



University of Plymouth
School of Computing

**Planning Sensory-Motor Associations: A Control Principle
Inspired by the Cerebellum.**

G. Bugmann

Research Report *CNAS-99-01*

June 1999

Centre for Neural and Adaptive Systems

<http://www.tech.plym.ac.uk/soc/research/neural>

Planning Sensory-Motor Associations: A Control Principle Inspired by the Cerebellum.¹

Guido Bugmann

Centre for Neural and Adaptive Systems
School of Computing, University of Plymouth, Plymouth PL4 8AA, UK
<http://www.tech.plym.ac.uk/soc/Staff/GuidBugm/Bugmann.htm>

10 June 1999

Abstract:

The Brain is a slow computer yet humans can skilfully play games where very fast reactions are required. A review of data on motor and cognitive effects of cerebellar lesions leads to propose that the brain minimises reaction time during skilled behaviour by eliminating on-line planning. Behaviour then results from direct mapping from perceptions to actions. The proposed role for the cerebellum is to i) learn sensory-motor associations planned off-line and ii) enable the set of sensory-motor associations most appropriate for a given goal. In that framework, planning defines both the stimulus and the action that will be triggered. In the performance of plans comprising sequences of actions, the cerebellum controls eye fixation to obtain the sensory trigger that initiates the next action in the plan. For producing effective behaviour, such a scheme requires in principle that, for each goal, a large set of sensory motor associations are prepared to cover a wide range of possible initial conditions and variations in execution. Hence the planning process may need to produce not a single optimal solution but a field of solutions. Such a scheme can then solve the paradox of so-called automatic behaviour often seeming to proceed via situation-dependent planning rather than simple stereotypical movements. Applied in the field of robotics, these principles unify purely behaviour-based approaches and approaches based on planning using internal representations. On one hand sensory-motor associations enable fast reactions and on the other hand, being products of planning, these associations enables flexible goal adaptation. Such a principle has been implemented using neural networks in the cases of a simulated arm control problem and the control of a real autonomous wheelchair. An interesting property of the proposed sensory-motor association neural net is its generalisation across initial conditions, which may obviate the need for planning multiple associations.

1. Introduction:

In the game of tennis, the time for the ball to cross the court is almost exactly equal to the reaction time of the player added to its displacement time². There does not seem to be time for planning and it is unclear how players manage to produce strategic shots under these conditions. The cerebellum is known to be involved in the co-ordination of action and its lesion results in slower reactions (Stein and Glickstein, 1992). Can we learn something from studying the cerebellum that may shed some light on the computational strategies used by

¹ This report is based on material initially presented at the IEE workshop on "Self-learning robots III Brainstyle robotics: The cerebellum beyond function approximation", London 15 February 1999.

² The length of the court is approximately 24 m and its width 8m. A ball flying at 100 km/h takes 860ms to travel 24m. Choice reaction time is approximately 370ms and the displacement time, to reach a corner of the court starting from the middle, is estimated at 540ms. It is unclear how much time planning would take.

tennis players? Can we learn something that may help design faster robots despite limitations in computational resources? These are the questions underlying this paper.

How can the cerebellum accelerate reactions? The answer is not straightforward, because the cerebellum does not initiate actions, its inputs arriving usually to the motor cortex around the time when the action starts (Butler, Horne & Hawkins, 1992). Several experiments will be described in section 2 that suggest that the cerebellum helps selecting actions to be executed. In section 3, a possible biological substrate is suggested. In section 4, a conceptual model based on these observations is proposed that combines planning and reactive behaviour. In section 5, a neural network implementation is described to illustrate the principle. It is concluded in section 6 that, the model proposed in this paper may enable skilled behaviour without on-line planning.

2. Role of the cerebellum in planning animal behaviour:

Several experiments show that the cerebellum is involved in planning. A more precise view on its role in that respect emerges from an analysis of the experiments described below.

- In one experiment, human subjects are asked to walk a path of irregularly placed stepping stones as fast as possible (Hollands & Marple-Horvat, 1996). Subjects usually fixate the next stone just before starting to lift the foot that will be placed on the stone. To assess the role of visual sampling, tests were also performed in the dark, with stone positions indicated by LED's turned off at different times. Visual deprivation had no effect when the goal was not illuminated during fixation time and swing phase of the movement. However, deprivation towards the end of the swing affected sometimes the accuracy of the next step. This indicated that subjects were planning the move to the next stepping stone, while fixating the current stone and executing the current stepping action. Interestingly, shift of fixation from the current to the next stone took also place in the dark, indicating that the fixation was more a reflection of a pre-existing plan than a source of information needed for its execution. The intermittent effect of visual deprivation could be due to compensation by spatial memory of the layout, which, in another experiment, has been shown to persist for several seconds in the dark (Thomson, 1980). The stepping-stone experiment done with cerebellar patients showed a reduced accuracy and intermittent delays (prolonged stance) correlated with inaccurate fixating saccades to the next stone (Marple-Horvat et al., 1998). Other experiments described below may allow narrowing down the numerous possible reasons for these effects.

- In another experiment, subjects were told to touch, upon hearing a GO sound, one of the elements of a display that is indicated by a light lit for a few seconds before the GO signal (Day, Thompson, Harding & Marsden, 1998). The reaching movements of healthy subjects had the same characteristics in the dark or in the light, indicating the execution of a pre-planned movement, or as noted above, the existence of an internal spatial memory (world model) that is replacing actual visual input. In contrast, cerebellar patients showed increased variability in initial movement direction and, for the final portion, performed worse in the darkness, indicating their reliance on visual feedback. However, the overall correct initial direction suggests that the problem was not one of spatial memory (not remembering the position of the target), but one of precise planning or plan execution.

- In another experiment, a monkey was trained to touch in the correct order the two elements lit simultaneously in a succession of displays. (Lu, Hikosaka & Miyachi, 1998). In such experiments, a trained animal shows anticipatory saccades to the next element, indicating knowledge of the sequence, and as suggested above, the existence of a plan as for which element to touch next. Lesions in the dorsolateral dentate nucleus, the output of the cerebellum to the premotor and supplementary motor area, reduced the number of anticipatory saccades and caused more errors. This could indicate a loss of sequence

memory, or errors in translating sequence knowledge into a correct plan or action. The next experiment shows that the latter is more likely.

- A patient with cerebellar lesions showed spatial dysgraphia characterised by a large number of stroke omissions or stroke repetitions when writing in upper case, even when copying a word (Silveri, Misciagna, Leggio & Molinari, 1997). Errors occurred despite a memory of the sequence of strokes not being required. The authors suggested that this was a case of neglected visual feedback (e.g. telling that all strokes of a letter have been produced), but as the performance worsened in darkness, visual feedback was being used. Somehow, cerebellar lesions prevented available knowledge (visible example or memory of sequence) to be correctly converted into a motor plan.

- Patients with cerebellar lesions tended to do more illegal moves in a Tower of Hanoi test which is a test of planning capability (Grafman et al., 1992). This effect was neither due to not understanding the instructions nor to some error in motor control. It was more like an inadequate plan was being selected and executed from start to end.

In summary, healthy animals plan ahead, which is revealed by anticipatory saccades that do not seem necessary for controlling the current action. The symptoms of cerebellar patients can be described as resulting from an inappropriate selection of subgoals.

3. Possible biological substrate:

The posterior parietal cortex (PPC) contains maps of possible movements. When given hand, arm or eye movements are intended, specific subgroups of neurones become active (Snider, Batista & Andersen, 1997). However, for actual movement to result, this activity must reach the primary motor cortex (M1). This can be done via the projections of PPC to the pre-motor (PM) and M1. It is speculated here that, for a given goal, movements could be pre-selected via the cerebellar input to superficial layers in PPC (Okuda, 1997). Such input would not be sufficient to activate PPC cells but could lower the firing threshold of selected PPC cells for responding to sensory cues. Thereby, movements would be generated not only with reduced latency but also in a purely feedforward circuit. The cerebellum could pre-select successive movements in a sequence by monitoring the execution of motor commands via various collaterals of descending fibres and ascending somatosensory feedback. The details of possible sequence learning circuits have been discussed elsewhere (see e.g. Braitenberg, 1997).

Plans are supposed to be elaborated in the pre-frontal cortex (although it is not clear how). Such plans can control movements in a voluntary way, via pre-motor areas, but seemingly at a slower rate and by addressing a small number of joints at a time (Day et al., 1998). In the proposed scheme, plans would be elaborated off-line, then downloaded to the cerebellum in the form of a series of sensory triggers and their corresponding actions, essentially equivalent to Stimulus-Reaction sequences. Such sequences would be replayed with timing determined by the rate at which new sensory stimuli appear. The details of the mechanism used to set up sensory-motor associations are unknown. At this stage, the proposed model is pitched at a functional level and further work is needed to determine if it can be supported by anatomical data. Many areas of the cerebellum have well documented role motor control but its lateral zone is reciprocally connected with various frontal areas (Middleton and Strick, 1998) and could be involve in the proposed function.

4. Conceptual model of the role of the cerebellum.

The reviewed data on cerebellar lesions and eye fixation are consistent with following hypotheses:

1. Visual fixation provides the trigger for pre-planned actions.
2. Planning is used to set-up set of perception-action associations.
3. The cerebellum is involved in priming neurones involved in the sensory and/or motor elements of the associations.

Figure 1 illustrates these hypotheses in the case of the stepping stone experiment of (Hollands & Marple-Horvat, 1996). During fixation of the target of the current swing, subjects would engage in planning the next swing. This entails (i) predicting the view at the end of the swing, which then is set as a perceptual key for the next swing and (ii) calculating the parameters of the next swing and the saccade to the next target.

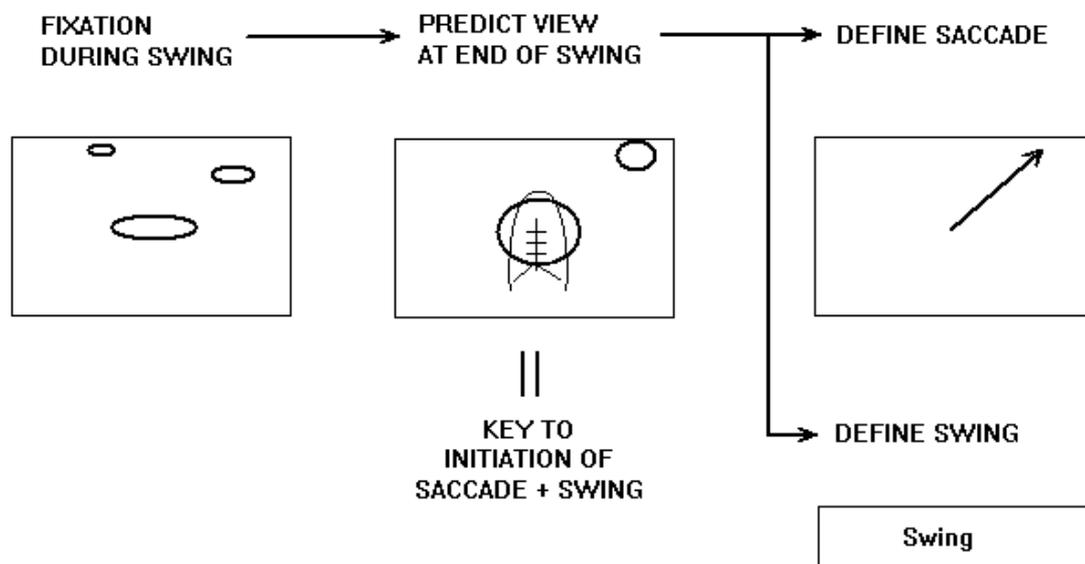


Figure 1. Model of the computational processes underlying the stepping stone experiment (Hollands & Marple-Horvat, 1996).

The properties of this model are: i) Priming causes faster response when the expected sensory pattern is seen, ii) The control loop is faster due to the elimination of the planning time. iii) If fixation is incorrect, the perceptual pattern does not precisely resemble the one defined in the plan and neuronal noise plays a greater part in action selection, which causes errors and slow initiation due to loss of the speed benefit of priming.

The hypotheses above are not specific enough on the biological substrate to constitute a model of cerebellar function, e.g. it is not clear which neurones are primed or how sensory motor association are set up by planning processes, however, this is irrelevant if the aim is the design of a biologically inspired robot control system.

Nevertheless, the model produces some testable predictions. For instance, swing and saccades should be initiated simultaneously. Now, regarding effects of cerebellar lesions, a correlation between swing and saccade aiming errors is probably not to be expected; that is if errors are due to noise in the action selection process.

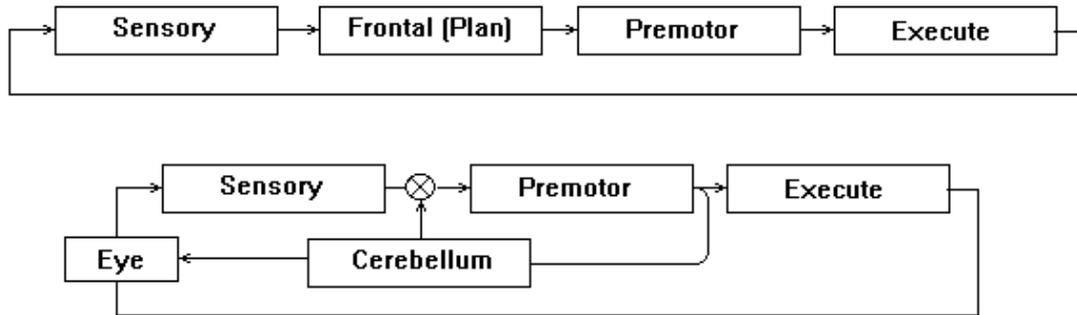


Figure 2. *Top diagram:* Serial view on action selection: Sensory information is used for planning and the resulting plan is executed. *Bottom diagram:* Proposed model: The cerebellum contributes to priming a pre-selected set of potential sensory-motor associations, and controls the eye fixation to ensure that appropriate sensory patterns are acquired.

From a control perspective, the key notion in the hypothesis is that actions are *selected*. This suggests that time can be saved if, instead of planning on the basis of current sensory information, pre-calculated sets of sensory-motor responses (a “plan”) were available and appropriate responses were selected on the basis of current sensory information. An important characteristic of a controller implementing such a principle is that a plan must be reflected in a variety of response being set up, to account for the variability of one’s action execution and the unpredictability of the environment. Choosing a plan therefore corresponds to the selection of a large set of possible sensory-motor associations. It might be possible to extend the model to the use of sensory patterns for the selection of plans themselves. Therefore, there is an interesting blurring of the distinction between planning sensory-motor associations and priming them. Hence plans or actions could be selected by the same principle.

A simple example of the combined use of sensory - plan association and sensory - action associations is the return of a serve in tennis. A tennis player waiting for a serve would have prepared several responses, e.g. forehand or backhand, cross court or along the line, depending on the class of the service trajectory. Then, from initial snapshots of the ball trajectory, an appropriate response could be rapidly selected (plan triggered by sensory patterns), then rapidly executed (actions triggered by sensory patterns).

A key feature of this approach is that off-line planning needs to define also a sensory trigger, not only an action as in the serial control approach in figure 2. This is not necessarily simple and practical implementation will enable to assess the usefulness of such an approach. Some implementations with simple sensory input have already been tested. The next section describes an artificial neural network architecture that demonstrates the proposed principle. A real-world application, the control of the trajectories of an autonomous wheelchair is described in (Bugmann et al., 1998).

5. Neural network implementation.

The neural network architecture described below was initially proposed in (Alhoefer and Bugmann, 1995). The principle of the implementation is: i) Use the output of a planner to train a sensory-motor (SM) association network; ii) Let the SM network control the sequence of actions.

As a planner, a neuro-resistive grid was used (Bugmann et al., 1995) that approximates Laplacian planning. In principle, any other techniques could be used, directed graphs, etc. For encoding sensory-motor associations, a network of Normalized Radial Basis Functions is used (Bugmann, 1998).

Section 5.1 details the implementation of a neuro-resistive grid and describes the two-dimensional robot arm control problem that is used to illustrate the proposed principle of

planning SM associations. Section 5.2 details the operation of the NRBF network and its application to the robot arm problem. Section 5.3 describes the generalisation properties of such a network.

5.1 Planning with a neural resistive grid.

For using a resistive grid for planning, the state space is divided into a set of small N -dimensional cubes. Each cube corresponds to a node in the resistive grid. Each node is connected to its $2N$ neighbours. Planning can be performed by setting the potential of the node corresponding to the goal state to a high value, e.g. 1. The nodes corresponding to obstacles, or forbidden states, are set to a potential zero. Electric current flows from the goal, through the grid and towards the obstacles. At any point in the grid, the direction of the current flow indicates the shortest path to the goal.

In a neural resistive grid (Bugmann et al.,1995), one neurone is assigned to each node. The neurone sets its output y_j to the average value of the outputs of its neighbours y_i :

$$y_j = f\left(\sum_i W_{ij} y_i\right) + I_j \quad (1)$$

where

$$f(x) = 0 \text{ for } x < 0; \quad f(x) = x \text{ for } x \in [0,1]; \quad f(x) = 1 \text{ for } x > 1.$$

After all neurones in the grid have been updated a large number of times, their outputs approximate the potential distribution in a physical resistive grid. It is shown in [2] that, in the limit of a very large number of nodes in the grid, the updating rule (1) is equivalent to the Poisson equation:

$$\nabla^2 y = -bI. \quad (2)$$

Where $b=2N/\Delta r^2$ and Δr is the horizontal and vertical distance between nodes. The use of a linear saturating transfer function $f(x)$ in (1) allows external control to be exerted on the resistive grid by the input I_j from the map[1] layer in Figure 3. Each neurone in that layer has only one output that is connected to the neurone in the resistive grid with corresponding location. By setting the output of the goal node in the map[1] to 1 (white node in the Figure), it forces the goal node in the neural resistive grid to have an output 1. Similarly, setting the activity of obstacles nodes in map[1] to -1 (grey nodes in the Figure) will force their counterparts in the neural resistive grid to have outputs zero. This flexible architecture allows dealing with any configuration of obstacles and goal(s).

The 2-dimensional resistive grid in Figure 3 represents the configuration space of a 2-joint robot-arm that is shown in Figure 4. A point in the configuration space represents a configuration of the robot arm. A path in the configuration space represents a sequence of arm configurations. The arm in this example has only two mobile joints, so its configuration space is two-dimensional. In Figure 3, the x-axis is the shoulder joint angle and the y-axis is the elbow joint angle. Those joint angles that correspond to collisions with walls and obstacles are represented by forbidden positions in the grid. The goal node corresponds to a configuration where the grip is in contact with the filled circle in Figure 4. This neural implementation of the resistive grid is very flexible due to the possibility of incorporating knowledge in the form of a temporary activation of nodes in the control layer (map[1] in Figure 3).

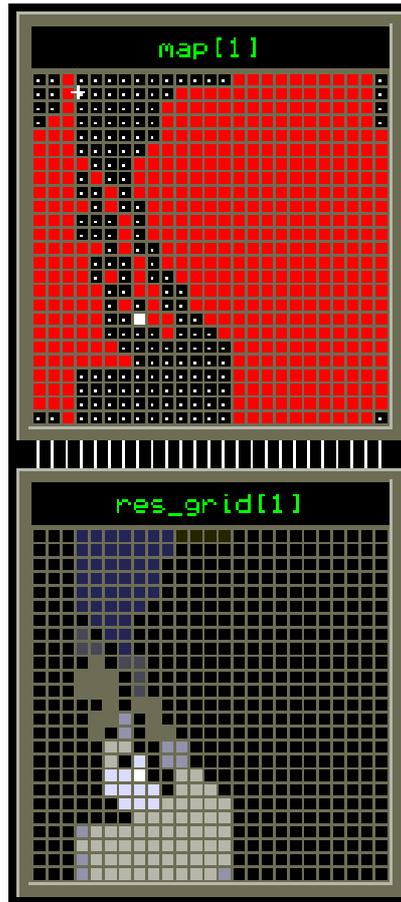


Figure 3. Neural resistive grid (bottom) and its control layer map[1] (top). The initial arm configuration is indicated by a cross in map[1]. The white node indicates the goal.

A planned sequence of movements is obtained by: i) reading out the direction of the potential gradient at the neurone corresponding to the current state, ii) selecting the next state, iii) executing the corresponding movement. The resulting trajectory in configuration space is shown in Figure 5, and the resulting arm trajectory in the workspace is shown in Figure 6.

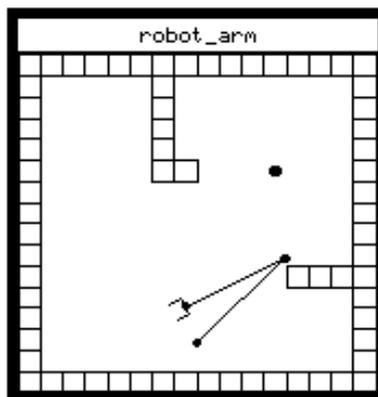


Figure 4. Robot arm in its initial position in the workspace.

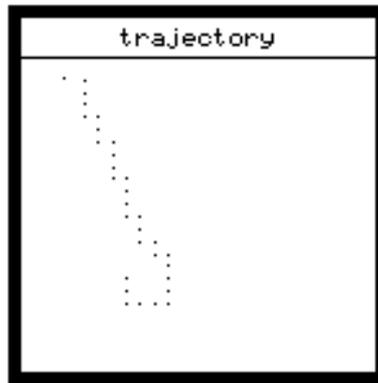


Figure 5. Trajectory in configuration space. The x-axis is the shoulder joint angle. The y-axis is the elbow joint angle.

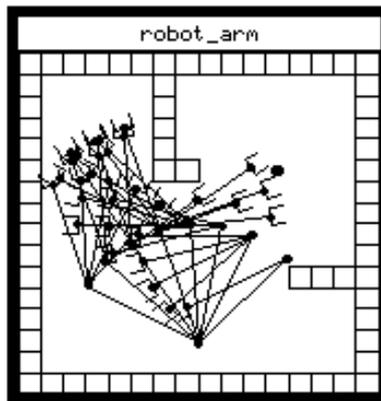


Figure 6. Sequence of arm positions corresponding to the trajectory in Figure 5.

The main advantage of the resistive-grid method is that it guarantees finding a solution to the planning problem, if there is one. However, it has also the limitations of grid-based methods: a limited resolution. In Figure 3, a 25 x 25 grid is used to represent all possible combinations of the two joint angles. As the joints can rotate by a full 360 degree, there is a 14.4 degree angular step change from one node to the next. This has two consequences, i) the planned movement is very saccadic (Figure 5 and 6), ii) the goal is not reached (Figure 6) because its exact location does not necessarily correspond to the centre of a node in the grid. It is indeed possible to increase the resolution of the grid but at the cost of a higher computational load. Another solution is to use the goal nodes (target_trans) in the sensori-motor association network presented in the next section.

5.2 A Normalized RBF net representing sensory-motor associations.

In this simplified example, sensory inputs are measurements of the joint angles and motor outputs are desired joint angles. This input corresponds to proprioceptive information rather than visual input, but the network operates in the same way in both cases³. The inputs and outputs are read out from the planner. For each node in the resistive grid (corresponding to the current values α_1 , α_2 of joint angles), the planner determines the best “next node” (corresponding to the next set of values α'_1 , α'_2). These transitions are learned by the sensory motor association network in figure 7.

³ Provided the representation of visual information generalises in such a way that the activity at the input of the net diminishes gradually as the image is transformed.

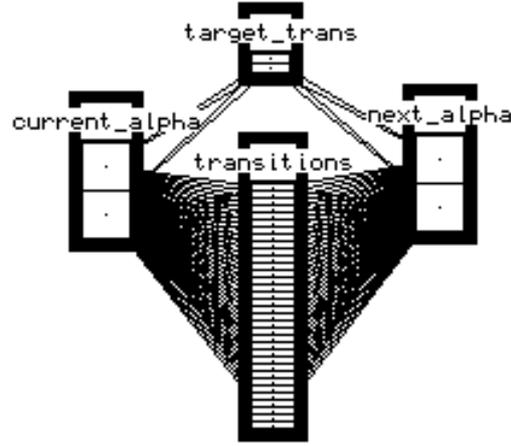


Figure 7. Neural network for Sensory Motor Association.

The sensory motor association network has two active layers, as depicted in Figure 7. The hidden layer, named "transitions", is composed of N_h neurones with Normalized Radial Basis Functions (NRBF) (Bugmann, 1998). In this network, these "neurones" have two joint angles α_1 and α_2 as input. They respond maximally when these angles are equal to pre-set values α_{10}^j and α_{20}^j (the centre of their "receptive field" in the configuration space). The width of the tuning of their response y_j (or size of the receptive field) is determined by σ_j :

$$y_j = \exp\left[-\frac{1}{2\sigma_j^2} \sum_i (\alpha_i^j - \alpha_{i0}^j)^2\right]. \quad (3)$$

These neurones project with weights W_{ij} to the output layer ("next_alpha" in the Figure 7). The output layer comprises two nodes with outputs α'_1 and α'_2 . These outputs are calculated as follows:

$$\alpha'_i = \frac{\sum_j W_{ij} y_j}{\sum_j y_j} \quad (4)$$

This operation corresponds to the weighted average over the input weights, where the input activities are used as weights for the weighted average. It may be possible to approximate such a function with biological neurones, assuming feedforward inhibition. The function (4) allows the encoding of the values of the next angles in the *weights* from hidden to output layer.

The network is trained during the read-out of a sequence of movements as determined by the resistive grid. At each transition from one node n_t to the next, n_{t+1} , a new RBF node is recruited, its receptive field is centred on the starting point in configuration space corresponding to the node n_t , and its *weights* to the output layer are set to the values of the joint angles corresponding to node n_{t+1} .

During replay, the output of this network alone is used to control the arm movement (The resistive grid is not used). As soon as one configuration has been reached, the network provides the next values of joint angles. As NRBFs have overlapping receptive fields, a small number of nodes in the hidden layer are simultaneously active and may point to different future configurations. However, the function (4) of the output nodes allows the most *active* of the inputs to have the largest *weight* in the decision.

To solve the problem of inaccurate end position, two extra-nodes "target_trans" are used (Figure 7). These nodes have receptive fields centred on the two possible final configuration of the arm, and have their output weights pointing to these exact final configurations. (There are two combinations of joint angles allowing the wrist to touch the filled circle, but one turns out to be forbidden, because the arm would have to go through the short wall protruding from the right in Figure 4). Thereby, when the arm approaches one of the selected final configurations, it becomes strongly attracted to the exact final position.

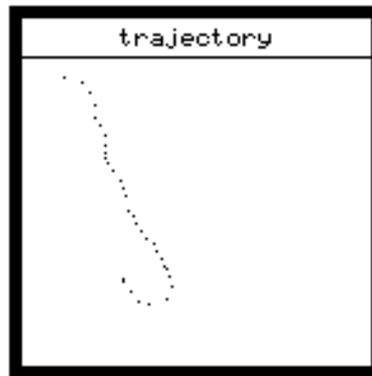


Figure 8. Trajectory in configuration space produced by the trained Sensory Motor Association network.

Figure 8 shows the trajectory in configuration space resulting from a movement controlled by the sensori-motor association network. Figure 9 shows the corresponding sequence of positions in workspace. One may note the much smoother trajectories and the correct final position of the arm.

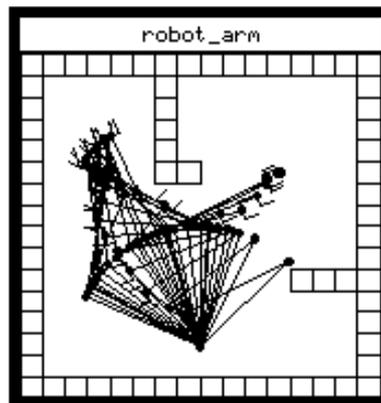


Figure 9. Sequence of arm positions corresponding to the trajectory in Figure 8.

The sequence is not replayed with the same number of steps as the sequence used for the training (Figure 5). In practice, the next position of the arm is calculated using current inputs. Then the arm is set to its new position, which modifies the inputs to the RBF nodes, and the new position is calculated. The speed at which the complete movement is performed depends on the size of the receptive fields of the RBF nodes. If these are too large, there is little change in the response of the RBF nodes after a movement and, conversely, little change in the calculated next position. If the receptive fields are too small, the smoothness of the trajectory is lost. For the results shown here $\sigma = 2\pi/25$ was used.

5.3 Generalisation

NRBF's nets behave as nearest neighbour classifier (Bugmann, 1998). This has the consequence that inputs far from any input value used during training produce the response ("next position") corresponding to the nearest input in the training set. Figure 10 illustrates the trajectories in configuration space generated by the sensory-motor association network for various initial positions.

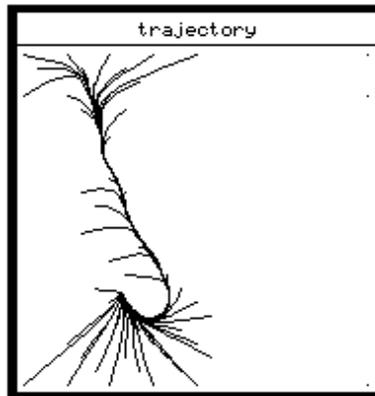


Figure 10. Example of generalisation: Trajectories in configuration space generated by the sensory-motor association network for various initial arm configurations. It should be noted that the network was trained solely with the trajectory data shown in figure 5.

The particular type of sensory-motor association network used here needs to be trained only on *one trajectory* produced by the planner, and can lead an arm to the goal from *any* initial configuration. Hence, the need for planning a field of sensory-motor associations that is mentioned in section 4 might not be a practical problem.

Generalisation also relates to the question of habitual behaviour. Habitual behaviour is generally noted when it leads to inappropriate actions, such taking the usual route to work after leaving home to go to the supermarket. When inspecting the details of the actions, those differ from previous day, being flexible responses to varying traffic conditions. Hence inadequate habitual behaviour seems to contain a planning component and could be described as resulting from planning towards an inadequate goal rather than being a set of inflexible stereotypical actions. Within the context of the model proposed here, the interpretation would be that, without specific goal being held actively in prefrontal areas, the more used (hence stronger) sensory-motor associations are dominating the behaviour. If the neural SM encoding mechanism had the generalisation properties described above, then no planning in the traditional sense needs taking place, but flexible responses to unpredicted situations are possible

6. Conclusions.

Observations of deficits related to cerebellar lesions have suggested that action plans may be selected from a pool of pre-calculated or task domain-specific plans. The positive side of such scheme is to enable fast skilled behaviour by eliminating the planning from the perception-to-action processing. On the negative side, it also enables the initiation of involuntary complex behaviour, as possibly revealed by inappropriate habitual behaviours. As for the underlying circuit, it is proposed that input from the cerebellum to the PPC may pre-select responses that are then triggered by sensory inputs. Inputs to frontal areas may possibly pre-select whole plans. This scheme is very speculative as a model of cerebellar function, but it can lead to the design of novel control system architectures with useful properties, as exemplified by an artificial neural network implementation.

From a control theory perspective, the proposed scheme unifies i) behaviour-based control that is fast but with limited flexibility in terms of goals and ii) classical plan-based control which is flexible but slow due to the computational overhead.

Due to the inclusion of off-line planning, the proposed approach differs from previous proposals viewing action selection as taken from a pre-existing pool of actions (Jeannerod, 1997) or schemata (Shallice, 1988).

References:

- Alhoefter K. & Bugmann G. (1995) "Planning and Learning Goal-Directed Sequences of Robot-Arm movements", Proc. of ICANN'95, Paris, Vol. 1, 449-454.
- Braitenberg V. (1997) "The detection and generation of sequences as a key to cerebellar function: Experiments and theory", *Behavior and Brain Science*, 20:2, pp. 229 et seq.
- Bugmann, G. (1998) "Normalized Radial Basis Function Networks", *Neurocomputing (Special Issue on Radial Basis Function Networks)*, 20, pp. 97-110.
- Bugmann, G., Taylor, J.G. and Denham M. (1995) "Route finding by neural nets" in Taylor J.G (ed) "Neural Networks", Alfred Waller Ltd, Henley-on-Thames, p. 217-230
- Bugmann G., Koay K.L., Barlow N., Phillips M. and Rodney D. (1998) "Stable Encoding of Robot Trajectories using Normalised Radial Basis Functions: Application to an Autonomous Wheelchair", Proc. 29th Intl. Symp. Robotics (ISR'98), 27-30 April, Birmingham, UK, pp. 232-235.
- Butler E.G., Horne M.K. & Hawkins N.J. (1992) "The activity of monkey thalamic and motor cortical-neurons in a skilled ballistic movement", *Journal of Physiology-London*, 445, pp. 25-48.
- Day B.L., Thompson P.D., Harding A.E. & Marsden C.D. (1998) "Influence of vision on upper limb reaching movement in patients with cerebellar ataxia", *Brain*, 121, pp. 357-372.
- Grafman J., Litvan I., Massaquoi S., Stewart M., Sirigu A. & Hallet M. (1992) "Cognitive planning deficits in patients with cerebellar atrophy", *Neurology*, 42:8, pp. 1493-1496.
- Hollands M.A. & Marple-Horvat D.E. (1996) "Visually guided stepping under conditions of step cycle-related denial of visual information", *Exp. Brain Res.*, 109, pp. 343-356.
- Jeannerod M. (1997) "The Cognitive Neuroscience of Action", Blackwell Pub;
- Lu X., Hikosaka O. & Miyach S. (1998) "Role of monkey cerebellar nuclei in skill for sequential movement", *J. of Neurophysiology*, 79:5, pp 2245-2254.
- Marple-Horvat D.E., Hollands M.A., Crowdy K.A. & Criado J.M (1998) "Role of cerebellum in visually guided stepping", *J. of Physiology*, 509P, pp. 28S.
- Middleton F.A. & Strick P.L. (1998) "Cerebellar output: motor and cognitive channels", *Trends in Cognitive Science*, 2:9, pp. 348-354.
- Okuda B. (1997) "Early morphological changes in the thalamocortical projection onto the parietal cortex following ablation of the motor cortex in the cat", *Brain Research Bulletin*, 44:3, pp. 281-287.
- Shallice, T. (1988) "From neuropsychology to mental structure", New York: Cambridge University Press
- Silveri M.C., Misciagna S., Leggio M.G. & Molinari M. (1997) "Spatial dysgraphia and cerebellar lesion: A case report", *Neurology*, 48:6, pp. 1529-1532.
- Snider L.H., Batista A.P. & Andersen R.A. (1997) "Coding of intention in the posterior parietal cortex", *Nature*, Vol. 386, no 6621, pp. 167-170.
- Stein J.F. & Glickstein M.G. (1992) "Role of the cerebellum in visual guidance of movement", *Physiological Reviews*, 72:4, pp. 967-1017.
- Thomson J.A. (1980) "How do we use visual information to control locomotion", *Trends in Neuroscience*, October issue, pp. 247-250.