

Neural network models of velocity storage in the horizontal vestibulo-ocular reflex

T. J. Anastasio

Department of Otolaryngology and Head and Neck Surgery, University of Southern California, Los Angeles, CA 90033, USA

Received March 19, 1990/Accepted in revised form July 18, 1990

Abstract. The vestibulo-ocular reflex (VOR) produces compensatory eye movements by utilizing head rotational velocity signals from the semicircular canals to control contractions of the extraocular muscles. In mammals, the time course of horizontal VOR is longer than that of the canal signals driving it, revealing the presence of a central integrator known as velocity storage. Although the neurons mediating VOR have been described neurophysiologically, their properties, and the mechanism of velocity storage itself, remain unexplained. Recent models of integration in VOR are based on systems of linear elements, interconnected in arbitrary ways. The present study extends this work by modeling horizontal VOR as a learning network composed of nonlinear model neurons. Network architectures are based on the VOR arc (canal afferents, vestibular nucleus (VN) neurons and extraocular motoneurons) and have both forward and lateral connections. The networks learn to produce velocity storage integration by forming lateral (commissural) inhibitory feedback loops between VN neurons. These loops overlap and interact in a complex way, forming both fast and slow VN pathways. The networks exhibit some of the nonlinear properties of the actual VOR, such as dependency of decay rate and phase lag upon input magnitude, and skewing of the response to higher magnitude sinusoidal inputs. Model VN neurons resemble their real counterparts. Both have increased time constant and gain, and decreased spontaneous rate as compared to canal afferents. Also, both model and real VN neurons exhibit rectification and skew. The results suggest that lateral inhibitory interactions produce velocity storage and also determine the properties of neurons mediating VOR. The neural network models demonstrate how commissural inhibition may be organized along the VOR pathway.

Introduction

The vestibulo-ocular reflex (VOR) stabilizes the visual image by producing compensatory eye movements that

are nearly equal and opposite to head movements (Wilson and Melvill Jones 1979). The VOR operates by using head rotational velocity signals from the semicircular canals to control contractions of the extraocular muscles. But despite its apparent simplicity, VOR is more than just a relay.

The VOR has been well studied in monkeys. In these animals, the eye movement response of horizontal VOR has a longer time course than the canal signal that drives it. The dominant time constant of canal afferents in monkeys is about six seconds (Fernandez and Goldberg 1971; Buettner and Waespe 1981). In contrast, the dominant time constant of primate VOR is about 20 s (Buettner et al. 1981; Paige 1983). This time constant lengthening is known as velocity storage (Raphan et al. 1979). Velocity storage is magnitude dependent; the VOR time constant decreases as stimulus magnitude increases (Paige 1983; Fetter and Zee 1988). The mechanism of velocity storage and the cause of its nonlinear behavior have not been elucidated.

The VOR is mediated by the three-neuron-arc, which is composed of semicircular canal afferents, vestibular nucleus (VN) neurons and motoneurons of extraocular muscles (Wilson and Melvill Jones 1979). In monkeys, many VN neurons have time constants that are similar in value to that of VOR (Buettner et al. 1978; Waespe and Henn 1979), indicating that velocity storage has occurred at the VN level. Compared to canal afferents, VN neurons have lower spontaneous discharge rates and higher head velocity sensitivities (gains) (Fuchs and Kimm 1975; Buettner et al. 1978; Waespe and Henn 1979). VN neurons also exhibit nonlinear properties such as rectification and skew (ibid.). Except for high gain (see below) the factors underlying these VN neuron response properties have not been identified.

The VOR is bilaterally symmetric (Wilson and Melvill Jones 1979). The semicircular canals work in push-pull pairs, and the extraocular muscles are arranged in agonist-antagonist pairs. The VNs are also arranged bilaterally and, in mammals, are interconnected by an extensive, inhibitory commissural system

(*ibid.*). Vestibular commissural connections are critical for velocity storage; sectioning them results in the loss of velocity storage in monkeys (Blair and Gavin 1981).

Early, lumped models represented the velocity storage mechanism as a leaky integrator, either in parallel with a direct pathway (Raphan et al. 1979) or in a feedback loop (Robinson 1981). Recent models demonstrated how VOR integration could be produced by a bilateral set of VN neurons that inhibit each other via vestibular commissures (Cannon et al. 1983; Galiana and Outerbridge 1984). Mutual inhibition results in net positive feedback that allows VN neurons to precess (integrate) the push-pull canal signals but not their spontaneous rates (*ibid.*). These models contained linear neurons only and interconnected them in arbitrary, predetermined ways. In contrast, real VN neurons have nonlinear response properties (Waespe and Henn 1979) and may be interconnected in highly complex ways.

Neural network learning algorithms, such as back-propagation, allow initially randomly connected sets of linear or nonlinear units to be trained to perform various transformations (Rumelhart et al. 1986). Previous learning network models of VOR employed nonlinear units but had forward connections only, and were therefore limited to static transformations without temporal dynamics (Anastasio and Robinson 1989 a, b; 1990). An extension of back-propagation has recently been developed (Williams and Zipser 1989) that trains continually running networks having feedback (recurrent) as well as forward connections. The purpose of this study was to construct neural network models of horizontal VOR that were composed of nonlinear units, with feedback as well as forward connections. The networks were trained to produce velocity storage using recurrent back-propagation.

The goal of modeling the relatively simple VOR is not to demonstrate the sophistication of learning algorithms. Instead, recurrent back-propagation is used to construct neural network models of VOR that can explain some of its properties. The networks learn to produce the VOR transformation by forming lateral inhibitory feedback loops among VN neurons that overlap and interact in complex ways. The response properties of the overall networks and individual VN neurons correspond closely to their real counterparts in monkey. The results suggest that lateral inhibitory interactions produce velocity storage and also determine the properties of VOR and VN neurons.

Methods

Network architecture

The VOR was modeled as a three layered, recurrent neural network. Units in the input layer corresponded to canal afferents. Output units represented extraocular motoneurons. The hidden (middle) layer contained units corresponding to VN neurons. No bias units were used. The input and output layers each contained

two units. The number of units in the hidden layer was varied for different simulations.

Network architectures reflected the known anatomy of mammalian VOR (Wilson and Melvill Jones 1979). To represent the three-neuron chain, every input and hidden unit projected to every hidden and output unit, respectively. Because canal afferents do not project to motoneurons in mammals (*ibid.*), no direct input-to-output connections were allowed. Anatomical studies indicate that many VN neurons mediating VOR send collateral projections to contralateral VN, but do not send recurrent collaterals to themselves or to ipsilateral VN (McCrea et al. 1987). Vestibular commissures were therefore modeled as interconnections between contralateral hidden units. Hidden units were not allowed to send recurrent collaterals to themselves or to other ipsilateral hidden units. The allowed connections are illustrated for a four-hidden-unit model in Fig. 1. Other possible connections, which have not been generally described in mammals, have been omitted.

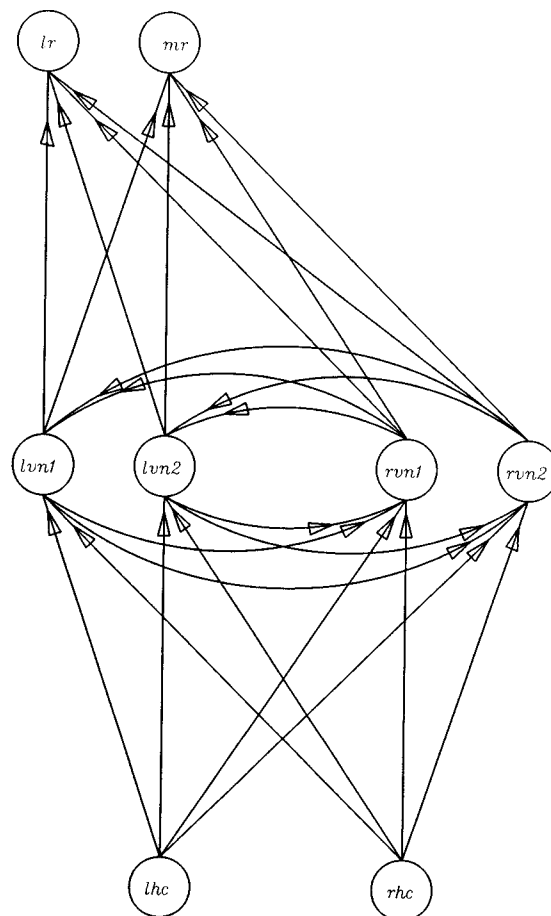


Fig 1. Schematic diagram of the four-hidden-unit, horizontal VOR neural network model. Both input units project to all hidden units and all hidden units project to both output units. Hidden units on one side project to hidden units on the other; same side hidden unit connections are not allowed. *lhc* and *rhc*, left and right horizontal canal afferents; *lvn* and *rvn*, left and right vestibular nucleus neurons, *lr* and *mr*, lateral and medial rectus motoneurons of the left eye

The operation of the networks was adapted from the formalism established by Williams and Zipser (1989). Input units, which do not receive any connections, are treated separately from hidden and output units. Let I designate the set of input units and U the set of hidden and output units, and let all the units be indexed by k . Then the state (output or firing rate) x , y and z of each unit at any time t is

$$z_k(t) = \begin{cases} x_k(t) & \text{if } k \in I \\ y_k(t) & \text{if } k \in U \end{cases} \quad (1)$$

where indices for x and y correspond to those of z .

Let W denote the weight matrix for each network. The element w_{ij} represents the weight (strength) of the connection to the i th hidden or output unit ($i \in U$) from any unit j ($j \in U \cup I$). For the purpose of modeling VOR, the weight matrices were constrained and did not allow every connection, as described above.

The net input $s_k(t)$ to any unit $k \in U$ is given by

$$s_k(t) = \sum_{i \in U} w_{ki}y_i + \sum_{i \in I} w_{ki}x_i = \sum_{i \in U \cup I} w_{ki}z_i(t) \quad (2)$$

and its output at the next discrete time step (network cycle or tick) is expressed in terms the net input by

$$y_k(t+1) = f(s_k(t)) = 1/(1 + \exp(-s_k(t))) \quad (3)$$

where f is the sigmoidal squashing function (Rumelhart et al. 1986; Anastasio and Robinson 1989a). The squashing function produces unit output responses which range from zero to one, and are linear for midrange inputs but are driven to cut-off (zero) or saturation (one) for large negative or positive inputs, respectively. The squashing function unit outputs are 0.50 for net inputs of zero. Equations (2) and (3) specify the entire (discrete-time) dynamics of the network.

Learning algorithm

The goal of the learning procedure is to adjust the weights so that the error between desired and actual network outputs is reduced. Like the more familiar back-propagation algorithm, this is done by computing the error gradient in weight space, and using this gradient to update the values of the weights. For dynamic trajectories, the gradient is also a function of time. The derivation of the learning algorithm is presented elsewhere (Williams and Zipser 1989). Only equations actually used to train networks will be reproduced.

Let T be the subset of output units in U for which desired target values $d_k(t)$ exist. Then the errors $e_k(t)$ are

$$e_k(t) = \begin{cases} d_k(t) - y_k(t) & \text{if } k \in T \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

For the simulations presented below, both output units had target values. The weight changes are computed as

$$\Delta w_{ij}(t) = \alpha \sum_{k \in U} e_k(t) p_{ij}^k(t) \quad (5)$$

where α is some fixed, positive learning rate. The $p_{ij}^k(t)$ values are computed iteratively for all $k \in U$, $i \in U$ and $j \in U \cup I$ by

$$p_{ij}^k(t+1) = f'(s_k(t)) \left[\sum_{i \in U} w_{ki} p_{ij}^i(t) + \delta_{ik} z_j(t) \right] \quad (6)$$

where δ_{ik} is the Kronecker delta ($\delta_{ik} = 1$ if $i = k$, 0 otherwise) and $f'(s_k(t))$ is the first derivative of the squashing function in (3). Finally, each weight is updated by

$$w_{ij}(t+1) = w_{ij}(t) + \Delta w_{ij}(t) \quad (7)$$

Training the networks began by assigning modifiable weights in W a value chosen at random from a uniform distribution between -1 and $+1$. All p_{ij}^k values were initially set to zero. The networks were trained on two sets of inputs and desired outputs, which were presented repeatedly in random order. The input/desired output sets consisted of continuous sequences of values, as described below.

At the start of each learning cycle, input units were initialized with the next values in the input sequence, and the response of the network to this new input, and to the previous states of the hidden units, was calculated (2 and 3). The errors between actual and desired outputs for the new inputs were computed (4). The new unit states were used to update the p_{ij}^k values (6), and the p_{ij}^k values were used to compute the weight changes (5). Finally, modifiable weights were updated (7) and the process was repeated for the next inputs and desired outputs in the sequence. Following training on a given set, the performance of the network on both sets was determined, and mean squared error between desired and actual output sequences was computed. Training continued until the mean squared error for six consecutive set presentations was less than a tolerance, which was 0.001 unless otherwise stated.

The actual output and hidden unit responses are characterized by their spontaneous rates (SRs), defined as a units response to constant inputs of 0.50. Units are further characterized by their gains and time constants in both excitatory and inhibitory directions. Gain is the absolute value of the maximum change in activity from SR of any unit divided by the change from 0.50 of the inputs. (The inputs have a gain of one, by definition.) Although actual output and hidden unit responses are not pure, single exponential decays, their time constants were estimated by finding the negative inverse of the slope of the line relating number of ticks to the semi-ln transformed response values.

Results

A VOR neural network model with four hidden units

A neural network model having only two hidden units could not learn the VOR transformation in a reasonable training time (less than 10,000 passes through the training set). The smallest, symmetrical network capable of consistently learning the VOR transformation

within 10,000 passes consisted of two input, four hidden and two output units (Fig. 1). Input units represented afferents from the left and right horizontal semicircular canals (*lhc* and *rhc*). Output units represented the motoneurons of the lateral and medial rectus muscles of the left eye (*lr* and *mr*). Hidden units represented VN neurons on the left (*lvn1* and *lvn2*) and right (*rvm1* and *rvm2*) sides of the brainstem.

To represent the neural control of VOR, input unit activities reflected head velocity signals from canal afferents, while output unit activities represented motoneuron eye velocity commands. Output responses were equal and opposite to the inputs, but had time constants four times longer, reflecting velocity storage. Only input and desired output activities were specified; actual output and hidden unit responses were the result of learning.

During VOR, real motoneurons have discharge components related to eye velocity and position, and to fast-phases of nystagmus (Skavenski and Robinson 1973). Canal afferents carry signals proportional to head velocity (Fernandez and Goldberg 1971). Canal velocity signals are converted to longer duration eye velocity commands by the velocity storage mechanism (Raphan et al. 1979). Eye position commands are derived from eye velocity commands by the velocity-to-position integrator, while fast-phases are introduced by a pulse generator (Robinson 1975). Only the velocity storage mechanism is considered in this study; the velocity-to-position integrator and pulse generator are beyond its focus. Therefore, desired output consists

only of the eye velocity component of motoneuron discharge, without eye position or fast-phase related activity.

The training set consisted of two impulse head accelerations, one to the left and the other to the right. Figure 2A shows the inputs (solid) and desired outputs (dashed) for *lhc* and *lr*; those for *rhc* and *mr* are shown in Fig. 2B. Each response sequence is thirty ticks long. Inputs and desired outputs were modulated in push-pull about an SR of 0.50. They had an amplitude of 0.10, which was within the linear range of the sigmoidal unit transfer function. Head acceleration to the left (Fig. 2, A and B, ticks 1 through 30) excited *lhc* and inhibited *rhc*. The corresponding compensatory eye movement would be produced by inhibition of *lr* and excitation of *mr*. The opposite pattern obtained for rightward head acceleration (Fig. 2, A and B, ticks 31 through 60). Inputs and desired outputs decayed exponentially and had time constants of one and four ticks, respectively. If each tick represents 5 s, then this corresponds to a 5 s time constant for canal afferents and a 20 s time constant for VOR. Because of the three layered architecture of VOR, a delay of one tick was introduced between input and output.

In initial simulations, all allowed connections (including all connections between hidden units) were modifiable. Learning required about 5,000 training set presentations. In all cases, the learned input-to-hidden connections were arranged in push-pull; each input unit excited one hidden unit and inhibited the other. The




Fig. 2A–D. Training the network to produce the horizontal VOR with velocity storage. In this and subsequent plots, unit responses to a left (network cycles 1–30) followed by a right (network cycles 31–60) impulse head acceleration are shown. **A and B:** Inputs (solid) and desired outputs (dashed) are plotted, along with the actual outputs of the initially randomized network (dotted). Waveforms are for *lhc* and *lr* in A and for *rhc* and *mr* in B. Inputs have a gain of one (by definition) and time constant of one network cycle (tick). Desired outputs have a gain of one and time constant of four ticks. **C and D:** Following training, actual outputs (solid) match those desired (dashed). Waveforms are for *lr* in C and for *mr* in D. Removing the commissural connections slightly increases the gain of the outputs and reduces their time constant to one tick (dotted), thus eliminating velocity storage. This demonstrates that velocity storage is subserved by commissural inhibition in the network. Abbreviations as in Fig. 1

hidden units also learned to reciprocally innervate the outputs. (Note that when any one unit both excites and inhibits target units, an interneuron with unity gain is implied to satisfy Dale's law). This arrangement reflects the actual organization of VOR (Wilson and Melvill Jones 1979) and is similar to the learned connections of purely feedforward VOR neural network models (Anastasio and Robinson 1989a, b; 1990). Because evidence to data suggests plastic modification of synapses only at the VN (Lisberger 1988), the hidden-to-output weights were fixed in subsequent simulations. The fixed hidden-to-output weights had the same reciprocal pattern as the learned weights, but all had equal, arbitrarily chosen absolute values.

In the four-hidden-unit network, the fixed hidden-to-output weights had the absolute value of 0.500; their reciprocal pattern is shown in Table 1. The hidden units (*lvn1* and *lvn2*) that are inhibitory to *lr* and excitatory to *mr* are by construction located on the left side of the brainstem; the hidden units on the right side (*rvn1* and *rvn2*) have the opposite output projection pattern. By lateralizing the hidden units in this way, the constraint could be placed that commissural connections occur only between contralateral hidden units. These learned commissural (hidden-to-hidden) weights were primarily inhibitory, but in some simulations small excitatory commissures could develop. Because the vestibular

commissures in mammals are believed to be inhibitory (Wilson and Melvill Jones 1979), commissural weights were constrained to be zero or less in subsequent simulations.

The input-to-hidden and commissural weights were randomized and then modified by the algorithm. The incorrect output unit responses of the initially randomized network are shown in Fig. 2, A and B (dotted). The network learned to produce the desired output after 5,369 training set presentations. The learned output responses (Fig. 2, C and D, solid) match those desired (Fig. 2, C and D, dashed) to within the required tolerance. The final weights are presented in Table 1; the weights for all other runs of this simulation are similar. The input units reciprocally innervate the hidden units, as before. The negative commissural connections form an approximately diagonally symmetrical submatrix.

Analyzing the four-hidden-unit VOR neural network model

In the mature, intact network, SR is 0.50, gain is 0.99 and time constant is 4.26 for both output units in response to impulse head accelerations in both directions (Table 2; Fig. 2, C and D, solid). Removing the commissural connections reduces output time constant

Table 1. Learned weight matrix for the four-hidden-unit, horizontal VOR neural network model with velocity storage. *lhc* and *rhc*, left and right horizontal canal afferents; *lvn* and *rvn*, left and right vestibular nucleus neurons; *lr* and *mr*, lateral and medial extraocular muscle motoneurons of the left eye

from:	<i>lhc</i>	<i>rhc</i>	<i>lvn1</i>	<i>lvn2</i>	<i>rvn1</i>	<i>rvn2</i>
to:						
<i>lvn1</i>	5.825	-5.958	0.000	0.000	-4.595	-1.383
<i>lvn2</i>	4.728	-6.707	0.000	0.000	-0.564	-0.001
<i>rvn1</i>	-6.225	5.820	-4.741	-0.940	0.000	0.000
<i>rvn2</i>	-6.452	4.902	-1.172	-0.002	0.000	0.000
<i>lr</i>	0.000	0.000	-0.500	-0.500	0.500	0.500
<i>mr</i>	0.000	0.000	0.500	0.500	-0.500	-0.500

Table 2. Analysis of unit responses in the four-hidden-unit, horizontal VOR neural network model with velocity storage. CI, total amount of commissural inhibition received by each unit; SR, spontaneous rate; Gex and Gin, impulse gains in the excitatory and inhibitory directions; Tex and Tin, time constants estimated from impulse responses in the excitatory and inhibitory directions. Other abbreviations as in Table 1

unit	CI	SR	Gex	Gin	Tex	Tin
<i>lhc</i>	—	0.50	1.00	1.00	1.00	1.00
<i>rhc</i>	—	0.50	1.00	1.00	1.00	1.00
<i>lvn1</i>	-5.978	0.21	2.67	1.76	4.23	4.42
<i>lvn2</i>	-0.566	0.25	2.61	1.54	3.87	3.87
<i>rvn1</i>	-5.681	0.19	2.56	1.56	4.22	4.43
<i>rvn2</i>	-1.174	0.26	2.63	1.61	4.08	4.08
<i>lr</i>	—	0.50	0.99	0.99	4.26	4.26
<i>mr</i>	—	0.50	0.99	0.99	4.26	4.26

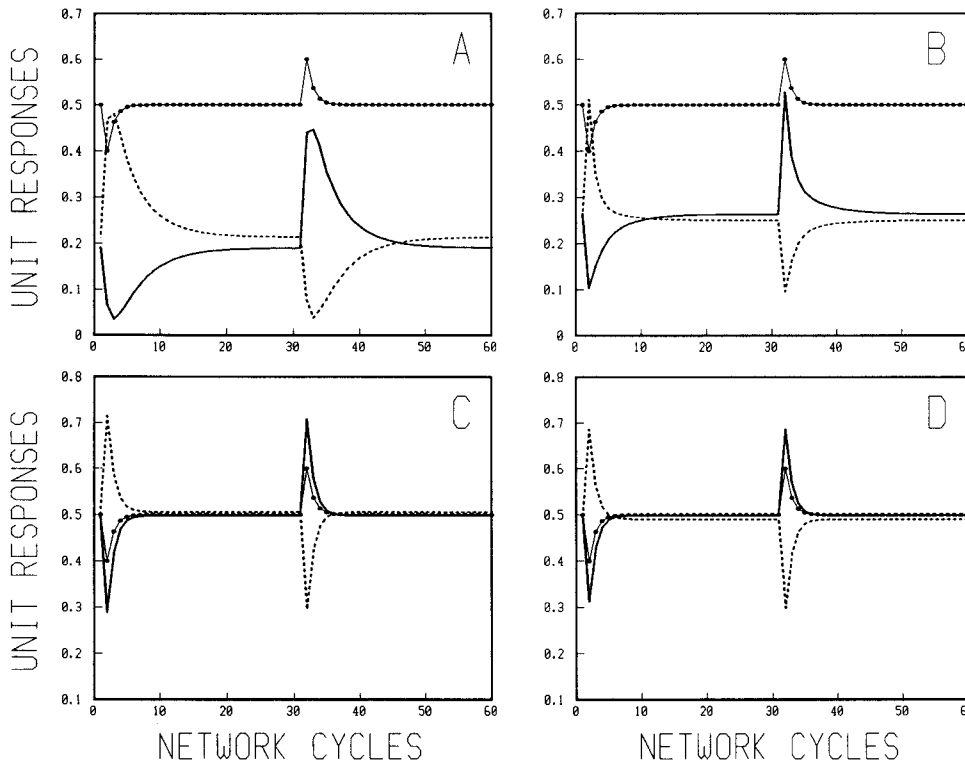


Fig. 3A–D. Hidden unit responses in horizontal VOR neural network models with and without velocity storage. A–D: Hidden unit response waveforms are for *rvn1* (solid) and *lvn1* (dashed) in A and C, and for *rvn2* (solid) and *lvn2* (dashed) in B and D. One input waveform is shown in all plots for reference (*rhc*, dotted). A and B: In models trained to produce VOR with velocity storage, hidden units have higher gains, lower spontaneous rates, longer time constants and asymmetric responses as compared to input units. Also, hidden units form separate slow and fast pathways, based on their initial responses (A, slow, rounded; B, fast, peaked). C and D: In models trained to produce VOR without velocity storage, hidden units also have higher gains, but other properties are similar to input waveforms. These results demonstrate that longer time constants, lower spontaneous rates and asymmetric responses of hidden as compared to input units are all associated with velocity storage. Abbreviations as in Fig. 1

to that of the canal inputs: 1.00 tick (Fig. 2, C and D, dotted). Thus velocity storage is completely abolished by commissurotomy. This demonstrates that velocity storage is subserved by commissural inhibition in the network. Commissurotomy also slightly increases output gain (to 1.20) but SR is unaffected. Similar effects are observed following commissurotomy in monkeys (Blair and Gavin 1980).

The impulse responses of hidden units in this network are shown in Fig. 3, A (*rvn1*, solid; *lvn1*, dashed) and B (*rvn2*, solid; *lvn2*, dashed); one input waveform is also shown for reference (*rhc*, dotted). A quantitative analysis of these responses is given in Table 2. In comparison with the inputs, these hidden units have higher gains, longer time constants, lower SRs and asymmetric response to oppositely directed inputs. The time constants of hidden and output units are similar. Low SR forces the hidden units to operate near the bottom of the sigmoidal unit transfer function. Thus hidden unit gain is higher in excitatory than inhibitory directions, and decay rate can be slightly prolonged following inhibitory responses (Table 2).

In another simulation, the four-hidden-unit network was trained to produce a VOR without velocity storage (inputs and outputs had the same time constant of one tick). This network was trained to a tolerance of 0.0001. Final input-to-hidden weights had the same organization as those in Table 1, but the commissural weights were all practically zero. The impulse responses of hidden units from this simulation are shown in Fig. 3, C (*rvn1*, solid; *lvn1*, dashed) and D (*rvn2*, solid; *lvn2*, dashed). In both VOR network models, whether trained

with or without velocity storage, hidden units had higher gains than the inputs. Higher hidden unit gain results in part from their push-pull innervation by input units (Table 1). In contrast, hidden units from VOR models without velocity storage (and without inhibitory commissures) are similar to inputs in time constant, SR and symmetry. These results demonstrate that, in VOR models with velocity storage, commissural inhibition accounts for the long time constants, low SRs and asymmetrical responses of the hidden units.

In the network trained to produce velocity storage, the two pairs of hidden units have different patterns of commissural connectivity. Hidden units *lvn1* and *rvn1* (Fig. 3A) have developed strong mutual inhibition (Table 1). Each unit in the pair inhibits the other and is inhibited by it in turn. When the inputs are both at SR, these hidden units inhibit each other approximately equally and their firing rates are constant. But when the inputs are modulated in push-pull, hidden units *lvn1* and *rvn1* exert net positive feedback on themselves, partially integrating the canal signals according to the values of their commissural weights. For example, during leftward head rotation, hidden unit *lvn1* is excited by the canal inputs while *rvn1* is inhibited (Fig. 3 A, ticks 1 through 30). Unit *lvn1* now inhibits *rvn1* more, and is inhibited less by it in turn, than with constant inputs. At the next network cycle, canal input is decaying rapidly, but *lvn1* continues to inhibit *rvn1* and so disinhibit itself, prolonging its excitatory response. The inhibitory response of *rvn1* is prolonged by the same process in reverse—*rvn1* inhibits *lvn1* less and is inhibited more by it in turn. This amounts to positive

feedback that allows *lvn1* and *rvn1* to persevere their activities for several network cycles. The opposite pattern applies for rightward head rotation (Fig. 3 A, ticks 31 through 60). The learned values of the commissural connections are such that the units only partially restore their firing rates after each cycle, making them leaky integrators. Thus the responses of *lvn1* and *rvn1* appear as low-pass filtered versions of the inputs, with rounded onsets and long decay times (Fig. 3A).

Hidden units *lvn2* and *rvn2* (Fig. 3B) have practically zero mutual inhibition (Table 1), so are not leaky integrators (low-pass filters) by themselves. However, they interact with *lvn1* and *rvn1* through moderately strong connections; *lvn2* is coupled to *rvn1*, and *rvn2* is coupled to *lvn1* (Table 1). The lack of mutual inhibition between *lvn2* and *rvn2* allows these units to have sharply peaked onsets, while their coupling to *lvn1* and *rvn1* provides them with longer decay times than those of the canal inputs. Based on their initial responses, hidden units form slow (*lvn1* and *rvn1*) and fast (*lvn2* and *rvn2*) pathways to output units. Removal of the coupling connections (between *lvn2* and *rvn1* and between *rvn2* and *lvn1*) reduces the time constant of fast pathway units to 1.00, but has little effect on the time constant of slow pathway units. Thus velocity storage in this network is subserved mainly by the strong, mutually inhibitory connections between *lvn1* and *rvn1*. The units receiving a small amount of commissural inhibition (*lvn2* and *rvn2*) nevertheless achieve time constant lengthening by coupling with the more strongly interacting units.

Nonlinear behavior of the four-hidden-unit VOR neural network model

In the network trained on a VOR with velocity storage, low SR forces the hidden units to operate closer to the bottom of the sigmoidal unit transfer function. Thus larger inputs can produce excitatory responses in the linear range, but drive inhibitory responses into the low-gain, nonlinear range or even into cut-off (rectification). Hidden unit cut-off temporarily disrupts commissural interactions, causing velocity storage to be reduced or eliminated. This in turn shortens the VOR time constant, increasing decay rate and reducing phase lag, and producing distortions of output responses.

The learned response of *lr* to the training inputs (amplitude 0.10) is reproduced in Fig. 4A (dotted). If the system were linear, it would be expected that output time constant should remain at about four ticks, regardless of input magnitude (Milsum 1966). The response of *lr* to inputs of amplitude 0.20 is shown in Fig. 4A (solid). This output response decays faster than would be expected, given a time constant of four ticks (Fig. 4A, dot-dash).

The more rapid decay of the actual response to the 0.20 amplitude inputs results because one slow path hidden unit is being driven into or near inhibitory cut-off by the larger amplitude stimulus. The response of *lr* to the 0.20 amplitude inputs is shown in Fig. 4C (solid). The excitatory phase of this response has about twice the amplitude of the response to the trained inputs (Fig. 4C, dotted) but is driven close to cut-off,

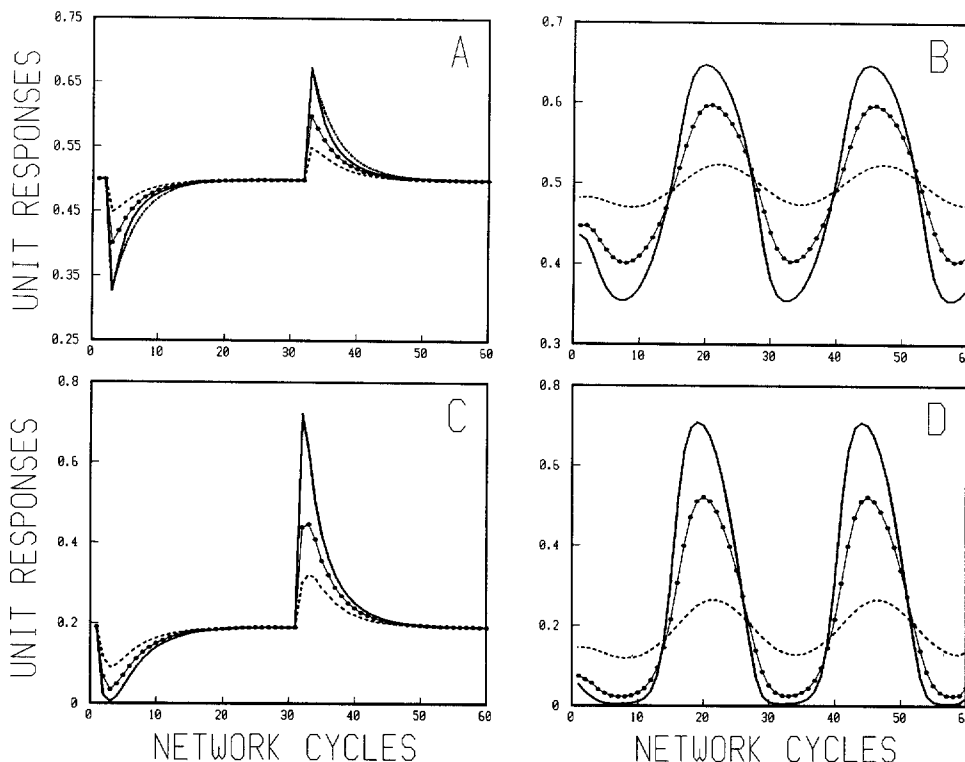


Fig. 4A–D. Nonlinear behavior of the horizontal VOR neural network model with velocity storage. **A** and **C**: Impulse responses of *lr* (**A**) and *rvn1* (**C**) to inputs of amplitude 0.05 (dashed), 0.10 (dotted) and 0.20 (solid). Also shown is the expected decay of *lr* from its peak response to the largest input with a time constant of four ticks (**A**, dot-dash). Output unit decay rate increases and hidden unit low-pass characteristics decrease as input magnitude increases. **B** and **D**: Responses of *lr* (**B**) and *rvn1* (**D**) to sinusoidal inputs of amplitude 0.01 (dashed), 0.05 (dotted) and 0.10 (solid). Phase lag decreases and skewness increases as input magnitude increases. Responses of *mr* and *lvn1* (not shown) are similar to those of *lr* and *rvn1*, but rotated by 30 network cycles. Fast pathway hidden units *rvn2* and *lvn2* (responses not shown) are less sensitive than slow pathway hidden units *rvn1* and *lvn1* to changes in input magnitude. Abbreviations as in Fig. 1

and into the low-gain region of the sigmoidal unit transfer function, at the peak, of the inhibitory phase. For a period of five ticks following the peak response, the time constant of the output units falls to 3.09 ticks. Following this period, output time constant returns to 4.29 ticks. The shorter time constant near peak results in a faster overall rate of decay for the outputs than expected. The response of *lr* to impulse inputs of amplitude 0.05 is shown in Fig. 4A (dashed). This response appears more rounded, and has a slightly increased time constant. The integrative properties of the slow path hidden units, revealed most clearly by their initial responses, are reduced as impulse input amplitude is increased. This is shown for unit *rvn1* in Fig. 4C, for input amplitudes of 0.05 (dashed), 0.10 (dotted) and 0.20 (solid).

The phase lag and shape of the response to sinusoidal inputs are also dependent upon input magnitude. Again, if the system were linear, the output response to a sinusoidal input at a given frequency should also be sinusoidal, and have the same phase difference regardless of input magnitude (Milsum 1966). The corner frequency of the VOR neural network models, having a time constant of four ticks, is 0.04 cycles per tick. At this frequency, the overall network should produce a phase lag at the output of 45 degrees, with the canals contributing about 14 deg at the front end. At 0.04 cycles per tick, each tick also equals about 14 deg. Since there is a pure delay of one tick between the input and output signals, the output responses should lag the inputs by about 45 deg, or three ticks.

The responses of *lr* and *rvn1* to sinusoidal inputs at a frequency of 0.04 cycles per tick is shown in Fig. 4, B and D, respectively, for input amplitudes of 0.01 (dashed), 0.05 (dotted) and 0.10 (solid). At an amplitude of 0.01, the output is sinusoidal and lags the input by three ticks, as expected for a linear system. Hidden units are not cutting-off at this amplitude level. For input amplitudes of 0.05 and 0.10, output responses are not sinusoidal. Phase lags relative to the inputs are reduced to two and one tick, respectively. Hidden units are driven into cut-off at these higher amplitude levels.

Besides loss of phase lag, output and hidden unit responses are skewed at higher input amplitude levels. For the low intensity inputs, hidden units never cut-off, and output responses are sinusoidal and phase lagged by three ticks. Output unit responses to higher intensity sinusoidal inputs can be thought of as alternating between two sinusoids, one phase lagged by three ticks (full velocity storage) and the other by one tick (no velocity storage) relative to the canal inputs. At midrange levels at and near output SR, while the hidden units are not cut-off, the outputs follow the more phase lagged sinusoid. But at peri-peak levels, during the hidden unit cut-off period, the outputs follow the less lagged sinusoid. This results in skewed sinusoids with the peaks shifted in lead relative to the midrange responses. Hidden units show a similar skewing, along with cut-off, at the higher input levels.

A VOR neural network model with eight hidden units

To explore a more distributed VOR neural network model, the hidden layer was expanded from four to eight units. The fixed hidden-to-output connections had the absolute value of 0.250. The eight-hidden-unit network was trained to produce the VOR transformation with velocity storage, requiring 2,567 training set presentations. Learned weights were generally similar to the previous case, with push-pull input-to-hidden connections and approximately symmetrical, inhibitory commissural connections. Commissural interactions included lateral inhibition and coupling to leaky integrators. But the commissural loops were more complex, and overlapped and interacted to a greater extent, in the eight- as compared to the four-hidden-unit network.

The learned outputs of the eight-hidden-unit network have parameters that are very close to those desired (Table 3). As in the previous example, hidden units in the eight-hidden-unit network have lower SRs, higher gains and longer time constants than input units. But hidden units are more nonuniform in their properties in the larger as compared to the smaller network (Tables 2 and 3). This is due in part to the more complex commissural interactions of the larger network. Hidden unit time constants vary over a broader range in the eight-hidden-unit network. As in the four-hidden-unit network, the estimated time constants of the hidden units in the larger network are not entirely dependent upon the amount of commissural inhibition each unit receives (for example, compare units *lvn4* and *rvn1* in Table 3). Again, this is due to inter-loop coupling. The hidden units in the larger network also form fast and slow pathways. Hidden units *lvn1* and *rvn2* have peaked initial responses and distinctly shorter time constants. The other hidden units have longer time constants and initial responses that range from peaked to rounded. Nonlinear behavior is similar in the eight- as in the four-hidden-unit network (responses not shown).

In initial eight-hidden-unit simulations, where the commissures were not constrained to be inhibitory, significant excitatory commissural connections sometimes developed. This resulted in hidden unit impulse responses having a peaked initial phase in one direction, which quickly reversed to a slowly decaying phase in the opposite direction. This type of VN neuron response has not been reported.

Discussion

Physiological veracity

A learning algorithm was used to construct neural network models of horizontal VOR that explain certain of its properties, specifically velocity storage. Of interest were the final networks resulting from learning, not the learning procedure itself. These networks conformed closely to the known anatomy and physiology of VOR. Network architectures were constrained

Table 3. Analysis of unit responses in the eight-hidden-unit, horizontal VOR neural network model with velocity storage. Abbreviations as in previous tables

unit	CI	SR	Gex	Gin	Tex	Tin
<i>lhc</i>	—	0.50	1.0	1.0	1.00	1.00
<i>rhc</i>	—	0.50	1.0	1.0	1.00	1.00
<i>lvn1</i>	-0.004	0.35	2.06	1.61	2.18	2.20
<i>lvn2</i>	-2.008	0.29	2.58	1.69	3.75	3.78
<i>lvn3</i>	-3.818	0.25	2.58	1.72	3.78	3.85
<i>lvn4</i>	-5.314	0.21	2.75	1.70	3.74	3.91
<i>rvn1</i>	-1.606	0.27	2.26	1.48	3.71	3.71
<i>rvn2</i>	-0.159	0.33	2.22	1.66	3.26	3.23
<i>rvn3</i>	-6.245	0.25	3.25	2.11	3.77	3.91
<i>rvn4</i>	-3.049	0.24	2.24	1.48	3.77	3.80
<i>lr</i>	—	0.50	0.98	0.99	3.79	3.78
<i>mr</i>	—	0.50	0.99	0.98	3.78	3.79

to allow only anatomically identified connections. Also, inputs and desired outputs represented actual physiological signals, and all neurons had the real properties of saturation and cut-off. Learning resulted in a novel synaptic organization at the VN which reproduced the static, dynamic and nonlinear features of VOR and VN neurons. The close correspondence between model and observed responses suggests that the neural network organization discovered by the learning algorithm may correspond to the real mechanism underlying velocity storage in the VOR.

The recurrent back-propagation learning algorithm (Williams and Zipser 1989), like regular back-propagation (Rumelhart et al. 1986), changes weights according to the value of network error (4, 5 and 7). In the case of VOR network models, output unit error is related to retinal slip, a signal that is known to produce plastic changes in the operation of VOR (Wilson and Melvill Jones 1979). Thus the algorithm is physiological in its use of error driven learning. However, the recurrent back-propagation algorithm used here is nonlocal in that computation of the change for any one synaptic weight requires knowledge of all other weights and neural firing rates network-wide (6). In this respect the algorithm may not be physiological. The algorithm is computationally intensive, and therefore practical only for networks containing a small number of units and simple training sets of low temporal resolution. Despite these limitations, the simulations presented above successfully capture many of the essential features of VOR and VN neurons.

Lateral inhibitory commissural interconnections

The VOR neural networks learned to produce velocity storage by developing closed, reverberatory circuits (Hebb 1972) that could hold (store) the canal velocity signals over many synaptic cycles. The networks employed lateral inhibition (via vestibular commissures) to create positive feedback loops among model VN neurons that allowed them to partially integrate the oppositely directed canal signals, but not their SRs. Using

lateral inhibition to produce this type of neural integration was first proposed by Cannon et al. (1983). In the networks, units receiving a small net amount of lateral inhibition, which do not integrate the canal signals, can still have long time constants if they are coupled to other units that are leaky integrators. This mechanism of time constant lengthening is reminiscent of feedback through "neural low-pass filters" (Robinson 1981; Galiana and Outerbridge 1984).

To produce the desired outputs, the network had to learn not only to integrate, but also to pass the peaked, initial portion of the afferent signal. Separate fast and slow pathways developed. A parallel arrangement, with direct (fast) and leaky integrated (slow) pathways was first proposed by Raphan et al. (1979). More recently, contributions from both fast and slow pathways have been discerned in monkey VOR at high temporal resolution (Lisberger 1988). Thus networks learn to produce the VOR transformation by employing lateral inhibition and feedback through neural low-pass filters, and also by developing parallel fast and slow pathways. It is possible that such a complex combination of mechanisms is also utilized by the brainstem in realizing the VOR transformation.

Hidden unit and nonlinear properties

Compared to input units, all hidden units in VOR models trained to produce velocity storage had higher gains, lower SRs, longer time constants and asymmetric responses (Fig. 3, A and B; Tables 2 and 3). While high gain is due in part to push-pull innervation by canal inputs (Wilson and Melvill Jones 1979), the other properties all result from commissural inhibition. In the larger network, where processing was more distributed, hidden units were nonuniform in their properties (Table 3). These features are also observed for actual VN neurons in monkeys (Fuchs and Kimm 1975; Buettner et al. 1978; Waespe and Henn 1979). Hidden units could be distinguished by their initial responses (peaked or rounded). The transient responses of VN neurons have not been studied.

In the networks, the following magnitude-dependent nonlinearities are observed: output unit decay rate increased and phase lag relative to the inputs decreased as input magnitude increased (Fig. 4). Similar magnitude-dependent nonlinearities are observed for monkey VOR at rotational stimulus magnitudes greater than about 120 deg/s (Paige 1983; Fetter and Zee 1988). For sinusoidal inputs, phase lag decrement is accompanied by a skewing of the output waveform such that excursions from baseline are always steeper than returns to baseline, both for output units (Fig. 4B) and hidden units (Fig. 4D). A similar skewing has been noted for monkey VOR (Paige 1983) and VN neurons (Fuchs and Kimm 1975). The models suggest that these magnitude-dependent nonlinearities are due to hidden unit cut-off. Many real VN neurons are known to exhibit rectification in monkeys (Fuchs and Kimm 1975; Buettner et al. 1978).

In conclusion, the results suggest that velocity storage in the horizontal VOR is subserved by complex, inhibitory commissural interactions among VN neurons. The interactions include leaky integration (low-pass filtering) via lateral inhibition, and feedback through these low-pass filters, in parallel slow and fast pathways. Commissural inhibition lowers the SRs of VN neurons, making them prone to rectification. VN neuron cut-off disrupts velocity storage, producing nonlinearities in the responses of VN neurons and VOR. Thus the low SR of VN neurons, and the integrative and nonlinear properties of VN neurons and VOR, may be commonly produced by commissural inhibition.

Acknowledgements. Special thanks to Drs. D. Zipser and R. J. Williams for helpful conversations on the implementation of the learning algorithm. This work was supported by the Faculty Research and Innovation Fund of the University of Southern California.

References

- Anastasio TJ, Robinson DA (1989a) The distributed representation of vestibulo-oculomotor signals by brain-stem neurons. *Biol Cybern* 61:79–88
- Anastasio TJ, Robinson DA (1989b) Distributed parallel processing in the vestibulo-oculomotor system. *Neural Comp* 1:230–241
- Anastasio TJ, Robinson DA (1990) Distributed parallel processing in the vertical vestibulo-ocular reflex: Learning networks compared to tensor theory. *Biol Cybern* 63:161–167
- Blair SM, Gavin M (1981) Brainstem commissures and control of time constant of vestibular nystagmus. *Acta Otolaryngol* 91:1–8
- Buettner U, Waespe W (1981) Vestibular nerve activity in the alert monkey during vestibular and optokinetic nystagmus. *Exp Brain Res* 41:310–315
- Buettner UW, Buettner U, Henn V (1978) Transfer characteristics of neurons in vestibular nuclei of the alert monkey. *J Neurophysiol* 41:1614–1628
- Buettner UW, Henn V, Young LR (1981) Frequency response of the vestibulo-ocular reflex (VOR) in the monkey. *Aviat Space Environ Med* 52:73–77
- Cannon SC, Robinson DA, Shamma S (1983) A proposed neural network for the integrator of the oculomotor system. *Biol Cybern* 49:127–136
- Fernandez C, Goldberg JM (1971) Physiology of peripheral neurons innervating semicircular canals of the squirrel monkey. II Response to sinusoidal stimulation and dynamics of peripheral vestibular system. *J Neurophysiol* 34:661–675
- Fetter M, Zee DS (1988) Recovery from unilateral labyrinthectomy in rhesus monkey. *J Neurophysiol* 59:370–393
- Fuchs AF, Kimm J (1975) Unit activity in vestibular nucleus of the alert monkey during horizontal angular acceleration and eye movement. *J Neurophysiol* 38:1140–1161
- Galiana HL, Outerbridge JS (1984) A bilateral model for central neural pathways in vestibuloocular reflex. *J Neurophysiol* 51:210–241
- Hebb DO (1972) *Textbook of psychology*. Saunders, Philadelphia
- Lisberger SG (1988) The neural basis for learning of simple motor skills. *Science* 242:728–735
- McCrea RA, Strassman A, May E, Highstein SM (1987) Anatomical and physiological characteristics of vestibular neurons mediating the horizontal vestibulo-ocular reflex of the squirrel monkey. *J Comp Neurol* 264:547–570
- Milsum JH (1966) *Biological control systems analysis*. McGraw-Hill, New York
- Paige GD (1983) Vestibuloocular reflex and its interaction with visual following mechanisms in the squirrel monkey. I. Response characteristics in normal animals. *J Neurophysiol* 49:134–151
- Raphan Th, Matsuo V, Cohen B (1979) Velocity storage in the vestibulo-ocular reflex arc (VOR). *Exp Brain Res* 35:229–248
- Robinson DA (1975) Oculomotor control signals. In: Lennerstrand G, Bach-y-Rita P (eds) *Basic mechanisms of ocular motility and their clinical implications*. Pergamon Press, Oxford, pp 337–374
- Robinson DA (1981) The use of control systems analysis in the neurophysiology of eye movements. *Ann Rev Neurosci* 4:463–503
- Rumelhart DE, Hinton GE, Williams RJ (1986) Learning internal representations by error propagation. In: Rumelhart DE, McClelland JL, PDP Research Group (eds) *Parallel distributed processing: explorations in the microstructure of cognition, vol 1: Foundations*. MIT Press, Cambridge, pp 318–362
- Skavenski AA, Robinson DA (1973) Role of abducens neurons in vestibuloocular reflex. *J Neurophysiol* 36:724–738
- Waespe W, Henn V (1979) The velocity response of vestibular nucleus neurons during vestibular, visual, and combined angular acceleration. *Exp Brain Res* 37:337–347
- Williams RJ, Zipser D (1989) A learning algorithm for continually running fully recurrent neural networks. *Neural Comp* 1:270–280
- Wilson VJ, Melvill Jones G (1979) *Mammalian vestibular physiology*. Plenum Press, New York

Thomas J. Anastasio, Ph.D.
 USC Vestibular Lab
 PMB C-103
 1420 San Pablo Street
 Los Angeles, CA 90033
 USA