Cerebellar Control of a Line Following Robot

David Collins and Gordon Wyeth
Computer Science and Electrical Engineering Department
University of Queensland
St Lucia, Queensland 4072
Australia

Abstract

This paper proposes a robot control architecture inspired by the biology of the cerebellum. The cerebellum is thought to be the location of movement coordination in the body. It is hypothesized that through the use of some form of learned internal model, the cerebellum is able to overcome inherent sensory latency and coordinate fast, accurate movement without the need of a complex mathematical algorithm. The study presented in this paper attempts to encapsulate these attributes in an introductory experiment based around a line following robot. The robot learns to negotiate a circuit using a crude supervising module as the teacher, and eventually becomes more proficient at performing the task than the example on which it based its experience.

1 Introduction

For conventional robotic systems to execute fast, smooth and accurate movements many difficult issues must be addressed. Tuning traditional controllers requires calculations based on accurate plant models or extensive experimentation. Such controllers degrade with environmental or plant variations and do not apply readily to non-linear systems. In many situations the latency of sensory data acquisition only serves to exacerbate the problem. Machine vision systems in particular have slow frame update rates, which if used as the source of feedback control can only produce at best slow, crude movements.

A possible solution could be found by observing the solution that biological systems have employed to combat these effects: a cerebellum. The challenges confronting modern robot designers are inherent on a much larger scale in all vertebrates. It is thought that the cerebellum, which resides at the base of the primate brain, coordinates the complex structural and temporal processes needed in voluntary movement [Kawato, 1995; Eccles *et al.*, 1967]

The biological cerebellum overcomes the problems inherent in many robotic systems. There is no need to mathematically define the plant model, it is learnt through experience, being taught via crude corrective movements. The resulting movement can be fast, accurate and compliant. Learning occurs on a continual basis providing infinite adaptability throughout the life of the application and sensory delays are overcome by predicting outcomes before they occur. Encapsulating these benefits would prove to be a significant advantage to any robotic system.

The cerebellum has inspired a series of successful network models, most notably Albus' CMAC (Cerebellar Model Articulation Controller) [1975a; 1975b]. In this model, inputs to the cerebellum arrive via two major fiber systems. Mossy fibers (MFs) are thought to convey state information, including that of the desired state, from the body and other parts of the brain. Climbing fibers (CFs) are thought to convey error information, again from proprioceptive sensors in the body and from other parts of the brain. Within the cerebellum, Albus [1971] postulated that MF signals went through a $N \rightarrow 100N$ expansive recoding operation onto parallel fibers (PFs) and then went on to synapse, via adjustable weights, on a series of large Purkinje cells (PCs) in a $200{,}000 \rightarrow 1$ convergent mapping. Each PC has a 1:1 connection with a CF which when activated causes the PC synapses that are currently excited by a PF to be modified. PCs are the only outputs from the cerebellum and act to modify efferent muscle commands.

A coordinated movement therefore is instigated via the current and desired states being presented to the cerebellum via MFs, the network forward propagates and depending on the synaptic strengths between PFs and PCs, a pattern of PCs will fire corresponding to a certain movement. If the PC activation pattern produced the wrong movement the CFs corresponding to the PCs responsible for the movement fire, causing the PC – PF synapses to change so that movement will be completed more accurately on subsequent attempts. This cycle is the basis of most artificial cerebellar models, and shares the mechanics with the model presented in this paper.

This paper outlines a cerebellar network structure that learns how to control a line following robot from a conventional controller. The purpose of this investigation is to see whether the generalisation and predictive qualities of cerebellar controller can learn to *out-perform* its non-learning counterpart. The line following task was chosen for:

- its readily obtainable quantitative results with relation to performance,
- simplicity of setup with existing resources, and
- the availability of standard control methods to use in the extra-cerebellar model.

The model specifically addresses some distinct aims:

1. To explore the effectiveness of a cerebellar control architecture within the wheeled mobile robot domain. Most cerebellar models investigated to date have been used in the robot arm environment.
2. To have a cerebellar control system assume autonomous control over the line follower, gaining experience through mistakes as exhibited by a crude corrective extra-cerebellar (EC) module.
3. To have the cerebellar control system out perform the supervising EC module by utilizing local generalization and temporal credit assignment techniques to smooth movements.

2 System Architecture

2.1 Problem Definition

The study was based around a mobile robot as illustrated in figure 1. The robot was primarily studied in simulation, with some results from a real robot. The simulation and real robot were found to produce similar results when controlled from the EC unit. Unfortunately, the cerebellar model could not fit in the limited memory space (32 kbytes) of the robot controller used.

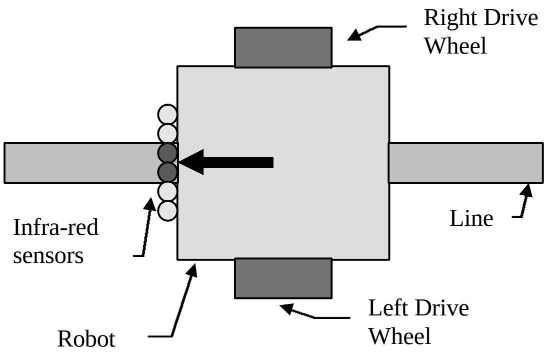

Figure 1 – The Line Following Robot

The robot repeatedly traverses a circuit defined by the guide line, acquiring more cerebellar experience with each lap. The circuit consists of a series of left and right-hand turns with varying degrees of severity. The robot is a differential (wheel chair) drive mobile robot with six infra-red (IR) emitter-detector pair sensors at the front of the robot to detect the robots relative orientation to the line.

The IR sensors are arranged such that two adjacent sensors will be active above the line at all times. To determine which two sensors are in the ON state, the IR detector exhibiting the highest intensity is chosen, followed by an interrogation of the two neighbouring detectors. The strongest of the neighbours is selected as the other active sensor. If either of the end sensors is the most active, then the adjacent inner sensor is automatically chosen to be the other active sensor.

With this configuration the IR sensors can assume five distinct states. Each of these states is assigned a numerical value as the line is positioned from far left to far right across the sensors. States 1 and 5 represent the most extreme cases, with the robot severely offset to either side of the line. These states also indicate the occurrence of a sharp turn in the track. States 2 and 4 represent a mild offset from the central position, or a slight turn on the track. State 3 verifies that the centre of the robot is directly above the line.

2.2 The CISM Model

The structure of the model presented in this paper is inspired, most significantly, by the model developed by Fagg *et al.*[1997a; 1997b; 1997c]. The Fagg model was developed to control a 2-DOF simulated planar arm. Significant differences between the two models were necessitated due to the contrasting applications. The Fagg model operated in two mutually exclusive phases, an execution phase and a correction phase. Each time the arm was instructed to move to the target position, the cerebellar module was given the first opportunity to reach the target successfully. If this attempt was not executed correctly the correction phase began with the EC module making approximate pulsing movements in the general direction of the target. Whilst this was happening CFs would become active, signaling the need for synaptic modification to improve the accuracy of subsequent attempts.

Although the above model proved to be effective in this application, it does not translate readily to some real time mobile robotic applications. The line follower robot does not have the luxury of being able to give the cerebellar module the chance to execute a turn properly. If the cerebellar module is not yet competent, the attempted movement could divert the robot to some point off the line where it would effectively become lost.

Because of the unrecoverable failure associated with an inexperienced attempted turn, the output of the cerebellar module must first be interrogated to judge whether it has acquired the necessary competence to successfully negotiate the given situation. If it is correctly reacting to the sensory information, the outputs of the cerebellar module can be passed to the motors in the form of a descending efferent command. If it is not yet competent, the EC module must intervene and produce a suitable set of motor commands, while the cerebellar module performs some correction to improve its performance in subsequent executions.

Some significant adaptations were required to allow judgements about the suitability of a motor signal to be made before the commands were sent to the motor units. Due to the mutually exclusive nature of the cerebellar and EC modules in the Fagg model, the descending motor command was simply the summation of the two modules. There was no possibility of “collisions” with both modules attempting to control the motors concurrently. In the line follower robot, both the cerebellar and EC modules are producing outputs concurrently. The most appropriate response has to be selected and if the EC response was judged to be the best, this activated a CF response to signal a correction in the cerebellum. A form of gating the outputs of the cerebellar and EC modules was needed which proved to be the most significant external difference to the Fagg model.

The model proposed to address these issues, called CISM (Cerebellar Interrogation and Selection Model), is displayed in figure 2**.** The model consists of two cerebellar modules, one for each wheel to be

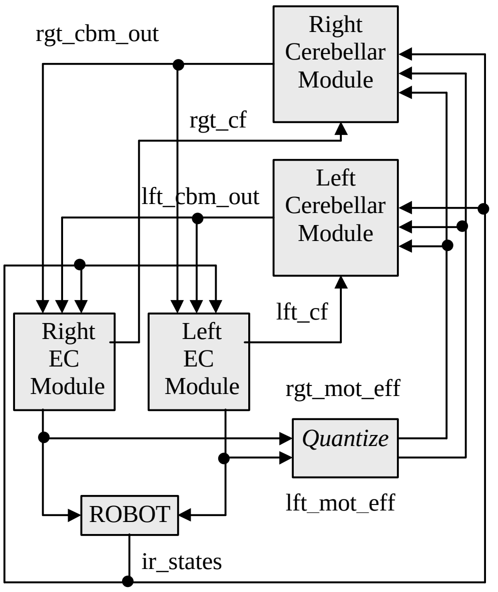

Figure 2 – The CISM model

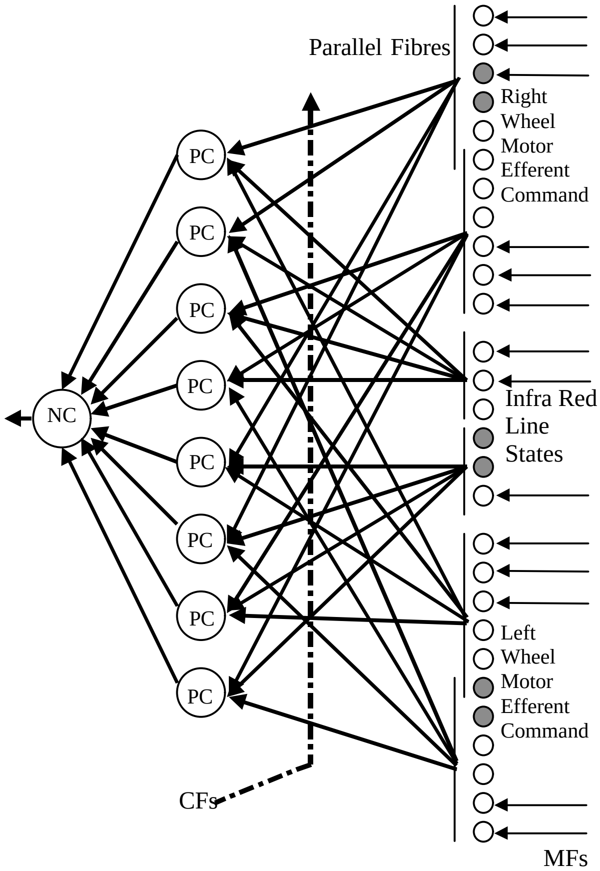

Figure 3 – The Cerebellar Network

controlled and two EC modules. If the cerebellar modules have produced commands deemed satisfactory by the EC modules, they will have their outputs passed through to become the descending motor efferent commands.

2.3 The Cerebellar Module

As stated previously, each cerebellar module receives inputs in the form of mossy fibers (MFs) or climbing fibers (CFs). The CFs in this model are used to signal an error condition, and the MFs are used to convey the IR states (proprioception) and the motor efferent copies. In the Fagg model, MFs also convey target information, but in the case of this line following robot, the goal of having the robot centred over the line is intrinsic to the model, and is accommodated by the EC module. The internal structure of the cerebellar module is shown in figure 3.

The MF inputs, after external quantization, project onto binary state receptive fields which are analogous to granule cells (GCs) in the biological cerebellum. The full range of motor efferent commands are quantized into 10 velocity levels, with a receptive field width of 2. The IR states are already in this format, due to their receptive field arrangement. This type of encoding is the same as the initial stages of the CMAC model (Cerebellar Model Articulation Controller). It allows two neighbouring input vectors to share a GC and thus, share the knowledge associated with that part of the input space during forward propagation through the network. From the GCs, an expansive recoding process occurs as the GCs give rise to several PFs, which in turn synapse upon PCs. During learning these PF-PC synapses are modified to alter the networks performance, all determined by the activation of the CFs. The PCs then converge on a single nuclear cell (NC) via a fixed weight synapse before leaving the cerebellar module.

The cerebellar module has eight PCs each acting as a quasi-feature detector in response to different MF input vectors. The GC to PC projection via the PFs is not a fully connected network. Unless a motor efferent command was zero or the IR states showed the robot to be centred, one of the eight PCs will respond more strongly than the other seven. For example, the top PC in figure 3 will respond maximally to the condition where the line is on the right side of the robot and both the left and right wheels have been commanded to go forward. Therefore the PCs act as state detectors and with a three dimensional input space, with a roughly binary nature (forward/reverse, left/right) to each dimension, 2^3 states can be distinguished, hence the use of eight PCs.

Because of the binary nature of the GC receptive fields, the signal through a PF will either be zero or the weight of the PF-PC synapse. Therefore a PC will have between zero and six of its fifteen PFs active at any time. There are a total of 120 PFs in the cerebellar module, each with a modifiable weight. This combined with the eight PC-NC synapses results in a total of 128 connections per cerebellar module, with no multiplication operations necessary for a forward pass of the network.

2.4 The Extra – Cerebellar Module

The EC module has three main functions. Firstly, it must be sufficiently capable to assume independent control of the line following robot, to ensure that the robot will operate effectively while the cerebellar module is learning. Secondly, the EC module must act to arbitrate between its own output and that of the cerebellar module. It should be qualified to judge the relative competency of the cerebellar module and allow the cerebellum to directly control the motors once it has gained the necessary knowledge for a given situation. Finally, the EC module must be able to provide a function similar to that of the inferior olive (IO) in biology and send meaningful CF activity patterns to the cerebellar module to aid in its learning.

To provide the arbitration and learning functions, some "measure of goodness" must be known for the line following application. A performance measure called the turning index (TI) was introduced to assist in the decision making process. Most forms of motion can be reduced down into two component parts, rotation and translation:

$$s_{trans} = \frac{(v_r + v_l)}{2}$$

$$q_{rot} = \tan^{-1}\left(\frac{(v_r - v_l)}{w}\right)$$

where s_{trans} is the translation displacement, q_{rot} is the rotation angle, v_r and v_l are the right and left wheel velocities respectively and w is the distance between the wheels.

Once the rotation and translation components of the robots motion has been resolved, the turning index is simply:

$$\mathrm{TI} = \frac{q_{rot}}{s_{trans}}$$

The turning index is a measure of how sharp a corner the robot can negotiate. A large turning index indicates a large proportion of rotation with respect to translation, and therefore shows that the robot is currently turning quite sharply. A small turning index indicates a large radius turn and a TI of zero verifies that the robot motion contains no rotation at all.

The EC module has a set response to the five possible IR states. They were acquired via a process of trial and error aided by a human teacher. When the cerebellar modules produce a new command, the turning index resulting from the command is calculated and compared with the turning index known to be successful for the current conditions. If the turning index for the cerebellar generated movement meets or exceeds the turning index known for the EC commands, the cerebellar commands are passed through to become the descending motor efferent commands. If not, the EC module selects the commands it knows to be reliable. A CF is then signaled to indicate that a weight update will be required to make the cerebellar module more competent, should the current conditions arise in future. The motor commands for various IR states and their respective turning indices are displayed in table 1.

Description	IR State	Right Motor Command	Left Motor Command	Turning Index
Large Right Offset	1	-12	20	9.142
Small Right Offset	2	8	20	0.982
Centred	3	15	15	0
Small Left Offset	4	20	8	-0.982
Large Left Offset	5	20	-12	-9.142

Table 1 – Hard Wired EC Motor Command and Turning Indices

3 Learning Algorithm

When the cerebellar module does not generate a motor command deemed adequate, a correction must occur to assist the module in becoming more competent. The credit assignment problem is the problem of assigning responsibility to specific synapses that were considered to be involved in the incorrect production of the motor command. The credit assignment problem can be decomposed into two components, structural credit assignment and temporal credit assignment.

Structural credit assignment is concerned with which of the synapses were responsible for an inadequate response. It produces a spatially coded pattern of PFs, that were active during the production of the erroneous movement.

Temporal credit assignment is used to tag the active PFs through time, so a brief history of a sequence of activation patterns can be recorded. This way the most recent pattern will not be the only one found responsible for an error. The preceding patterns will also be reprimanded, but the extent of their responsibility will diminish as time proceeds. In Faggs model this is used to address the substantial time delays which simulate the physiological delays inherent in any biological system. In the CISM model these delays are not a problem, so it is expected that temporal learning will aid the robot in developing smoother predictive movements.

An eligibility trace (ET) is attached to each PF in the cerebellar module. Each time the output of the network is computed, the PFs that were active in the process are assigned an eligibility of 100%. If they are not active again on the next forward pass their eligibility is downgraded to 60% and so on, until their eligibility to claim responsibility for an error is reduced to zero. When the time comes to make weight adjustments, the eligibility trace for each PF is examined and the PFs with non zero eligibility are subjected to a weight modification. A graphical representation of an eligibility trace is shown in figure 4.

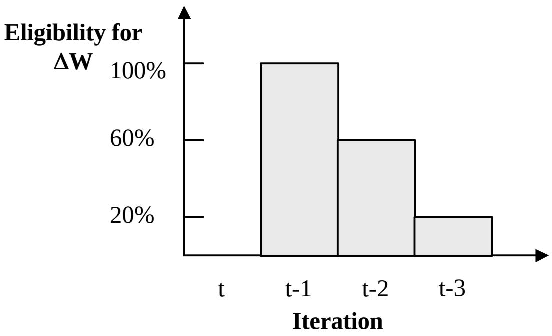

Figure 4 – The Credit Assignment Eligibility Trace

In the CISM model the EC module will determine the suitability of the cerebellar modules outputs. If they are not adequate, the module will send off a CF activation to the cerebellar modules. Each module receives one CF that can assume the values of -1,0 or +1. A zero on the CF means that the previous response of the cerebellar module was sufficient and no correction is required. A positive or negative one means that a correction is needed to make the make to robot turn more or less towards a certain direction in subsequent movements. The direction of the error is conveyed in the sign of the CF and will either cause an increase or decrease in the synaptic efficacy of the eligible PFs.

The weight update rule is shown below for each PF:

$$\Delta W_i = k \cdot CF \cdot ET$$

where K is a constant that is proportional to the rate of convergence, CF is the climbing fiber activation and ET is the eligibility trace value attached to the PF.

4 Experimental Procedure

The experiments conducted with the simulated robot contrasted three different configurations. The first has the robot operating without any cerebellar involvement, being controlled solely by the EC module. The other two configurations allow the cerebellar models to gradually assume control of the robot. One of these arrangements disables the eligibility trace, effectively eliminating temporal credit assignment. Therefore only the PFs active on the last iteration are modified during learning.

5 Simulation Results

The simulated line following mobile robot could navigate the circuit under full cerebellar control within 25 laps, as shown in figure 5. Convergence to this state was not accelerated with the omission of the eligibility trace, as both cerebellar variants approximately learnt at the same rate.

With the results for the mean offset distance recording, the cerebellar configurations both converged to an average offset distance of 0.22 units. The cerebellar model with the eligibility trace enabled faired no better than the model with the trace disabled. These results were still significantly better than those of the EC configuration, which had an average offset of 0.39 units.

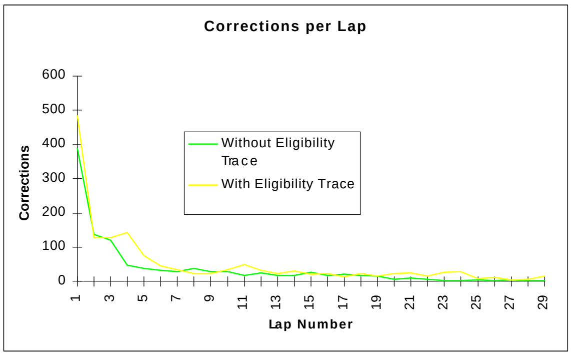

Figure 5 – Number of Corrections needed per Lap

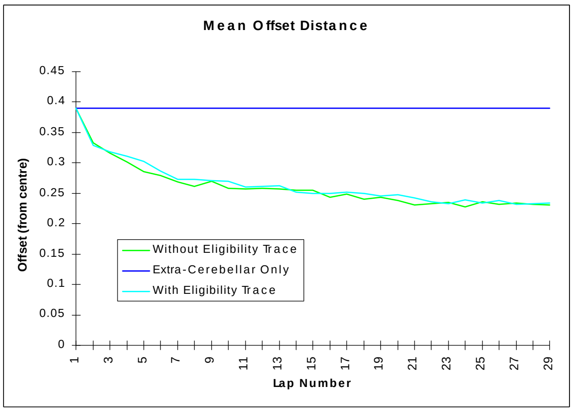

Figure 6 – Mean Offset Distance from the Centre of the Line

The outputs learnt by the cerebellar configurations in response to the IR states are shown in table 2. Note that the "learnt" turning index for the large right offset has not changed due to lack of exposure of the system to large left-hand turns. This implies that large right offset condition was never controlled by the cerebellum.

Description	IR State	Right Motor Command	Left Motor Command	Turning Index
Large Right Offset	1	-12	20	9.142
Small Right Offset	2	5	20	1.374
Centred	3	15	15	0
Small Left Offset	4	20	-4	-3.432
Large Left Offset	5	20	-13	-10.773

Table 2 – Learnt Cerebellar Responses to IR states.

6 Discussion

Comparing the mean offset distance for the cerebellar and EC configurations indicates an apparently excellent learning result. This measure was significantly better for the cerebellar configurations, than the EC configuration alone. The fact that the cerebellar configurations performed the same, regardless of the contribution from an eligibility trace, meant that this was not the result expected. It was thought that the inclusion of an eligibility trace would be the source of any performance improvements, allowing the cerebellar module to predict the required response needed to negotiate a turn smoothly.

The superior results obtained with the cerebellar configurations can be explained by viewing the resulting motor commands of the trained cerebellar modules in table 2. The lack of sharp right hand turns on the circuit caused an inexperienced response to this condition, as the circuit was executed in a counter clockwise manner, dominated by left turns. Apart from this, it can be seen that all the responses have resulted in turning indices, that are over qualified for the turns encountered. During a lap of the circuit, the cerebellar modules should have ceased modifying their weights, once they had become proficient at negotiating a corner for a given set of motor commands. So how did they manage to keep increasing their turning responses for a given set of input vectors?

The answer is evident when the mechanisms involved in the overlapping receptive fields are examined. Consider the case where the cerebellar module becomes competent at negotiating the condition where the IR states equal 4, and both motor commands are +15. Upon future presentation of this particular input pattern, the cerebellar modules will no longer need corrections. Suppose now that the left motor command to a situation is +14, causing the receptive field for the left motor efferent input to shift one position across. Due to the overlapping receptive fields, 5 out of the 6 previously active PFs, will still be active, but their influence may not be enough to produce the correct command. This causes all 6 to be tagged for update, to make them perform better the next time. If the original input pattern was again presented to the network, it would now find that 5 out of the 6 original synapses have since been strengthened causing it to respond more strongly to the input pattern than previously. This is why the cerebellar configurations of the study managed to out perform the EC configuration.

The eligibility trace was not responsible for the increased performance, as it was thought it would. The trace was used by Fagg mainly because it overcame the simulated sensory delays, by storing a copy of the activity constellations, until the afferent signals reporting the effectiveness of an attempt arrived. It was not really relevant in this case.

Although the performance gains came through a different mechanism than that originally hypothesized, they are no less encouraging. There appears to be potential to exploit the local generalization observed in this model in future applications.

7 Conclusions

The cerebellar control architecture presented in this paper was an introductory study into the effectiveness of such a model in the wheeled mobile robot domain. The simple task of line following aided in accomplishing all of the original aims:

- The cerebellar modules, after sufficient training, assumed control of the robot. This was achieved with the aid of a simple teaching module that controlled the robot as the cerebellar modules learnt by example.
- The cerebellar modules did not require any intrinsic knowledge of the task or the robot. It built up all knowledge through observing the actions of the crude extra-cerebellar module.
- The cerebellar modules managed to out perform the teaching module. It achieved this via the local generalization inherent in the cerebellar network, not via temporal credit assignment as hypothesized. This interesting result could be exploited in future applications.

A clear area for immediate work is to test the simulated results on the real robot that it simulates. This will require investigations into improved coding techniques, or expenditure on hardware. The study showed there is much potential for the application of cerebellar control in mobile robotics. The robot soccer domain is the next field of investigation to expose the real potential of these biologically inspired techniques.

References

[Albus,1971] J.S.Albus. A Theory of Cerebellar Function. *Mathematical Biosciences* 10, pp. 25-61, 1971.

[Albus, 1975a] J.S.Albus. A New Approach to Manipulator Control: The Cerebellar Model Articulation Controller (CMAC). *Transactions of the ASME, Journal of Dynamic Systems, Measurement and Control*, pp. 220-227, September 1975

[Albus, 1975b] J.S.Albus. Data Storage in the Cerebellar Model Articulation Controller (CMAC). *Transactions of the ASME, Journal of Dynamic Systems, Measurement and Control*, pp. 228-233, September 1975.

[Eccles *et al.*, 1967] J.C.Eccles, M.Ito and J.Szentagothai. The Cerebellum as a Neuronal Machine. Springer, Berlin, 1967.

[Fagg *et al.*, 1997a] A.H.Fagg, N.Sitkoff, A.G.Barto and J.C.Houk. A Model of Cerebellar Learning for control of Arm Movements Using Muscle Synergies. *Proceedings of the IEEE International Symposium on Computational Intelligence in Robotics and Automation*, pp. 6-12, 1997

[Fagg *et al.*, 1997b] A.H.Fagg, N.Sitkoff, A.G.Barto and J.C.Houk. Cerebellar Learning for Control of a Two-Link Arm in Muscle Space. *Proceedings of the IEEE Conference on Robotics and Automation*, pp. 2638-2644, May 1997.

[Fagg *et al.*, 1997c] A.Fagg, L.Zelevinsky, A.Barto and J.Houk. Using Crude Corrective Movements to Learn Accurate Motor Programs for Reaching. *Extended Abstracts of the NIPS*97 Workshop, Can Artificial Cerebellar Models Compete to Control Robots ?*, Chapter 6, pp. 20-24, 1997.

[Kawato, 1995] M.Kawato. Cerebellum and Motor Control. *The Handbook of Brain Theory and Neural Networks*, MIT Press, Cambridge, Massachusetts, pp. 172-178, 1995.