Biol. Cybern. 61, 79–88 (1989)

Biological Cybernetics

The Distributed Representation of Vestibulo-Oculomotor Signals by Brain-Stem Neurons

T. J. Anastasio and D. A. Robinson

Departments of Ophthalmology and Biomedical Engineering, The Johns Hopkins University, School of Medicine, Baltimore, MD 21205, USA

Abstract. The vestibuloocular reflex and other oculomotor functions are subserved by populations of neurons operating in parallel. This distributed aspect of the system's organization has been largely ignored in previous block diagram models. Neurons that transmit oculomotor signals, such as those in the vestibular nucleus (VN), actually combine the different types of signals in a diverse, seemingly random way that could not be predicted from a block diagram. We used the backpropagation learning algorithm to program distributed neural-network models of the vestibulo-oculomotor system. Networks were trained to combine vestibular, pursuit and saccadic eye velocity command signals. The model neurons in these neural networks have diverse combinations of vestibulo-oculomotor signals that are qualitatively similar to those reported for actual VN neurons in the monkey. This similarity implicates a learning mechanism as an organizing influence on the vestibulo-oculomotor system and demonstrates how VN neurons can encode vestibulo-oculomotor signals in a diverse, distributed manner.

Introduction

Modelling in the oculomotor system has largely utilized block diagrams in which the properties of neural populations are lumped into single black-box operators. These models approximate overall oculomotor behavior but indicate little about how signals would be processed and carried by real networks of neurons. Recordings from single neurons indicate that oculomotor signals are combined and transmitted in a diverse manner. Each neuron carries components of signals with apparently randomly chosen weights so that each total signal is distributed over the population. These findings would be impossible to predict from block diagrams and indicate a gap between black-box modelling and actual neural behavior. This study is a preliminary attempt to bridge that gap.

Some pools of neurons do carry a single message. Cells in the reticular formation burst during ipsilateral saccades and are otherwise silent (ls, rs, Fig. 1). If E is

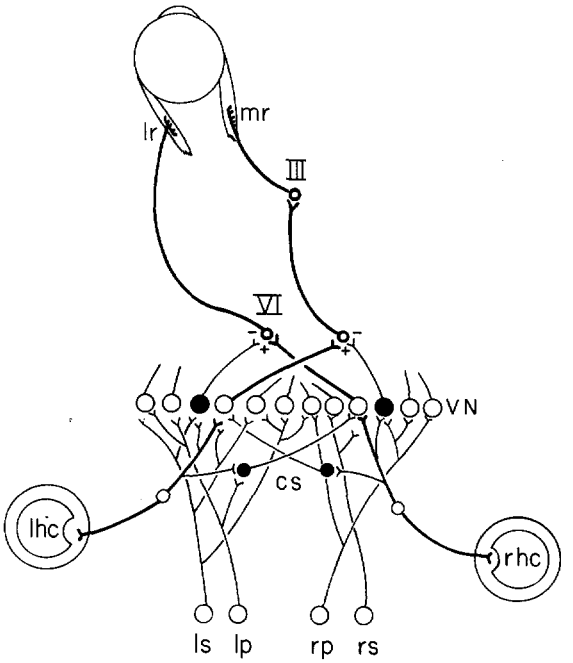

Fig. 1. A schematic of the vestibuloocular reflex (VOR) originating in the left (l) and right (r) horizontal (hc) semicircular canals and ending on the motoneurons of the left lateral and medial recti (lr, mr) via the cranial nuclei III, VI. The three-neuron arc of the horizontal VOR from rhc to lr is indicated by heavy lines. The projection to mr is a four-neuron arc, an exception to the general rule. A pool of neurons in the vestibular nucleus (VN) receive eye-velocity commands from the canals, from left and right smooth pursuit neurons (lp, rp) and saccadic burst neurons (ls, rs) where they are intermixed and passed to the motoneurons. A few inhibitory neurons are represented by filled cells to illustrate reciprocal innervation. Two of them illustrate a feedforward commissural system (cs)

eye position and $\dot{E}_r$ eye velocity during a rapid movement (saccade, quick phase), their discharge rate is roughly described by $r\dot{E}_r(t)$, r varying from cell to cell (van Gisbergen et al. 1981). Similarly, vestibular fibers from the semicircular canals carry the signal $(R_0 + v\dot{E}_v(t))$ where R_0 is the spontaneous discharge rate (with the head still) and $v\dot{E}_v$ is proportional to the vestibular eye-velocity command (Miles and Braitman 1980). R_0 and v vary over the population. Finally, cells in the flocculus carry a pursuit eye-velocity command $(R_0 + p\dot{E}_p(t))$ where $\dot{E}_p$ is eye velocity during pursuit and R_0 and p vary from cell to cell (Miles et al. 1980). $\dot{E}_r$, $\dot{E}_p$, and $\dot{E}_v$, are distributed here but only by variations in the scalars r, p, and v; an analytically trivial distribution.

When, however, one records from neurons in the region where these signals are combined and integrated ($\int$ dt) to be passed on to the motoneurons (Fig. 1), one finds a diversity of signal combinations indicating a truly distributed representation. The discharge rate, R, of most cells can be described by

$$R = R_0 + kE + p\dot{E}_p + v\dot{E}_v + r\dot{E}_r \tag{1}$$

(Tomlinson and Robinson 1984) where kE is proportional to the eye position command produced by neural integration of the eye velocity signals. All combinations of k, p, v, and r can be found among premotor neurons. Yet when these signals converge onto motoneurons, the signal

$$R_m = R_0 + kE + \varrho\dot{E} \tag{2}$$

emerges, where R_m is the discharge rate of the motoneuron and $\varrho\dot{E}$ is proportional to eye velocity regardless of the source of command (Robinson 1970). Thus, the velocity commands spread out over the premotor population and then converge on the motoneurons to produce the correct final output. The seemingly random distribution of the signal components suggests that the network began with a partially randomized set of synaptic weights and then organized itself into whatever pattern of connectivity got the job done based on trial and error.

Meanwhile, theorists have been exploring various learning networks such as backpropagation schemes. Learning is driven by the error between the current and desired outputs which is used to adjust the synaptic weights between the layers until, after many iterations, the errors for all the inputs are acceptably small. Initially, the weights are randomized. This scheme is remarkable at learning complex behaviors (Rumelhart et al. 1986; Sejnowski and Rosenberg 1987; Zipser and Andersen 1988) and its units always show distributed contributions to the solutions. This suggested to us that such learning networks could represent the distributed aspect of the oculomotor system. Because the analysis of learning networks with dynamic elements and feedback loops, such as would be required to simulate the neural integrator, is still very difficult, this preliminary study only concerns the nature of the velocity signals in (1).

Methods

Network Architecture

The vestibulo-oculomotor system (Fig. 1) was modelled as a three-layer network (Fig. 2C). The input layer represents neurons carrying eye-velocity commands from canal primary afferent, pursuit and burst (saccade) neurons. Output units represent motoneurons of the muscles of the left eye. Units in the hidden layer correspond to neurons in the caudal pontine nuclei such as the vestibular nuclei and nuclei prepositus hypoglossi (all abbreviated as VN in Fig. 1) that combine the velocity commands and send them to the motoneurons. Every input unit projected to every hidden unit and every hidden unit projected to every output unit. There were no intra-layer or feedback connections. The number of units in each layer differed and are described in each section.

The operation of the network follows the description of Rumelhart et al. (1986). The activity of each model neuron is a sigmoidal function of the weighted sum of its inputs (Fig. 2A, B):

$$s_i = S(E_i) = \frac{1}{1 + e^{-E_i}}, \tag{3}$$

where s_i is the output or firing rate of the ith unit and the weighted input is:

$$E_i = \sum_j w_{ij} s_j, \tag{4}$$

where w_{ij} is the synaptic strength from the jth unit, firing at s_j, to the ith unit. The weights can be positive or negative representing excitation or inhibition. Like neurons, the output is roughly linear over a midrange and can be driven to cut-off or saturation with large negative or positive inputs (Fig. 2B).

Learning Algorithm

The symbol $s_i^{(n)}$ denotes the firing rate of the ith unit in the nth layer. The output is the Nth layer. Learning begins with randomization of all the synaptic weights; their values had a uniform distribution between ± 1.0. Training proceeds according to an input/output table that specifies the desired output pattern for each input pattern. The desired input s_i for each pattern is forward-propagated through the network, layer by layer; the s_i of all the hidden and then output units are computed for this input by successive applications of (3) and (4).

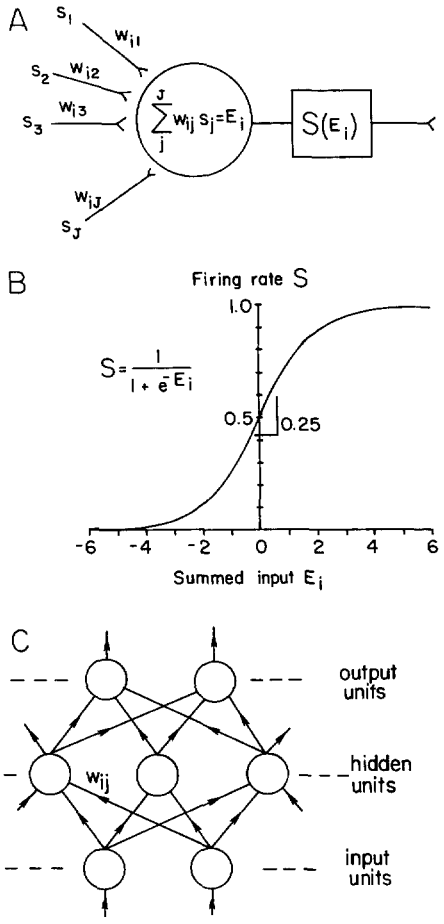

Fig. 2A–C. Properties of the units and networks. **A** The output of each processing unit is a function of the weighted sum of all its inputs. **B** This function is sigmoidal. A zero input causes a spontaneous rate of 0.5; moderate modulation would be over a quasilinear midrange where the slope is about 0.25. Large inputs can drive the unit into cutoff or saturation. **C** General network architecture. Normally each unit projects to all units in the subsequent layer

The goal of the learning procedure is to minimize the average, squared error between the computed values of the output and the desired output s_i^*:

$$\text{Error} = \sum_i (s_i^* - s_i^{(N)})^2 . \tag{5}$$

The backpropagation algorithm does this by adjusting the synaptic weights via a gradient descent that approaches a set of weights yielding the minimum error. The error gradient on the output layer is:

$$\delta_i^{(N)} = (s_i^* - s_i^{(N)}) S'(E_i^{(N)}), \tag{6}$$

where $S'(E_i)$ is the first derivative of the function $S(E_i)$ in (3). The error is then back-propagated through the network, layer by layer, by computing the error gradients on previous layers:

$$\delta_i^{(n)} = \sum_j \delta_j^{(n+1)} w_{ji}^{(n)} S'(E_i^{(n)}), \tag{7}$$

where $w_{ji}^{(n)}$ is a weight from layer n to $n+1$. These gradients determine the changes in each weight to reduce the error for a given input. Since there are several input patterns, the learning of one must not destroy the learning of others. To do this, a running average is formed of the previous weight changes with the current weight change:

$$\Delta w_{ij}^{(n)}(u) = \alpha \Delta w_{ij}^{(n)}(u-1) + (1-\alpha)\delta_i^{(n+1)} s_j^{(n)}, \tag{8}$$

where α is a smoothing parameter (we used 0.9) and u is a running index of the sequence of input patterns. The weights are then updated by:

$$w_{ij}^{(n)}(t+1) = w_{ij}^{(n)}(t) + \varepsilon \Delta w_{ij}^{(n)}, \tag{9}$$

where t is the index of the weight updates and ε controls learning rate (we used a value of 10).

Following forward-propagation of the current input, the output-unit gradients were calculated using (6) and then used to calculate adjustments for the hidden-to-output weights using (8) and (9). Next the hidden-unit gradients were calculated using (7) incorporating the newly adjusted hidden-to-output weights. Finally, the hidden-unit gradients were used to calculate adjustments for the input-to-hidden weights using (8) and (9). The algorithm cycled through the input/output table repeatedly until the difference between the actual and desired values of the outputs was less than a tolerance T (we used 0.01).

For most simulations, the hidden-to-output weights were fixed at specific values. In those cases, the hidden-unit gradient was calculated using the output unit gradient without ever changing the hidden-to-output weights. Then the input-to-hidden weights were modified as before.

Results

Simplest Model of the Horizontal Vestibuloocular Reflex

The three-neuron arc of the vestibuloocular reflex (VOR) is the backbone of the networks to be described (Fig. 1). To examine the most elementary aspects of network behavior, we began with the simplest arrangement – a network with only two neurons in each layer. The input units represent afferents from the left and right horizontal semicircular canals (lhc and rhc, Fig. 3) and the output units represent the motoneurons of the left lateral and medial recti (lr and mr, Fig. 3). The input/output requirements are given in Table 1A. Each input and output unit has a spontaneous rate (SR) of 0.5. The input/output table specifies a compensatory VOR; head rotation left increases the rate of the lhc and decreases rhc by the same amount. The compensatory eye movement to the right would then

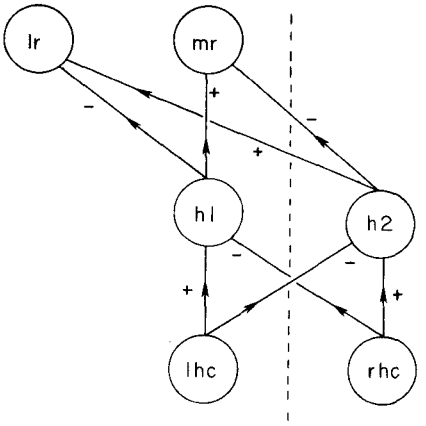

Fig. 3. The simplest model of the horizontal VOR with only two hidden units, h1 and h2. Other abbreviations as in Fig. 1. All weights were modifiable. The result of training is shown in Table 1; the system discovers Sherrington's law of reciprocal innervation and utilizes a feedforward version of the inhibitory vestibular commissural system indicated by the signs shown

occur with an equal increase in the rate of mr and a decrease in lr, as specified.

In a typical example, a network, as in Fig. 3, finished training after about 200 passes through the input/output table. Final weights are in Table 1B. The result is similar to the basic VOR circuitry (Fig. 1). Each hidden unit excites one motoneuron and inhibits the other, thus discovering Sherrington's law of reciprocal innervation. In dozens of trials we have never seen this fail to occur and have indicated this by putting signs on the weights in Fig. 3. (Note that when any unit both excites and inhibits target units, an interneuron of fixed weight is implied to satisfy Dale's law.) This push-pull organization also applies to the canals; a hidden unit excited by one canal is always inhibited by the other. This arrangement is similar to the commissural connections actually present in the VOR. A nontrivial difference is that it is generally believed that second-order vestibular neurons inhibit each other across the midline via type II interneurons (Shimazu and Precht 1966), not shown in Fig. 1. Such an arrangement would involve projections between hidden units creating feedback loops. In this exploratory study we did not attempt to incorporate feedback connections but settle, for now, for a feedforward inhibitory commissural system. Nevertheless, the learning network does automatically adopt some form of commissural inhibition.

The responses of the hidden and output units for this run appear in Table 1C. The outputs match those desired (Table 1A) within the tolerance of 0.01. The SRs and gains of the hidden and output units appear in Table 1D. Ipsilateral and contralateral refers to the *left* side in this and all other examples. The SR of h1 (0.55) is greater than 0.50 because its excitation (1.63 from lhc) is greater than its inhibition (1.21 from rhc), and the SR of h2 (0.45) is less than 0.50 because its inhibition (2.64 from lhc) is greater than its excitation (2.25 from rhc). The learning algorithm ensures that the output SRs end up near 0.5; although the hidden-to-output weights are greater from h2 than h1, the SR of h2 is lower so that, with the head still, their influences cancel.

Gains are taken as $(s_i - \mathrm{SR}_i)/(s_j - \mathrm{SR}_j)$ where the numerator is the change in s_i from the SR of a hidden or output unit, and the denominator is the change in an ipsilateral (left) input unit. During ipsilateral head rotation, the firing rate of h1 increases due to excitation from the lhc and disinhibition from the rhc with a net ipsilateral vestibular gain, iV, of 0.69, while the rate of h2 decreases due to inhibition from the lhc and disfacilitation from the rhc with an iV of −1.16. Similarly, the gains of the lr and mr are −0.89 and 0.95 respectively. To test linearity, gains were calculated for

Table 1. Backpropagation uses reciprocal innervation in solving the horizontal VOR problem. lhc and rhc, left and right horizontal canals; lr and mr, left lateral and medial recti; h1, h2, hidden units 1 and 2; SR, spontaneous discharge rate; iV, cV, vestibular gains for ipsilateral and contralateral head movements

A. Input/output table

	Input		Output	
	lhc	rhc	lr	mr
head still	0.50	0.50	0.50	0.50
head left	0.60	0.40	0.40	0.60
head right	0.40	0.60	0.60	0.40

B. Final weights

to		h1	h2	to		lr	mr
from	lhc	1.63	−2.64	from	h1	−1.71	1.84
	rhc	−1.21	2.25		h2	2.08	−2.23

C. Hidden and output unit responses

Input		Hidden		Output	
lhc	rhc	h1	h2	lr	mr
0.50	0.50	0.552	0.450	0.498	0.502
0.60	0.40	0.621	0.334	0.409	0.597
0.40	0.60	0.481	0.572	0.590	0.403

D. Spontaneous rates and gains of hidden and output units

	Hidden		Output	
	h1	h2	lr	mr
SR	0.55	0.45	0.50	0.50
iV	0.69	−1.16	−0.89	0.95
cV	0.71	−1.22	−0.92	0.99

contralateral head rotations (*cV*, Table 1D). The small differences are due to h1 and h2 working around SRs not exactly in the middle of the range. Such differences will be ignored in subsequent examples. Other runs with different initializations will come to different solutions – but they will only differ quantitatively – the general ideas remain as illustrated.

Models of the VOR With Many Hidden Units

Since there are many interneurons in the VOR, one must examine models with more hidden units. Fully-connected, fully-adaptable models with six hidden units were trained on the input/output regimen of Table 1A. Most hidden units were innervated in push-pull by the canals and reciprocally innervated the motoneurons but some hidden units had anomalous connections. The properties of a typical network appear in Table 2. Hidden units h2 and h3 are miswired (*). Units h4 and h5 are excited by the lhc, inhibited by the rhc, inhibit the lr and excite the mr, as expected. Similarly, units h1 and h6 are inhibited by the lhc, excited by the rhc, excite the lr and inhibit the mr, also as expected. In contrast, units h2 and h3 receive same-sign connections from both canals and do not reciprocally innervate the motoneurons. Note, however, that they have low gains (V) so their non-reciprocal connections will not adversely affect the overall response.

This example illustrates a general trend: in VOR network models with expanded hidden layers, some units will be miswired and have low gains. This occurs because simple functions such as the VOR require only a few hidden units to represent them. After the initial randomization, the learning algorithm produces a few appropriately connected, high gain hidden units that support the function after which the other (often miswired) low gain hidden units remain essentially unused. Evidently, the fully-connected, fully-adaptable network has more modifiable weights than it needs and, for our purposes, it is desirable to reduce them in order to deal with more than just three or four hidden units. Moreover, it is unrealistic to assume that the VOR develops from an entirely random network. Because the evidence to date suggests synaptic plasticity at premotor stages rather than at the motoneurons, we fixed the hidden-to-output weights. These fixed connections had a uniform, specified value and were arranged to reciprocally innervate the outputs. An example is shown in Table 3A. Note that this polarizes the hidden units into ipsi- and contralateral. In Fig. 3,

Table 2. Miswirings can occur in horizontal VOR network models incorporating more than two hidden units. Nomenclature as in Table 1. Asterisk * indicates units that are miswired. V, vestibular gain

Synaptic weights		h1	h2	h3	h4	h5	h6
from inputs	lhc	−1.89	−0.66	−0.27	1.79	0.06	−1.18
	rhc	1.79	−0.53	−0.02	−1.82	−1.02	0.79
to outputs	lr	1.53	−0.61	0.49	−1.91	−0.52	0.87
	mr	−1.81	−0.55	0.39	2.07	0.35	−0.58
miswirings			*	*			
Spontaneous rates							
and gains	SR	0.49	0.36	0.46	0.50	0.38	0.45
	V	−0.90	−0.03	−0.06	0.89	0.26	−0.48

Table 3. Hidden unit gain and learning rate depend upon output strength. Nomenclature as in Table 2. how, hidden-to-output weights; cy, number of cycles through the input/output table to learn task

A. Fixed hidden-to-output connectivity pattern

	from	h1	h2	h3	h4	h5	h6
to	lr	−1.00	−1.00	−1.00	1.00	1.00	1.00
	mr	1.00	1.00	1.00	−1.00	−1.00	−1.00

B. Vestibular gains V *of hidden and output units*

how	h1	h2	h3	h4	h5	h6	lr	mr	cy
10.00	−0.01	0.08	0.07	−0.05	−0.09	−0.12	−0.97	0.97	192
5.00	0.16	0.40	0.14	−0.07	0.07	−0.16	−1.10	1.10	16
2.00	0.51	0.30	0.45	−0.32	−0.06	−0.21	−0.91	0.90	32
1.00	0.66	0.52	0.61	−0.33	−0.58	−0.94	−0.90	0.90	125
0.50	1.08	1.29	1.34	−1.43	−1.05	−1.08	−0.90	0.90	545

with the signs as shown, h1 could have been on the right as easily as on the left and vice versa for h2. After fixing the hidden-to-output weights, activation of ipsilateral (left) hidden units always drove the eye right and vice versa. These networks still had adaptable input-to-hidden weights.

The properties of these networks depended on the values of the fixed hidden-to-output weights. Several models with six hidden units, but different hidden-to-output weights, were trained according to Table 1A. Gains of the hidden and output units for typical trials appear in Table 3B, along with the hidden-to-output weights and the number of cycles, or passes through the input/output table, needed to train each network to a T of 0.01.

Due to the fixed pattern in Table 3A, the gains of h1 to h3 should be positive, those of h4 to h6 negative, to be connected appropriately. The two networks with the highest hidden-to-output weights each had one wrong gain which was, as usual, small. No inappropriate gains occurred for the other cases. The number of cycles had a minimum in the region of hidden-to-output weights around 2.0 to 5.0. Thus, fixing the hidden-to-output weights produces fewer hidden units with inappropriate connections. Such weights were used in all subsequent examples. The choice of these weights was varied in a trade-off between adequate hidden-unit utilization and learning speed.

Pursuit-Vestibular-Saccadic Compound Representations

Pursuit and saccadic eye-velocity commands also funnel through the VN and we wanted to see how they would be represented on the hidden units in a combined oculomotor-vestibular network. We began with just pursuit and vestibular signals in networks with six hidden units. These models had four inputs; one pair represented the horizontal canals as before, the other represented pursuit velocity commands from left and right pursuit neurons (lp and rp, Fig. 1) located on the left and right sides of the brain respectively. The latter also operated in push-pull and spanned the same range as the vestibular inputs but required responses in the same rather than the opposite direction; activation of lp and suppression of rp should produce a leftward eye movement whereas activation of lhc and suppression of rhc produced a rightward movement. The input-/output patterns are given in Table 4A. The models had fixed hidden-to-output connections as in Table 3A but with weights of 2.0. These networks learned the pursuit-vestibular input/output patterns (Table 4A) after about 50 cycles.

The SRs and gains of the hidden and output units of a typical, mature network appear in Table 4B. Gains are calculated as before with the input taken to be the left input (lp or lhc). Consequently, for an appropriately-connected left hidden unit, the vestibular gain, V, will be positive, the pursuit gain, P, negative and vice versa on the right. Thus, the pursuit and vestibular gains of a hidden (or output) unit should normally have opposite signs. Table 4B mainly represents appropriately-connected hidden units (only h5 shows an inappropriate, but very small, pursuit gain). These data show that pursuit and vestibular gains varied greatly not only between hidden units but also within individual units. Some carried one eye-velocity signal almost exclusively. For example, h1 had a negligible vestibular gain and is essentially a pure-pursuit neuron, while h5 had a negligible pursuit gain and would appear as a pure-vestibular neuron. The other hidden units carried both signals to a significant extent but with widely differing gains. Thus, members

Table 4. Pursuit and vestibular compound representations. Nomenclature as in previous tables. lp, rp, left and right pursuit velocity commands. P, pursuit gain

A. Input/output table

Inputs				Outputs	
lp	lhc	rhc	rp	lr	mr
0.50	0.50	0.50	0.50	0.50	0.50
0.60	0.50	0.50	0.40	0.60	0.40
0.40	0.50	0.50	0.60	0.40	0.60
0.50	0.60	0.40	0.50	0.40	0.60
0.50	0.40	0.60	0.50	0.60	0.40

B. Gains and spontaneous rates of hidden and output units

	h1	h2	h3	h4	h5	h6	lr	mr
SR	0.45	0.41	0.57	0.34	0.51	0.58	0.50	0.50
P	−0.29	−0.36	−0.27	0.33	−0.02	0.62	0.91	−0.91
V	0.08	0.51	0.41	−0.17	−0.40	−0.26	−0.91	0.91

of such a distributed system can participate quite differently in subserving different functions.

To explore this further the hidden layer was expanded to 40 units. The fixed hidden-to-output weights were 0.35; ipsilateral h1–h20 inhibited the lr and excited the mr, the reverse for contralateral h21–h40. This model was trained with the same input/output patterns in Table 4A and learned after about 200 cycles. The properties of the hidden units were evaluated, and scatter plots of the vestibular versus the pursuit gains for all 40 hidden units were made. Figure 4A is a typical example. Since V and P should normally have opposite signs, units in the upper left and lower right quadrants are correctly wired.

Figure 4A shows the striking variety of combinations of hidden-unit pursuit and vestibular gains. Units falling near the vertical axis would appear as pure-vestibular neurons; units near the horizontal axis as pure-pursuit. The two output units (+, Fig. 4A), with gains of ± 0.90, lie on the 45 deg line. Hidden units along this line have equal pursuit and vestibular gains. Most units have significant but unequal pursuit and vestibular gains. Some units even fall into quadrants one and three. These apparently anomalous units fire faster for a vestibular eye movement in one direction and a pursuit movement in the opposite. Such pursuit-vestibular opposite neurons have been reported in the vestibular nuclei (Fuchs and Kimm 1975; Lisberger and Miles 1980; Chubb et al. 1984).

Rather than forming discrete subpopulations, hidden units seemed to fall along a continuum with regard to pursuit and vestibular gains. When the weights were randomized and the network retrained, the distribution of points in the scatter plots could be very different from one mature network to another. No general patterns emerged even from an ensemble of such plots. This emphasizes an important point: there is no unique solution in such learning networks and hidden units can participate quite differently from one network to another.

We added saccades to this model which now employed six input units. The additional input pair represented left and right burst neurons (ls and rs, Fig. 1) that carry saccadic eye-velocity commands, ls being ipsilateral, rs contralateral. These neurons were either on or off; during a left saccade, for example, the output of ls was 1.0, that of rs 0.0. Otherwise, the units were silent (output 0.0). At this elementary level, we made no attempt to simulate saccades of various sizes; all inputs commanded a saccade of fixed velocity (left or right). Besides contacting all 40 hidden units, the burst neurons also made direct, reciprocal connections with the motoneurons to reflect known anatomical pathways. These input-to-output weights were fixed at 2.5.

Networks were trained according to the input/output patterns in Table 5A; training required about 1,000 cycles. The SRs and gains of the mature networks were evaluated. Data were further reduced for illustration by combining the pursuit and vestibular gains of each unit into a single number by $(|P| - |V|)/(|P| + |V|)$, where $|P|$ and $|V|$ are the absolute values of P and V. This characterizes units by the relative strengths of pursuit and vestibular gains. It is -1.0 for pure-

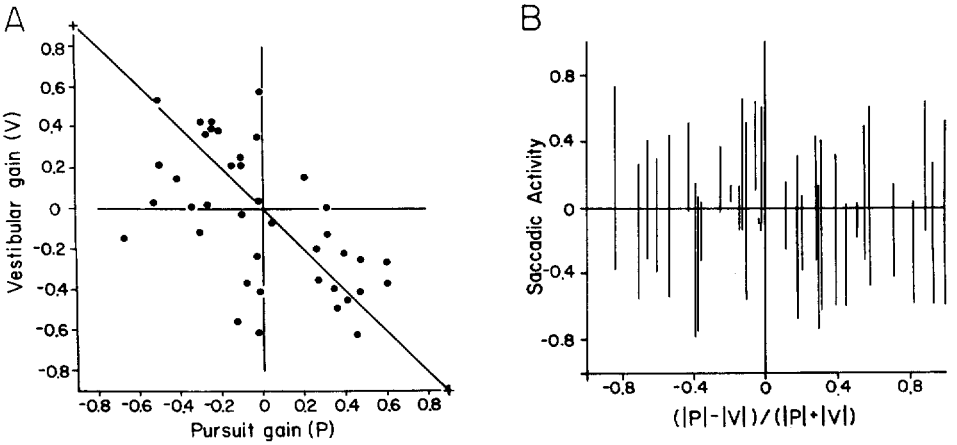

Fig. 4. **A** A typical scattergram of a network with forty hidden units and pursuit and vestibular inputs. Gains P and V are changes in firing rate for the hidden units (filled circles) or output units (+) divided by the change in the left unit of an input pair for pursuit or vestibular commands respectively. Appropriately connected units have pursuit and vestibular gains of opposite signs. Note that the pursuit and vestibular gains of each unit can be very different. **B** Behavior of a similar network with saccadic commands added. $|P|$ and $|V|$ are the absolute values of pursuit and vestibular gains so the abscissa is a measure of the relative role a unit plays in each activity; pure vestibular units on the left at -1.0, pure pursuit on the right at $+1.0$, equal absolute gains, which includes the output units, in the center. Each unit is represented by a vertical line indicating activity during saccades to the left and right. Positive activity is a burst, negative a pause. Many units burst in one direction and paused in the other but other units had more complicated behaviors. Again, the units show a great diversity

Table 5. Pursuit, vestibular, and saccadic compound representation. Nomenclature as in previous tables. ls, rs, left and right saccadic inputs. iSA, cSA, ipsilateral and contralateral saccadic activity, burst are positive, pauses are negative

A. Input/output table

Input						Output	
lp	lhc	ls	rs	rhc	rp	lr	mr
0.50	0.50	0.00	0.00	0.50	0.50	0.50	0.50
0.60	0.50	0.00	0.00	0.50	0.40	0.60	0.40
0.40	0.50	0.00	0.00	0.50	0.60	0.40	0.60
0.50	0.60	0.00	0.00	0.40	0.50	0.40	0.60
0.50	0.40	0.00	0.00	0.60	0.50	0.60	0.40
0.50	0.50	1.00	0.00	0.50	0.50	1.00	0.00
0.50	0.50	0.00	1.00	0.50	0.50	0.00	1.00

B. Properties of selected units from Fig. 4B

	VB	PB	PVP	PV	PVOB	lr	mr
SR	0.46	0.52	0.58	0.47	0.47	0.50	0.50
P	−0.06	−0.36	0.29	0.30	−0.51	0.91	−0.91
V	0.75	0.00	−0.21	−0.64	0.58	−0.92	0.92
iSA	−0.37	−0.58	0.00	0.02	0.11	0.98	−0.98
cSA	0.51	0.54	−0.69	0.00	0.65	−0.98	0.98

vestibular neurons, +1.0 for pure-pursuit neurons and 0.0 for equal absolute gains. The properties of a typical network are shown in Fig. 4B. Each unit is represented by a vertical line located horizontally by the above ratio. Because of the extreme values of the burst, the hidden-unit responses were often nonlinear so they were analyzed separately for saccades in each direction. Saccadic activity, SA, was taken as $(s_i - \mathrm{SR}_i)/(0.5)$. A positive response is a burst, negative a pause. Because most units had a SR near 0.5, the rate change was normalized by 0.5 so that a full burst or pause had a range of ± 1.0. The end points of the vertical lines in Fig. 4B show this saccadic activity.

The long line down the center of Fig. 4B represents the output units where $|P|$ and $|V|$ are practically equal (Table 5B); their saccadic activities are close to ± 1.0 as demanded by training. The other 40 lines for the hidden units are distributed fairly evenly along the horizontal axis, indicating an equal probability with regard to their relative pursuit and vestibular contributions. Fewer than half the lines extend well above and below zero. These units burst for a saccade in one direction and paused in the other and would correspond to so-called burst-tonic cells. Many cells burst but did not pause, or paused but did not burst. Some burst or paused in both directions. Such patterns are not uncommon experimentally (Fuchs and Kimm 1975; Lisberger and Miles 1980; Chubb et al. 1984; Tomlinson and Robinson 1984).

The properties of some selected hidden units are detailed in Table 5B. SA has been separated into ipsilateral (iSA) and contralateral (cSA); bursts are still positive and pauses negative. The first two hidden units burst for right saccades (cSA; 0.51, 0.54) and paused for left (iSA; −0.37, −0.58) with roughly equal intensity. The first (VB) had a small pursuit gain (−0.06) but a good vestibular gain (0.75) and could be called a vestibular-burst unit. The other (PB) had no vestibular gain and would be a pursuit-burst unit. The last three hidden units represented saccades differently. The pursuit-vestibular-pause (PVP) neuron was inhibited for pursuit and vestibular movements to the right, paused for saccades to the right (−0.69) but did nothing for saccades to the left. The pursuit-vestibular (PV) neuron did nothing for saccades. Finally, the pursuit-vestibular-omni-burst (PVOB) neuron burst (0.11, 0.65) for saccades in both directions. These units were chosen to emphasize the variability in such a distributed model because it is similar to what can be found in the caudal pons and VN (ibid.).

Discussion

We used a backpropagating learning network because we knew a priori that its hidden units would carry distributed oculomotor commands. A major criticism of this scheme is that it is unphysiological. Fortunately, we avoided this by fixing the hidden-to-output weights so that learning was restricted to the input-to-hidden connections. For the VOR and pursuit, the output error was the difference between actual and desired eye velocity which is the velocity with which visual images slip across the retina. This signal abounds in the nervous system and its appearance in the accessory

optic system and flocculus are important for motor learning in the VOR (e.g. Ito 1982). Most of our units worked over their middle range for the VOR and pursuit, so their $S'(E)$ values in (6) and (7) were nearly constant at 0.25 (Fig. 2B). Thus δ_i^3 is proportional to retinal slip and, since the hidden-to-output weights, w_{ji}^2, were all fixed, so was δ_i^2. Consequently, the delta rule (8) simply indicates that the change in weights, Δw_{ij}^1, be proportional to retinal slip (δ_i^2) and current firing rate s_j^1. This is not unphysiological. Thus, our reduction in network architecture also removed the unphysiological features of backpropagation and the role of the external teacher could be played by retinal slip. These considerations do not apply to saccades. While saccadic plasticity is well documented, its error involves eye position, a signal not present in our models. Thus, our saccadic training scheme is unphysiological.

One should recognize that a close correspondence between the VOR and a three-layered network is not necessary. Relay neurons such as internuclear neurons (Pola and Robinson 1978) could have been added to the model. Many third- and higher-order neurons in the pons help carry eye-velocity commands to the motoneurons suggesting more layers. More inputs could have represented the canal nerve population and more motoneurons could have been added. Not all input units need contact all hidden units. Like bursts inputs, pursuit inputs could also have projected directly to motoneurons. All these variations are equivalent to the architecture used. None of these changes would have altered the basic result that the three eye-velocity commands become distributed, seemingly haphazardly, over the hidden units.

Any of these changes would, however, change the distribution of signals. For example, studies in the VN of monkeys show a higher percentage of pure vestibular cells than in our model (15%, Fuchs and Kimm 1975; 15%, Tomlinson and Robinson 1984; 50%, Chubb et al. 1984; 10%, our model). This was adjusted in some models (not presented) by allowing the pursuit inputs to project direcrly to the motoneurons with a fixed weight. This, of course, decreased the average pursuit gain in the hidden layer and increased the number of pure-vestibular units. In other simulations (not presented) we deleted some connections from pursuit inputs to hidden units which also, of course, increased the number of pure-vestibular units. We also deleted vestibular and saccadic inputs to arbitrarily selected hidden units. While all these signals were still distributed over the hidden layer, these deletions caused pockets in which one of the input signals was absent.

Our results suggest that error-driven learning shapes the vestibulo-oculomotor network. Presumably, VOR pathways are formed during embryonic development, but prenatal synaptic strengths are unlikely to be correct. The gain of the VOR in many newborn animals is abnormal (e.g. Collewijn 1977). Postnatal experience, especially retinal slip, is needed to adjust these weights. Evolution probably establishes a basic network topology but genes can only predetermine statistical order. The initially randomized synaptic connections must then be modified by visual experience, which proceeds to sculpture the network into one of the many possible solutions that works.

A serious shortcoming in the backpropagation model we used is that all the connections can project only forward. The vestibular commissural system allows second-order vestibular cells to inhibit their contralateral fellows via an inhibitory interneuron (Shimazu and Precht 1966). This is equivalent to allowing hidden units to form lateral connections and feedback loops. The mathematics for dealing with this is still rudimentary and beyond the scope of this study. This is unfortunate, however, because a strictly feedforward model can only create zero-memory input/output functions without interesting temporal dynamics. Actual neural networks can differentiate and integrate signals (with respect to time) and generate patterns such as saccadic pulses. Without developing such abilities, learning networks will be of limited use in studying motor control.

In summary, this model did well in simulating the distributed representation of vestibulo-oculomotor signals by VN and related neurons. It quickly discovered Sherrington's law of reciprocal innervation and invented a feed-forward vestibular, inhibitory commissural system. Just as is found experimentally, the hidden units form a heterogeneous pool of units, within which signal parameters vary essentially as a continuum, which nevertheless can bring order out of seeming chaos to produce organized eye movements. These results suggest a learning-network mechanism for the variability found in the neural organization of the vestibulo-oculomotor system.

Acknowledgements. This work was supported by a NRSA Fellowship, EY05901, to T.J.A., Grant EY00598 to D.A.R. and Grant EY01765 for computer support, all from the National Eye Institute of the National Institutes of Health, Bethesda, Maryland, USA. We thank A. L. McCracken for preparing the manuscript and C. V. Bridges for the illustrations. We also thank Dr. T. J. Sejnowski for advice and encouragement.

References

Chubb MC, Fuchs AF, Scudder CA (1984) Neuron activity in monkey vestibular nuclei during vertical vestibular stimulation and eye movements. J Neurophysiol 52:724–742

Collewijn H (1977) Optokinetic and vestibulo-ocular reflexes in dark-reared rabbits. Exp Brain Res 27:287–300

Fuchs AF, Kimm J (1975) Unit activity in the vestibular nucleus of the alert monkey during horizontal angular acceleration and eye movement. J Neurophysiol 38:1140–1161

Gisbergen JAM van, Robinson DA, Gielen S (1981) A quantitative analysis of generation of saccadic eye movements by burst neurons. J Neurophysiol 45:417–442

Ito M (1982) Cerebellar control of the vestibulo-ocular reflex – around the flocculus hypothesis. Ann Rev Neurosci 5:275–296

Lisberger SG, Miles FA (1980) Role of primate medial vestibular nucleus in long-term adaptive plasticity of vestibuloocular reflex. J Neurophysiol 43:1725–1745

Miles FA, Braitman DJ (1980) Long-term adaptive changes in primate vestibuloocular reflex. II. Electrophysiological observations on semicircular canal primary afferents. J Neurophysiol 43:1426–1436

Miles FA, Fuller JH, Braitman DJ, Dow BM (1980) Long-term adaptive changes in primate vestibuloocular reflex. III. Electrophysiological observations in flocculus of normal monkeys. J Neurophysiol 43:1437–1476

Pola J, Robinson DA (1978) Oculomotor signals in the medial longitudinal fasciculus of the monkey. J Neurophysiol 41:245–259

Robinson DA (1970) Oculomotor unit behavior in the monkey. J Neurophysiol 33:393–404

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

Sejnowski TJ, Rosenberg CR (1987) Parallel networks that learn to pronounce english text. Compl Syst 1:145–168

Shimazu H, Precht W (1966) Inhibition of central vestibular neurons from the contralateral labyrinth and its mediating pathway. J Neurophysiol 29:467–492

Tomlinson RD, Robinson DA (1984) Signals in vestibular nucleus mediating vertical eye movements in the monkey. J Neurophysiol 51:1121–1136

Zipser D, Anderson RA (1988) A back propagation programmed network that simulates response properties of a subset of posterior parietal neurons. Nature 33:679–684

Received: November 4, 1988

Accepted: January 16, 1989

Thomas J. Anastasio, Ph.D.
USC Vestibular Laboratory
Parkview Medical Building, C-103
1420 San Pablo Street
Los Angeles, CA 90033
USA