Distributed Parallel Processing in the Vestibulo-Oculomotor System

Thomas J. Anastasio
Vestibular Laboratory, University of Southern California,
Los Angeles, CA 90033, USA

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

The mechanisms of eye-movement control are among the best under-stood in motor neurophysiology. Detailed anatomical and physiolog-ical data have paved the way for theoretical models that have unified existing knowledge and suggested further experiments. These mod-els have generally taken the form of black-box diagrams (for example, Robinson 1981) representing the flow of hypothetical signals between idealized signal-processing blocks. They approximate overall oculomo-tor behavior but indicate little about how real eye-movement signals would be carried and processed by real neural networks. Neurons that combine and transmit oculomotor signals, such as those in the vestibular nucleus (VN), actually do so in a diverse, seemingly ran-dom way that would be impossible to predict from a block diagram. The purpose of this study is to use a neural-network learning scheme (Rumelhart et al. 1986) to construct parallel, distributed models of the vestibulo-oculomotor system that simulate the diversity of responses recorded experimentally from VN neurons.

1 Introduction

The primary function of the VN is to relay head-velocity signals from the semicircular canals to extraocular-muscle motoneurons (Fig. 1). The purpose of this relay — the three-neuron-arc of the vestibuloocular re-flex (VOR) — is to stabilize images on the retina by producing eye movements that compensate for head movements (Wilson and Melvill Jones 1979). The VOR operates in all three rotational dimensions. Since canal and muscle coordinate frames are non-orthogonal and rotated from one another, a sensory/motor transformation must occur (Pelionisz and Llinás 1980) by combining inputs from the various canal pairs. The three-dimensional behavior of any VOR neuron can be specified by its

activation-vector, defined as the axis of head (or eye) rotation for which its change of activity is a maximum. The activation-vectors of the canal inputs are specified by canal geometry, as are those of the motoneurons by the pulling directions of the muscles (Robinson 1982). In contrast, single-unit recordings have shown that the activation-vectors of VN neurons are dispersed in various directions and do not align with either the canal or muscle vectors (Baker et al. 1984). Thus, the transformation is distributed over the VN.

In addition to vestibular signals, the VN also relays eye-velocity command signals from the smooth pursuit and saccadic oculomotor subsystems (Fig. 1). The pursuit system enables certain foveate animals, such as primates, to visually follow smoothly moving targets, while the saccadic system is used to rapidly change the direction of gaze from one target to another (Carpenter 1978). The pursuit and saccadic systems probably evolved on top of the basic, ubiquitous vestibulooocular reflex as the eyes developed foveas and moved to a forward-looking position. Saccades may have utilized the pre-existing quick-phase system, while pursuit possibly originated as an effort to suppress the basic vestibular reflex when not wanted. These later systems do, however, have some direct projections to motoneurons (not shown) as well as the VN-mediated connections. Thus, most vestibulo-oculomotor neurons (this includes cells in the nucleus prepositus hypoglossi and other areas in the caudal pons — but we will use VN for short) carry parts of all three commands. The change in discharge rate divided by the change in eye-velocity command is called the gain of each cell for each function: vestibular (V), pursuit (P) and saccadic (S). The three systems intermix on the cells of the VN on which can be found seemingly any possible combination of V, P, and S (Chubb et al. 1984; Fuchs and Kimm 1975; Miles 1974; Tomlinson and Robinson 1984). Thus, the three systems are diffusely represented and, again, their signals are distributed over this neuronal pool.

One job of the VN neural network is to ensure that eye velocity, represented by the activity of motoneurons, is actually that called for by the three input velocity commands. The network can do this job by adjusting synaptic weights, according to some optimization procedure such as error-driven learning. It will be appreciated at once that the vestibulo-oculomotor tasks to be learned are simple. Consequently, the purpose of this study is not to illustrate the sophistication of learning network models where the neurophysiological basis for the input, hidden, or output layers, and often all three, is incomplete or nonexistent. We have the opposite problem: the task to be performed is simple, but twenty years of research has provided a solid basis for the behavior of cells in all three layers. Thus, this study presents learning network models of simple functions that simulate known, but heretofore unexplained, neurophysiological properties of real VN neurons.
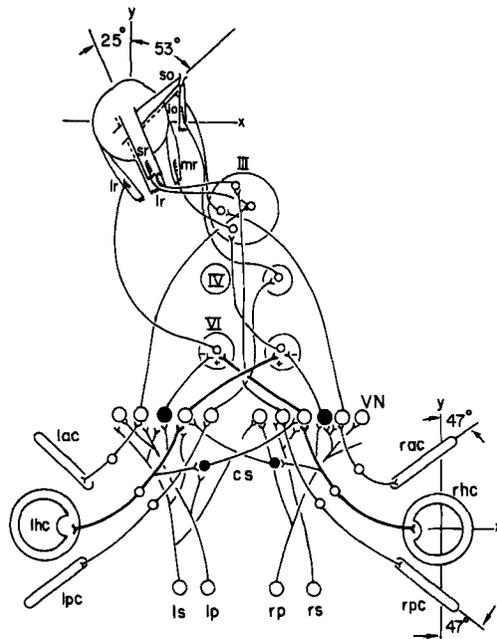
Figure 1: Schematic skeleton of the vestibulo-oculomotor system. Inputs to the vestibular nucleus (VN) include afferents from the six canals and projections from the pursuit (lp, rp) and saccadic (ls, rs) oculomotor subsystems. Outputs from the VN project to the extraocular muscle motoneurons of which are shown the six to the left eye. The initially random, divergent input projections onto VN neurons are depicted. The subsequent convergence of VN projections onto motoneurons is not shown for clarity. It is proposed that the cells in the VN act as a hidden layer in which the three eye-velocity command signals are distributed and, via modifiable synapses, maintain correct eye-velocity responses. Some of the major connections of the horizontal VOR are shown by heavy lines. The VN neuron that receives an excitatory connection from the rhc makes an excitatory connection to the motoneuron of the lr. The excitatory relay from the lhc to the VN and on to the motoneuron of the mr occurs through an internuclear interneuron connecting the abducens and oculomotor nuclei. The principal connections of the vertical VOR are also shown. A feedforward, inhibitory commissural system (cs) and some inhibitory VN neurons are shown by filled cells. Canal and muscle geometry for humans is taken from Robinson (Robinson 1982). III, IV, VI, oculomotor, trochlear and abducens nuclei; sr, ir, lr, mr, so, io, left superior, inferior, lateral and medial recti, and superior and inferior oblique muscles or motoneurons; lac, lpc, lhc, rac, rpc, rhc, left or right anterior, posterior and horizontal semicircular canal or primary afferent.

2 Neural Network Models

The vestibulo-oculomotor system was modelled as a network with three layers. The input layer carries the three eye-velocity commands into the network, the output layer represents the motoneurons of the muscles of the left eye, while units in the middle (hidden) layer represent VN neurons. The activity of each model unit, interpreted as its firing rate, is a sigmoidal function of the weighted sum of its inputs. Units respond almost linearly to midrange inputs but can be driven to cut-off or saturation with large negative or positive inputs, respectively. All input units projected to all hidden units. All hidden units projected to all output units and there were no direct input-to-output connections unless otherwise stated. There were also no intra-layer or feedback connections. Each input and output is represented as a single pair of units; only in the hidden layer, which can contain up to 40 units, is information represented in a distributed manner. The activities of vestibular and pursuit input pairs are modulated in push-pull within the linear range of the sigmoidal transfer function $(0.4 - 0.6)$ about a spontaneous (firing) rate (SR) of 0.5, while the burst (saccadic) input pair is modulated over the entire range $(0.0 - 1.0)$ and is spontaneously silent (SR = 0.0). All modifiable connections are randomized prior to training (range: -1.0 to 1.0, uniform distribution) and the networks are trained using backpropagation learning (Rumelhart et al. 1986). Briefly, the pattern of current outputs is compared to the desired outputs to form an error. This error is effectively propagated backward through the network to generate error patterns at each layer which are then used to modify the weights of the interlayer projections. The process is repeated for all inputs until, after many iterations, the errors for all outputs are less than a tolerance, in our case, 0.01.

3 The Simplest Arrangement

The three-neuron-arc of the VOR (heavy lines, Fig. 1) serves as a backbone for the networks to be described. To understand the basic manner in which the network organizes itself, we began with the simplest arrangement; the horizontal VOR was modelled with only two neurons in each layer. The input units represent primary afferents from the left and right horizontal canals (lhc and rhc, respectively), and the output units represent the motoneurons of the left lateral and medial rectus muscles (lr and mr, respectively, Fig. 1). The network was trained on the input/output table given in table 1. This table specifies inputs and outputs consistent with a compensatory VOR. For example, head rotation to the left causes the firing-rate (representing head velocity) of the lhc to increase (to 0.6) and that of the rhc to decrease (to 0.4). The compensatory eye movement to the right would then be produced by an equal increase in the

firing-rate (representing eye velocity) of the mr (to 0.6) and a decrease in that of the lr (to 0.4). The opposite pattern applies for head rotation to the right. The network finished learning this simple task after about 200 cycles through the input/output table and the values of the final weights are shown in table 2 where h1 and h2 are the hidden units. This arrangement, which is representative of every other run of this simulation, is similar in many ways to the actual circuitry of the VOR (Fig. 1). Hidden unit h1, which receives excitation from the lhc (1.63) and inhibition from the rhc (−1.21), is excitatory to the mr (1.84) and inhibitory to the lr (−1.71) while the opposite pattern is observed (with different values of weights) for hidden unit h2 (Table 2). These features loosely resemble the reciprocal innervation of motoneurons by VN neurons and the inhibitory commissural connections between the two vestibular nuclei (Wilson and Melvill Jones 1979). (Excitation and inhibition by the same model neurons simply implies the existence of an interneuron.) Thus, the learning algorithm clearly produces a set of connections consistent with the known organization of the VOR. The hidden and output unit responses for this run are given in table 3. The activities of the input units (lhc and rhc) are simply copied from table 1. The responses of the ouput units (lr and mr) match those desired (Table 1) to within the tolerance of 0.01, while the responses of the hidden units are more variable. For example, the SR of h1 (0.55) is its activity due to the head-still inputs (lhc = 0.5 and rhc = 0.5) via the weights in table 2. Similarly, the response of h1 to leftward head rotation (0.62) is its activity following forward propagation of the head-left input (lhc = 0.6 and rhc = 0.4). The SRs and vestibular gains, V, are presented in table 4. V is the change in activity from SR of a hidden or output unit divided by the change from 0.5 in the left input unit. For example, V for h1 would be: $(0.62 - 0.55)/(0.6 - 0.5) = 0.70$. (Gains P and S for units during pursuit and saccades, see below, are computed in an analogous way.)

Although the learning algorithm produces realistic connections in networks having only two hidden units, the rules of reciprocal innervation are sometimes violated in networks having more. In larger networks, such nonreciprocally-connected units always have very small gains. In large networks having all connections modifiable, the learning algorithm appears to converge to a configuration wherein the VOR transformation is supported by a few reciprocally connected, high gain units, after which the others, many nonreciprocally-connected, remain essentially unused. To prevent this development, the hidden-to-output connections were fixed in all subsequent examples. These fixed connections are arranged to provide reciprocal innervation of the motoneurons and all had the same absolute value. Its choice depended on the simulation being run.

	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

Table 1: Vestibular input/output table.

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

Table 2: Final weights of simple VOR model.

Input		Hidden		Output	
lhc	rhc	h1	h2	lr	mr
0.50	0.50	0.55	0.45	0.50	0.50
0.60	0.40	0.62	0.33	0.41	0.60
0.40	0.60	0.48	0.57	0.59	0.40

Table 3: Hidden and output unit responses for simple VOR model.

4 Distributed Representations

To explore the representation of multiple signals in the VN, we began by considering the interaction of vestibular and pursuit signals only. This representation is simulated on networks having 40 hidden units. The fixed hidden-to-output weights have the value 0.35 and the following reciprocal pattern: hidden units 1–20 inhibit the lr and excite the mr while the reverse holds for hidden units 21–40. The two input unit pairs represent vestibular inputs (from the lhc and rhc, as before) and left and right pursuit inputs (lp and rp, respectively, Fig. 1). The pursuit pair operates in the sense opposite to that of the vestibular pair in that

	Hidden		Output	
	h1	h2	lr	mr
SR	0.55	0.45	0.50	0.50
V	0.70	−1.20	−0.90	1.00

Table 4: Gain, V, and SRs of hidden and ouput units in simple VOR model.

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

Table 5: Vestibular-pursuit input/output table.

a left pursuit command (lp activated, rp inhibited) produces a left eye movement (lr activated, mr inhibited) and vice-versa. The vestibular-pursuit input/output table is given in table 5. The network learned this transformation after about 200 cycles. The vestibular and pursuit gains, V and P, of the hidden and output units were evaluated and V versus P is plotted for each unit in a typical run in figure 2. The most striking feature of this scatter-plot is the variety of combinations with which hidden units can carry the pursuit and vestibular signals, just as observed experimentally (Chubb et al. 1984; Fuchs and Kimm 1975; Miles 1974; Tomlinson and Robinson 1984).

To explore the interaction of vestibular, pursuit and saccadic eye-velocity commands, a third input pair was added to the 40 hidden unit network representing left and right saccadic inputs (ls and rs, respectively, Fig. 1). The saccadic pair operates in the same sense as the pursuit pair. The hidden-to-output connections are the same as for the vestibular-pursuit network. In addition, to reflect known anatomy, there are also direct connections from the saccadic inputs to the outputs; these are also fixed in a reciprocal pattern and have the value of 2.5. The vestibular-pursuit-saccadic input/output table is given in table 6. The network learned this transformation after about 1,000 cycles. The V, P, and S gains of the hidden and output units were calculated. The V and P gains were distributed much as is shown in figure 2. Saccadic behavior consisted of bursts and pauses of activity. Many hidden units burst for saccades in one direction (almost always the same as that for which pursuit activity

	Input						Output	
	lp	lhc	ls	rs	rhc	rp	lr	mr
head still	0.50	0.50	0.00	0.00	0.50	0.50	0.50	0.50
pursue left	0.60	0.50	0.00	0.00	0.50	0.40	0.60	0.40
pursue right	0.40	0.50	0.00	0.00	0.50	0.60	0.40	0.60
head left	0.50	0.60	0.00	0.00	0.40	0.50	0.40	0.60
head right	0.50	0.40	0.00	0.00	0.60	0.50	0.60	0.40
saccade left	0.50	0.50	1.00	0.00	0.50	0.50	1.00	0.00
saccade right	0.50	0.50	0.00	1.00	0.50	0.50	0.00	1.00

Table 6: Vestibular-pursuit-saccadic input/output table.

increased) and paused in the other direction. Some units would burst or pause for saccades in one direction and do nothing in the other direction. Several burst for saccades in both directions and others paused for saccades in both directions. This diversity of behavior is not unlike that seen experimentally (Chubb et al. 1984; Fuchs and Kimm 1975; Miles 1974; Tomlinson and Robinson 1984).

To explore the distributed representation of spatial information, the sensory/motor transformation between the vertical semicircular canals and the cyclovertical muscles of the left eye was modelled using a network having four input, four output and 40 hidden units. One pair of input units represented afferents from the right anterior (rac)-left posterior (lpc) canal pair and the other from the right posterior (rpc)-left anterior (lac) canal pair (Fig. 1). One pair of output units represented motoneurons of the left superior (so)-inferior oblique (io) muscle pair and the other of the left superior (sr)-inferior rectus (ir) muscle pair (Fig. 1). The hidden-to-output connections were fixed so that hidden units 1–10 and 21–30 reciprocally innervated so and io, and 11–20 and 31–40 reciprocally innervated sr and ir; their weights had the value 0.25. This vertical VOR network model was trained to produce input/output patterns that were consistent with compensatory eye-movements for head rotations about eight axes spaced 45 degrees apart in the horizontal plane. The activity of each input pair was modulated according to the projection of each head-rotation vector onto the activation-vectors of the canal pairs, while that of each output pair corresponded to the muscle pair components calculated to produce the compensatory eye movements using the angles for humans in figure 1. The network learned the transformation after about 500 cycles. Activation-vectors for the hidden and output units are plotted in figure 3. Those of the hidden units are dispersed in many directions and are qualitatively similar to those of cat VN neurons (Baker et al. 1984).
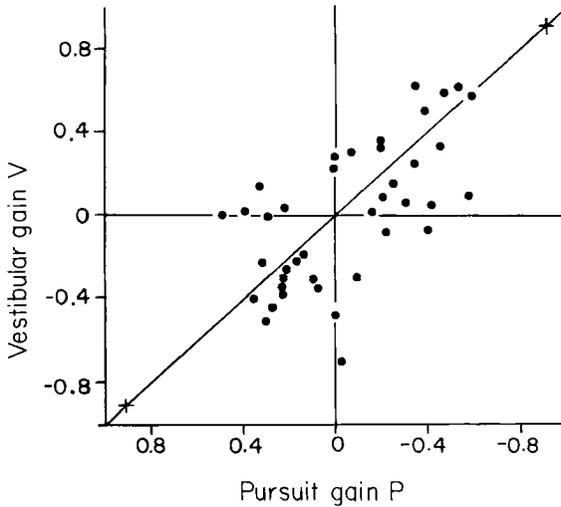
Figure 2: Plot of vestibular (V) versus pursuit (P) gains of hidden and output units in a 40 hidden unit, vestibular-pursuit neural network model. Because correctly wired units would have vestibular and pursuit gains of opposite signs, the polarity of the P axis is reversed for illustrative purposes. Hidden units (filled circles) falling on or near the horizontal or vertical axis can be considered as pure-pursuit or pure-vestibular units, respectively. The output units (+), having equal and opposite V and P values (by design), lie along the diagonal line. Hidden units falling on or near this line also have V and P values that are approximately equal and opposite. Most units are scattered throughout quadrants 1 and 3 and, having V and P of opposite sign, encode eye movements in the same direction. However, a small percentage of hidden units fall in quadrants 2 and 4 and have V and P of the same sign. Such seemingly anomalous units have been occasionally observed experimentally (Chubb et al. 1984; Fuchs and Kimm 1975; Miles 1974; Tomlinson and Robinson 1984).

5 Conclusions

This study indicates that a learning-network model provides a good explanation for the variability found in the neural organization of the vestibulo-oculomotor system in monkey and cat, and shows how neurons in the caudal pons combine and carry vestibulo-oculomotor signals in a distributed manner. Note that these models, however simple, are circuits equivalent to other multisynaptic, more realistic versions of the vestibulo-oculomotor system. No external teacher is needed since the error is retinal image motion which is well represented by signals in the brain stem and cerebellum. Further, the ability of retinal slip to produce plastic changes in the operation of the vestibulo-oculomotor system is
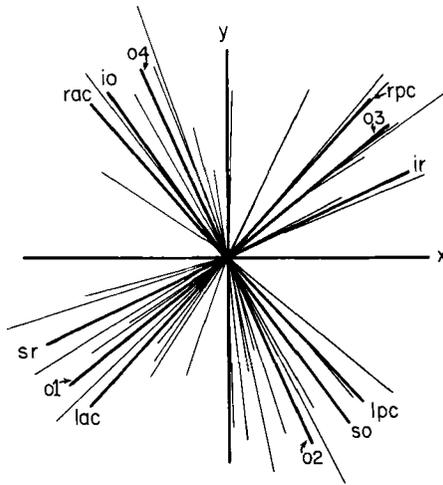
Figure 3: Activation-vectors of hidden and output units in a 40 hidden unit neural network model of the vertical VOR. The activation-vectors of the hidden and output units were determined by testing the mature network with the eight vectors of head-rotation used in training. A cosine function was fit to the eight responses of each unit; the magnitude and direction of the activation-vector was taken as the amplitude and phase of this cosine function. The canal activation vectors are lac, lpc, rac and rpc; those of the muscles (of the left eye) are o1, o2, o3 and o4. The latter do not coincide with the rotation axes of these four muscles, sr, so, ir and io, respectively, because the muscles form a skewed coordinate system (Robinson 1982). Note that the activation-vectors of the hidden units vary in magnitude and are dispersed in many directions. Activation-vector divergence is also characteristic of cat VN neurons (Baker et al. 1984).

well documented (Wilson and Melvill Jones 1979). To date, vestibulo-oculomotor modelling has been confined to the lumped, black-box approach. We hope that the results summarized here and presented in detail in two forthcoming articles (Anastasio and Robinson 1989; Anastasio and Robinson, in review) provide an alternative analysis technique that recognizes the distributed nature of vestibulo-oculomotor signals.

Additionally, it is interesting to note that when every model had its weights rerandomized and was retrained, it reached a different solution. In such networks there is no unique solution. Thus, each run of the 40 hidden unit signal combination model ended with a different distribution of V, P, and S gains. Rather than forming discrete subpopulations, hidden units seemed to fall along a continuum with regard to the relative strengths of these gains. In the two-dimensional spatial transfor-

mation model (Fig. 3), a 40-dimensional vector on the hidden units is collapsed to a two-dimensional motoneuron vector (sr − ir, so − io). This is sometimes referred to as the overcomplete problem since, in a closed form, there would be too many equations for too few variables to allow a unique solution. Our learning network is not bothered by this problem precisely because it does not seek a unique solution. Thus, the learning network approach suggests that, as far as the nervous system is concerned, the overcomplete problem may not be a problem at all.

Acknowledgments

References

Anastasio, T.J. and D.A. Robinson. 1989. The distributed representation of vestibulo-oculomotor signals by brain-stem neurons. *Biological Cybernetics,* in press.
——. Sensorimotor coordinate transformations: Learning networks and tensor theory. In review.
Baker, J., J. Goldberg, G. Hermann, and B. Peterson. 1984. Optimal response planes and canal convergence in secondary neurons in vestibular nuclei of alert cats. *Brain Research,* **294**, 133–137.
Carpenter, R.H.S. 1978. *Movements of the Eyes.* London: Pion Press.
Chubb, M.C., A.F. Fuchs, and C.A. Scudder. 1984. Neuron activity in monkey vestibular nuclei during vertical vestibular stimulation and eye movements. *Journal of Neurophysiology,* **52**, 724–742.
Fuchs, A.F. and J. Kimm. 1975. Unit activity in the vestibular nucleus of the alert monkey during horizontal angular acceleration and eye movement. *Journal of Neurophysiology,* **38**, 1140–1161.
Miles, F.A. 1974. Single unit firing patterns in the vestibular nuclei related to voluntary eye movements. *Brain Research,* **71**, 215–224.
Pelionisz, A. and R. Llinás. 1980. Tensorial approach to the geometry of brain function: Cerebellar coordination via a metric tensor. *Neuroscience,* **5**, 1125–1136.
Robinson, D.A. 1981. The use of control systems analysis in the neurophysiology of eye movements. *Annual Review of Neuroscience,* **4**, 463–503.
——. 1982. The use of matrices in analyzing the three-dimensional behavior of the vestibulo-ocular reflex. *Biological Cybernetics,* **46**, 53–66.
Rumelhart, D.E., G.E. Hinton, and R.J. Williams. 1986. Learning internal representations by error propagation. *In:* Parallel Distributed Processing: Explorations in the Microstructure of Cognition, 1, *Foundations,* eds. D.E. Rumelhart, J.L. McClelland, and The PDP Research Group, 318–362. Cambridge: MIT Press.

Tomlinson, R.D. and D.A. Robinson. 1984. Signals in vestibular nucleus mediating vertical eye movements in the monkey. *Journal of Neurophysiology*, **51**, 1121–1136.

Wilson, V.J. and G. Melvill Jones. 1979. *Mammalian Vestibular Physiology*. New York: Plenum Press.