse models for motor control,” *Neural networks*, vol. 11, no. 7, pp. 1317–1329, 1998.
- [16] W. T. Miller III, R. P. Hewes, F. H. Glanz, and L. G. Kraft III, “Real-time dynamic control of an industrial manipulator using a neural network-based learning controller,” *Robotics and Automation, IEEE Transactions on*, vol. 6, no. 1, pp. 1–9, 1990.
- [17] G. A. Bekey, *Autonomous robots: from biological inspiration to implementation and control*. MIT press, 2005.
- [18] R. R. Carrillo, E. Ros, C. Boucheny, and J.-M. C. Olivier, “A real-time spiking cerebellum model for learning robot control,” *Biosystems*, vol. 94, no. 1, pp. 18–27, 2008.

- [19] C. Casellato, A. Antonietti, J. A. Garrido, R. R. Carrillo, N. R. Luque, E. Ros, A. Pedrocchi, and E. D'Angelo, "Adaptive robotic control driven by a versatile spiking cerebellar network," 2014.
- [20] M. C. Capolei, N. A. Andersen, H. H. Lund, E. Falotico, and S. Tolu, "A cerebellar internal models control architecture for online sensorimotor adaptation of a humanoid robot acting in a dynamic environment," *IEEE Robotics and Automation Letters*, vol. 5, no. 1, pp. 80–87, 2019.
- [21] S. Tolu, M. Vanegas, J. A. Garrido, N. R. Luque, and E. Ros, "Adaptive and predictive control of a simulated robot arm," *International journal of neural systems*, vol. 23, no. 03, p. 1350010, 2013.
- [22] A. Antonietti, D. Martina, C. Casellato, E. D'Angelo, and A. Pedrocchi, "Control of a humanoid nao robot by an adaptive bioinspired cerebellar module in 3d motion tasks," *Computational intelligence and neuroscience*, vol. 2019, 2019.
- [23] E. Massi, L. Vannucci, U. Albanese, M. C. Capolei, A. Vandesompele, G. Urbain, A. M. Sabatini, J. Dambre, C. Laschi, S. Tolu, *et al.*, "Combining evolutionary and adaptive control strategies for quadruped robotic locomotion," *Frontiers in Neurorobotics*, vol. 13, p. 71, 2019.
- [24] P. F. Verschure and T. Voegtlin, "A bottom up approach towards the acquisition and expression of sequential representations applied to a behaving real-world device: Distributed adaptive control iii," *Neural Networks*, vol. 11, no. 7, pp. 1531–1549, 1998.
- [25] C. Hofstoetter, M. Mintz, and P. F. Verschure, "The cerebellum in action: a simulation and robotics study," *European Journal of Neuroscience*, vol. 16, no. 7, pp. 1361–1376, 2002.
- [26] I. Herreros, G. Maffei, S. Brandi, M. Sanchez-Fibla, and P. F. Verschure, "Speed generalization capabilities of a cerebellar model on a rapid navigation task," in *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*, pp. 363–368, IEEE, 2013.
- [27] J. L. McKinstry, G. M. Edelman, and J. L. Krichmar, "A cerebellar model for predictive motor control tested in a brain-based device," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 103, no. 9, pp. 3387–3392, 2006.
- [28] A. Lenz, S. R. Anderson, A. G. Pipe, C. Melhuish, P. Dean, and J. Porrill, "Cerebellar-inspired adaptive control of a robot eye actuated by pneumatic artificial muscles," *IEEE T. Syst. Man. Cy. B*, vol. 39, no. 6, pp. 1420–1433, 2009.
- [29] L. Vannucci, E. Falotico, S. Tolu, V. Cacucciolo, P. Dario, H. H. Lund, and C. Laschi, "A comprehensive gaze stabilization controller based on cerebellar internal models," *Bioinspiration & biomimetics*, vol. 12, no. 6, p. 065001, 2017.
- [30] S. R. Anderson, M. J. Pearson, A. Pipe, T. Prescott, P. Dean, and J. Porrill, "Adaptive cancelation of self-generated sensory signals in a whisking robot," *Robotics, IEEE Transactions on*, vol. 26, no. 6, pp. 1065–1076, 2010.
- [31] S. Tolu, M. C. Capolei, L. Vannucci, C. Laschi, E. Falotico, and M. V. Hernandez, "A cerebellum-inspired learning approach for adaptive and anticipatory control," *International Journal of Neural Systems*, vol. 30, no. 01, p. 1950028, 2020.
- [32] M. Fujita, "Adaptive filter model of the cerebellum," *Biol. Cybern.*, vol. 206, pp. 195–206, 1982.
- [33] P. Dean, J. Porrill, C.-F. Ekerot, and H. Jörntell, "The cerebellar microcircuit as an adaptive filter: experimental and computational evidence," *Nature Reviews Neuroscience*, vol. 11, no. 1, pp. 30–43, 2010.

- [34] P. Dean, S. Anderson, J. Porrill, and H. Jörintell, “An adaptive filter model of cerebellar zone C3 as a basis for safe limb control?,” *The Journal of Physiology*, vol. 591, no. 22, pp. 5459–5474, 2013.
- [35] E. D. Wilson, T. Assaf, M. J. Pearson, J. M. Rossiter, S. R. Anderson, J. Porrill, and P. Dean, “Cerebellar-inspired algorithm for adaptive control of nonlinear dielectric elastomer-based artificial muscle,” *Journal of The Royal Society Interface*, vol. 13, no. 122, p. 20160547, 2016.
- [36] E. D. Wilson, T. Assaf, M. J. Pearson, J. M. Rossiter, P. Dean, S. R. Anderson, and J. Porrill, “Biohybrid control of general linear systems using the adaptive filter model of cerebellum,” *Frontiers in neurorobotics*, vol. 9, 2015.
- [37] T. Assaf, E. D. Wilson, S. Anderson, P. Dean, J. Porrill, and M. J. Pearson, “Visual-tactile sensory map calibration of a biomimetic whiskered robot,” in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 967–972, 2016.
- [38] R. A. Grant, A. L. Sperber, and T. J. Prescott, “The role of orienting in vibrissal touch sensing,” *Frontiers in behavioral neuroscience*, vol. 6, p. 39, 2012.
- [39] L. W. Bosman, A. R. Houweling, C. B. Owens, N. Tanke, O. T. Shevchouk, N. Rahmati, W. H. Teunissen, C. Ju, W. Gong, S. K. Koekkoek, *et al.*, “Anatomical pathways involved in generating and sensing rhythmic whisker movements,” *Frontiers in integrative neuroscience*, vol. 5, p. 53, 2011.
- [40] M.J.Pearson, T.Assaf, and J.Rossiter, “High speed switched, multi-channel drive for high voltage dielectric actuation of a biomimetic sensory array,” in *3rd International conference on Biomimetic and Biohybrid Systems*, pp. 414–416, Springer International Publishing, 2014.
- [41] B. Mitchinson, T.-S. Chan, J. Chambers, M. Pearson, M. Humphries, C. Fox, K. Gurney, and T. J. Prescott, “Brahms: Novel middleware for integrated systems computation,” *Advanced Engineering Informatics*, vol. 24, no. 1, pp. 49–61, 2010.
- [42] K. Doya, “What are the computations of the cerebellum, the basal ganglia and the cerebral cortex?,” *Neural networks*, vol. 12, no. 7, pp. 961–974, 1999.
- [43] J. Porrill and P. Dean, “Silent synapses, ltp, and the indirect parallel-fibre pathway: computational consequences of optimal cerebellar noise-processing,” *PLoS Comput Biol*, vol. 4, no. 5, p. e1000085, 2008.
- [44] B. Widrow and S. D. Stearns, *Adaptive signal processing*. Englewood Cliffs, NJ: Prentice-Hall, 1985.
- [45] J. Porrill, P. Dean, and J. V. Stone, “Recurrent cerebellar architecture solves the motor-error problem,” *Proceedings of the Royal Society of London-B*, vol. 271, no. 1541, pp. 789–796, 2004.
- [46] P. Dean and J. Porrill, *The importance of Marr’s three levels of analysis for understanding cerebellar function*. Oxford University Press, 2016.
- [47] E. D. Wilson, S. R. Anderson, P. Dean, and J. Porrill, “Sensorimotor maps can be dynamically calibrated using an adaptive-filter model of the cerebellum,” *PLoS computational biology*, vol. 15, no. 7, 2019.
- [48] M. Ito, “Cerebellar circuitry as a neuronal machine,” *Progress in neurobiology*, vol. 78, no. 3, pp. 272–303, 2006.
- [49] C.-F. Ekerot and H. Jömtell, “Parallel fiber receptive fields: a key to understanding cerebellar operation and learning,” *The Cerebellum*, vol. 2, no. 2, pp. 101–109, 2003.

- [50] P. Isope and B. Barbour, “Properties of unitary granule cell- purkinje cell synapses in adult rat cerebellar slices,” *The Journal of neuroscience*, vol. 22, no. 22, pp. 9668–9678, 2002.
- [51] P. Dean, J. Porrill, and J. V. Stone, “Decorrelation control by the cerebellum achieves oculomotor plant compensation in simulated vestibulo-ocular reflex,” *Proceedings of the Royal Society of London B: Biological Sciences*, vol. 269, no. 1503, pp. 1895–1904, 2002.
- [52] S. R. Anderson, J. Porrill, M. J. Pearson, A. G. Pipe, T. J. Prescott, and P. Dean, “An internal model architecture for novelty detection: implications for cerebellar and collicular roles in sensory processing,” *PloS one*, vol. 7, no. 9, pp. e44560–e44560, 2012.
- [53] B. Widrow, J. R. Glover Jr, J. M. McCool, J. Kaunitz, C. S. Williams, R. H. Hearn, J. R. Zeidler, E. Dong Jr, and R. C. Goodlin, “Adaptive noise cancelling: Principles and applications,” *Proceedings of the IEEE*, vol. 63, no. 12, pp. 1692–1716, 1975.
- [54] R. Apps, “Movement-related gating of climbing fibre input to cerebellar cortical zones,” *Progress in neurobiology*, vol. 57, no. 5, pp. 537–562, 1999.
- [55] J.-M. C. Olivier, M. P. Arnold, T. J. Sejnowski, and M. A. Jabri, “Parallel fiber coding in the cerebellum for life-long learning,” *Autonomous robots*, vol. 11, no. 3, pp. 291–297, 2001.
- [56] E. Oja, “Neural networks, principal components, and subspaces,” *International journal of neural systems*, vol. 1, no. 01, pp. 61–68, 1989.
- [57] R. M. Kelly and P. L. Strick, “Cerebellar loops with motor cortex and prefrontal cortex of a nonhuman primate,” *Journal of neuroscience*, vol. 23, no. 23, pp. 8432–8444, 2003.
- [58] H. Zhou, Z. Lin, K. Voges, C. Ju, Z. Gao, L. W. Bosman, T. J. Ruigrok, F. E. Hoebeek, C. I. De Zeeuw, and M. Schonewille, “Cerebellar modules operate at different frequencies,” *Elife*, vol. 3, p. e02536, 2014.
- [59] N. L. Cerminara, E. J. Lang, R. V. Sillitoe, and R. Apps, “Redefining the cerebellar cortex as an assembly of non-uniform purkinje cell microcircuits,” *Nature Reviews Neuroscience*, vol. 16, no. 2, pp. 79–93, 2015.