

Distributed Parallel Processing in the Vertical Vestibulo-Ocular Reflex: Learning Networks Compared to Tensor Theory

T. J. Anastasio¹ and D. A. Robinson²

¹ Department of Otolaryngology – Head and Neck Surgery, University of Southern California, Los Angeles, CA 90033, USA

² Department of Ophthalmology and Biomedical Engineering, The Johns Hopkins School of Medicine, Baltimore, MD 21205, USA

Received November 21, 1989/Accepted in revised form February 26, 1990

Abstract. The vestibulo-ocular reflex (VOR) is capable of producing compensatory eye movements in three dimensions. It utilizes the head rotational velocity signals from the semicircular canals to control the contractions of the extraocular muscles. Since canal and muscle coordinate frames are not orthogonal and differ from one another, a sensorimotor transformation must be produced by the VOR neural network. Tensor theory has been used to construct a linear transformation that can model the three-dimensional behavior of the VOR. But tensor theory does not take the distributed, redundant nature of the VOR neural network into account. It suggests that the neurons subserving the VOR, such as vestibular nucleus neurons, should have specific sensitivity-vectors. Actual data, however, are not in accord. Data from the cat show that the sensitivity-vectors of vestibular nucleus neurons, rather than aligning with any specific vectors, are dispersed widely. As an alternative to tensor theory, we modeled the vertical VOR as a three-layered neural network programmed using the back-propagation learning algorithm. Units in mature networks had divergent sensitivity-vectors which resembled those of actual vestibular nucleus neurons in the cat. This similarity suggests that the VOR sensorimotor transformation may be represented redundantly rather than uniquely. The results demonstrate how vestibular nucleus neurons can encode the VOR sensorimotor transformation in a distributed manner.

Introduction

The primary function (Fig. 1) of many vestibular nucleus (VN) neurons is to relay head-velocity signals from semicircular canal primary afferents to extraocular muscle motoneurons. The purpose of this relay, commonly known as the three-neuron-arc of the vestibulo-ocular reflex (VOR), is to stabilize the retinal image by producing eye movements that compensate for head movements (Wilson and Melvill Jones 1979). The VOR operates in three dimensions. In frontal-eyed mammals,

the VOR produces almost compensatory eye rotations for yaw and pitch head rotations, but eye movements in response to roll are less than compensatory (*ibid.*). To simplify the development to follow, however, it will be assumed that the VOR produces perfectly compensatory eye movements in all three rotational dimensions.

In Fig. 1, the geometrical arrangement of the six canals and the extraocular muscles of the left eye in the cat are shown; geometrical data are from Ezure and Graf (1984). Because the coordinate frames defined by the geometry of the semicircular canals and extraocular muscles are not orthogonal and differ from one another, a sensorimotor transformation must occur within the relay. The tensor theory of the VOR (Pellionisz and Llinas 1980; Pellionisz and Graf 1987) utilizes certain mathematical forms to develop a linear, matrix transformation that can describe the three-dimensional behavior of this reflex. Tensor theory specifies unique connectivity patterns for the synaptic projections from the canal afferents to the VN neurons and from the VN neurons to the motoneurons. By seeking a unique solution, it describes the connections from the six semicircular canals to the six extraocular muscles of one eye (Pellionisz and Graf 1987) as though they were mediated by only six VN neurons. Given that there are many hundreds of such interneurons, this mathematical description indicates little about how individual VN neurons actually behave.

The three-dimensional behavior of any vestibulo-ocular neuron can be described by its sensitivity-vector, defined as the axis of head (or eye) rotation for which its activity is a maximum. The sensitivity-vectors of the canal afferents are specified by canal geometry just as those of the motoneurons are specified by the pulling directions of the muscles (Robinson 1982). Since tensor theory describes a unique pattern of connectivity from the canal afferents to the VN neurons, it also specifies a unique set of sensitivity-vectors for these neurons. If such a unique connectivity pattern existed, then the sensitivity-vectors of real VN neurons should line-up along these specified directions. Instead, single-unit recordings in the cat have shown that the sensitivity-




Fig. 1. Schematic diagram of part of the vestibulo-ocular reflex (VOR). The vestibular nucleus (VN) is represented as a group of neurons arranged in a line near the center of the figure. Inputs to these neurons are semicircular canal primary afferents from each of the six canals. Outputs from the VN neurons project to the motoneurons of the six extraocular muscles of the left eye. Head-velocity signals from all of the canals converge onto the VN neurons which relay the results of this convergence to the motoneurons. The initially random, divergent input projections onto VN neurons are depicted. The subsequent convergence of VN neuron projections onto motoneurons is not shown for clarity. Some of the principal connections of the horizontal VOR are shown in boldface. A feed-forward, inhibitory commissural system (cs) and some inhibitory VN neurons are also shown; filled cells are inhibitory. Canal and muscle geometry for cat (Ezure and Graf 1984) is shown. III, IV, VI, oculomotor, trochlear and abducens nuclei; sr, ir, lr, mr, left superior, inferior, lateral and medial rectus muscle or motoneuron; so, io, left superior and inferior oblique muscle or motoneuron; lac, lpc, lhc, rac, rpc, rhc, left or right anterior, posterior and horizontal semicircular canal or primary afferent

vectors of VN neurons are dispersed in various directions rather than aligning uniquely with any particular vectors (Baker et al. 1984). These data indicate that tensor theory is not sufficient to describe the behavior of individual VN neurons and suggest that another mechanism is responsible for the neural organization of the VOR sensorimotor transformation.

We modeled the vertical VOR as a three-layered neural network programmed using the back-propagation learning algorithm (Rumelhart et al. 1986). The units in mature networks had divergent sensitivity-vectors which resembled those of actual VN neurons in the cat. Also, different runs of the same transformation

resulted in different distributions of sensitivity-vectors, indicating that there is no unique network solution to the problem. These results suggest that error-driven learning exerts an organizing influence on the neural circuitry of the VOR. They extend tensorial descriptions to the level of individual neurons.

Model Architecture

The vestibulo-ocular system is modeled as a neural network consisting of three layers of processing units (Anastasio and Robinson 1989). Units in the input layer represent canal afferents that carry head-velocity signals into the network. Units in the output layer represent the motoneurons of the extraocular muscles of the left eye, while those in the middle (hidden) layer represent VN neurons. The activity of each unit, interpreted as its firing-rate, is a nonlinear, sigmoidal function of the weighted sum of its inputs (Rumelhart et al. 1986). As such, units respond linearly to midrange inputs but can be driven into cut-off or saturation with large negative or positive inputs, respectively. Network architectures are kept as simple as possible, with complete connectivity from input-to-hidden layers and no direct input-to-output connections. There are no intra-layer or feedback connections. Inputs to the hidden or output layers do not include a bias (true) unit. Inputs and outputs are represented as single pairs of units; only in the hidden layer, which can contain up to 40 units, is information represented in a truly distributed manner. All modifiable connections are randomized prior to training (range: -1.0 to 1.0 , uniform distribution) and the networks are trained using back-propagation learning (Rumelhart et al. 1986). Briefly, an input pattern is forward propagated through the network and the pattern of actual 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 inter-layer projections. The process is repeated for all input/output patterns until, after many iterations, the errors for all patterns are less than a tolerance, in our case, 0.01.

To explore the sensorimotor transformation between the vertical semicircular canals and the cyclovertical muscles of the left eye, we started with a network having four input, four output and four hidden units. One pair of input units represented afferents from the right anterior(rac)-left posterior(lpc) canal pair (ralp) and the other afferents from the right posterior(rpc)-left anterior(lac) canal pair (rpla) (Fig. 1). One pair of output units represented motoneurons of the left superior-inferior oblique (sio) muscle pair and the other motoneurons of the left superior-inferior rectus (sir) muscle pair (Fig. 1). Because evidence to date suggests synaptic plasticity only at the level of the VN (Miles and Lisberger 1981), a constraint was placed on the model which allowed only the input-to-hidden synapses to be modifiable. This was done by fixing the hidden-to-output connections so that, for the four-hidden-to-

Table 1. Canals and muscles labelled as in Fig. 1. h1, h2, h3 and h4 are the four hidden units. Matrices are based on canal and muscle angles for the cat shown in Fig. 1

A. Weights of the fixed hidden-to-output connections

		from	h1	h2	h3	h4
to	sr		1.00	0.00	-1.00	0.00
	so		0.00	1.00	0.00	-1.00
	ir		-1.00	0.00	1.00	0.00
	io		0.00	-1.00	0.00	1.00

B. Canal matrix

$$\begin{vmatrix} \text{rac-lpc} \\ \text{rpc-lac} \end{vmatrix} = \begin{vmatrix} -0.682 & 0.731 \\ 0.682 & 0.731 \end{vmatrix} \begin{vmatrix} \dot{H}_x \\ \dot{H}_y \end{vmatrix}$$

C. Motor matrix (inverse)

$$\begin{vmatrix} \text{so-io} \\ \text{sr-ir} \end{vmatrix} = \begin{vmatrix} 0.486 & -0.876 \\ -0.840 & -0.546 \end{vmatrix} \begin{vmatrix} \dot{H}_x \\ \dot{H}_y \end{vmatrix}$$

D. Examples of final input-to-hidden weights

		from	lac	lpc	rpc	rac
to	h1		3.63	-1.19	-3.53	1.09
	h2		0.94	2.78	-0.57	-3.14
	h3		-3.52	0.21	4.31	-1.00
	h4		-1.04	-4.41	1.27	4.19

model, hidden units 1 and 3 reciprocally innervated the sir motoneurons and 2 and 4 reciprocally innervated the sio motoneurons; hidden-to-output weights had the absolute value of 1.0. This pattern of hidden-to-output weights is given in Table 1A. Fixing the hidden-to-output weights in this way restricted the convergence of canal inputs to the hidden layer.

Training the Network

The coordinate transformation to be learned was based on the geometry in Fig. 1, for the cat (Ezure and Graf 1984), where x is the pitch axis (positive to the right) and y is the roll axis (positive anterior). The actions of the canals and muscles are described by sensitivity and action vectors perpendicular to their planes of activity (Robinson 1982). The canals are maximally activated by head rotations about their sensitivity-vectors. The muscle action-vectors are the axes of rotation of the eye when each muscle acts alone. These vectors are shown in Fig. 2 for the four canals and the four left muscles as seen looking down on the horizontal plane. The direction of the canal vectors obey the right-hand rule. Those of the muscle vectors are reversed because the VOR requires an opposite eye movement. Thus, in Fig. 2, when rpc is stimulated, ir should contract, etc. This figure illustrates the transformation required of the hidden units.

The network was trained to make a compensatory response for rotations about eight horizontal axes at every multiple of 45 deg. There was a training vector in each of the cardinal directions (anterior, posterior, lateral and medial) and at 45 deg to these vectors. The input and output neurons were required to modulate in




Fig. 2. The sensitivity-vectors and muscle vectors for the two-dimensional model of the vertical VOR with four hidden units as seen looking down on the horizontal plane. x is right, y is anterior. The sensitivity-vectors of the canals are labeled as in Fig. 1. The sensitivity-vectors of the hidden units are h1, h2, h3, and h4, those of the output motoneurons are o1, o2, o3 and o4. The action-vectors of the muscles are labeled as in Fig. 1. The latter do not coincide with the sensitivity-vectors of the motoneurons because the muscles form a skewed coordinate system. The polarity of the muscle vectors are chosen to show which muscle is maximally activated by stimulation of each canal. In this example, the sensitivity-vectors (magnitudes scaled) of the four hidden units align themselves roughly near those of the output units

push-pull within the linear range of the sigmoidal transfer function (0.4–0.6) about a spontaneous (firing-) rate (SR) of 0.5. The input/output values were calculated for each rotation as follows. The rotation vector, of length 0.1, was projected onto the canal axes to find their modulation. This was done by multiplying each head rotation vector by the canal matrix for cat (Table 1B), yielding the required amount of modulation of the ralp and rpla canal afferent pairs, respectively. For the muscles one must find a combination of the sio and sir pairs that combine vectorially to produce an eye movement equal and opposite to a given head movement. This was done by multiplying each head rotation vector by the inverse of the motor matrix (Robinson 1982) for cat (Table 1C), yielding the required amount of modulation of the sio and sir motoneuron pairs, respectively. For canal afferents and motoneurons alike, the modulation amount calculated in this way was added to the SR (0.5) of the first named unit in each pair and subtracted from the SR of the second named unit in each pair. For example, head rotation about the positive x axis (up in pitch) would be represented by the vector (0.1, 0.0). Multiplication of the canal and muscle matrices (Table 1B and C) by this vector yields the amounts of modulation of the canal afferents and motoneurons that should

be produced by pitch-up head rotation. Thus, the firing rate of *rac* decreases and *lpc* increases by 0.068, while *rpc* increases and *lac* decreases also by 0.068. Similarly, *sr* increases and *io* decreases by 0.049, while *sr* decreases and *ir* increases by 0.084. This procedure determined the input/output table for all eight axes. All networks were trained to within the required tolerance after about 250 cycles.

Network Analysis

The final weights of a typical network appear in Table 1D. With the configuration of the hidden-to-output weights (Table 1A), *h2* and *h4* must subserve the main *ralp* to *sio* transformation while *h1* and *h3* must subserve the main *rpla* to *sir* transformation. Therefore, *h2* and *h4* should receive strong, push-pull connections from *rac* and *lpc*, and weaker ones from *rpc* and *lac* while *h1* and *h3* should receive strong, push-pull connections from *rpc* and *lac*, and lesser ones from *rac* and *lpc*. The learned weights shown are in accord. (Note that when any one unit both excites and inhibits target units, an interneuron with unity gain is implied to satisfy Dale's law.)

The values of the weights in the mature network (Table 1D) are in good qualitative agreement with the connection strengths that are predicted using tensor theory (Pellionisz and Graf 1987). When all of the proposed tensorial transformations occurring at each stage are combined into one overall "brainstem matrix" (Pellionisz and Graf 1987), the resulting connections agree with the present results in sign and relative value. Slight disagreements are due to the push-pull interactions of the hidden units, which allow some indeterminacy. Thus back-propagation and tensor analysis compute similar transformations in the non-redundant case, as expected.

The mature model was evaluated by rotations about the eight training axes. Unit responses were a trigonometric function of axis angle and a cosine was fit to them using a least squares algorithm. The magnitude and direction of the sensitivity-vector for each unit is the amplitude and phase of the cosine fit. The sensitivity-vectors (magnitudes scaled) of all the hidden and output units appear in Fig. 2. The sensitivity-vectors of the muscles (the axis of greatest modulation: *o1*–*o4*) do not coincide with their action-vectors (around which they rotate the eye: *sio*–*sir*) because the action-vectors do not form a rectangular coordinate system (ie projections and components are not the same vectors). The sensitivity-vectors of the hidden units lie close to those of the output units, indicating that the sensorimotor transformation has taken place at the hidden unit level.

The individual hidden unit vectors do not exactly coincide with their associated output unit vectors. This occurred because each hidden unit pair was constrained to reciprocally innervate its agonist/antagonist motoneuron pair. Thus it was the push-pull interaction of the hidden unit responses that had to result in the




Fig. 3. Sensitivity-vectors of hidden and output units (magnitudes scaled) in a neural-network model of the vertical VOR with 40-hidden-units. The sensitivity-vectors of the hidden units vary in magnitude and are dispersed in many directions. This divergence is qualitatively similar to that observed for actual VN neurons in the cat (Baker et al. 1984)

desired output vector; the individual hidden unit vectors could vary somewhat. In all cases, the vector representing the average of any hidden unit vector and the reflection of the vector of its reciprocal partner always exactly coincided with the associated output unit vector.

To be more realistic, the hidden layer was expanded to 40 units. This simulated the abundance of VN neurons which actually represent the vertical VOR sensorimotor transformation. The fixed hidden-to-output weights were similar to those in Table 1A except that ten hidden units had each pattern of connectivity and their absolute value was 0.25. The network was trained as before, requiring about 500 cycles. Sensitivity-vectors for all 40 hidden and the four output units are shown in Fig. 3 for a typical network.

The vectors of the hidden units are dispersed. Instead of aligning with canal or motor vectors, their sensitivity-vectors pointed in many directions indicating that they had received canal convergence in a variety of combinations. The vectors do tend to cluster somewhat around those of the muscles and canals, but are spread widely. This dispersion of hidden unit sensitivity-vectors is qualitatively similar to that reported for cat VN neurons (Baker et al. 1984).

The results of the four- and 40-hidden-unit networks presented above are representative of all other runs of these models, but no two runs were the same. Every time these models had their input-to-hidden weights randomized and were retrained, they reached

different final weight matrices and sensitivity-vector distributions. This underscores the fact that these distributed, neural-network models of the vertical VOR have no unique solutions (see below).

Testing the Network

The VOR produces the correct compensatory eye movement regardless of the rotational axis of the imposed head rotation. To test whether the distributed representations constructed by back-propagation could generalize, both four- and 40-hidden-unit, mature network models were tested using axes of head rotation for which they had not been trained. The testing set consisted of the eight head rotation axes lying midway between the eight training axes. The mature networks were tested using both the training and testing sets and the errors between actual and desired outputs were computed. In both the four- and 40-hidden-unit models tested on the training set, worst case error for all output units and all patterns was less than 0.0100 (by design), and the average error was less than 0.0060 for both. Worst case error for the testing set was 0.0106 and 0.0107 for the four- and 40-hidden-unit models, respectively, and the average error was again less than 0.0060 for both. Thus, vertical VOR models utilizing the back propagation algorithm can extrapolate their learning to novel vectors. Extrapolation of learning is a feature of distributed representations in general (Hinton et al. 1986).

The conclusions reached here and supported by experimental evidence suggest that the sensorimotor transformation of the vertical VOR is subserved in a redundant way with an abundance of neurons. One may wonder why the nervous system would represent such a simple transformation with so many neurons. A simple explanation involves fault tolerance. A transformation subserved by few units would suffer more from the loss of one unit than would a transformation subserved by many. This was demonstrated by subjecting mature, four- and 40-hidden-unit models to the removal of the second and the 20th hidden units, respectively. Given the observed variability in the vector contributions of each unit in different runs, the average error due to the removal of the designated unit in the magnitude and direction of the vector of output unit 4 was computed over ten runs (the magnitude and direction of output unit 2 was affected in a similar way, but output units 1 and 3 were unaffected due to the fixed hidden-to-output connections). On average, the magnitude of the sensitivity-vector of output unit 4 decreased by 51.3% and 5.2% in four- and 40-hidden-unit models, respectively. The error in output unit magnitude is ten times greater in four- than in 40-hidden-unit models, as would be expected. The average change in vector direction (angle) was 24.4% and 5.2% in four- and 40-hidden-unit models, respectively. Directions are also affected more in four- than in 40-hidden-unit models.

Neurophysiological Validity

Our model of a sensorimotor transformation between the vertical canals and cyclovertical muscles shows that the transformation is distributed. The sensitivity-vectors of the hidden units (model VN neurons) are widely dispersed, having a variety of angles and amplitudes. This behavior is qualitatively similar to what is actually observed for VN neurons in the cat. Using head rotations in the planes of the various canal pairs, Curthoys and Markham (1971) reported that, of all VN neurons tested which received horizontal canal pair input, approximately 50% also responded unambiguously to small (4–6 deg/s/s) brief (3–4 s) rotational accelerations in the planes of one or both of the vertical canal pairs. By using sinusoidal rotations in a number of different planes, Baker et al. (1984) determined the axis of head rotation for which the sensitivity of each VN neuron was maximal. These sensitivity-vectors were found to be dispersed over a wide region around those of the canals, just as in Fig. 3. This behavior requires that canal signals converge. Markham and Curthoys (1972) found that monosynaptic activation of a second-order VN neuron rarely occurred upon electrical stimulation of more than one ampullary nerve in the cat. This suggests that convergence on a given second-order unit is mediated by interneurons and could be accommodated in our model by feed-forward onto hidden units from several canals via relay units.

Graf and Ezure (1986) in the cat showed that VN neurons projected to more than one pool of motoneurons, indicating that convergence can occur on motoneurons as well. (Whether these weights are modifiable is not known.) The implications of convergence onto motoneurons was investigated in four- and 40-hidden-unit models. A four-hidden-unit simulation was run having the input-to-hidden weights fixed in a reciprocal pattern (as in Table 1A) and the hidden-to-output weights modifiable. This constraint required that the sensorimotor transformation take place at the output unit level, rather than at the hidden unit level as before. The hidden-to-output weight matrix of the mature network was similar to the previous input-to-hidden weight matrix given in Table 1D. This hidden-to-output matrix was then used to fix the hidden-to-output weights in a 40-hidden-unit simulation; ten hidden units had each pattern of connectivity, but scaled down by a factor of ten. Thus the hidden-to-output weights were fixed in a pattern capable of producing the sensorimotor transformation, while the input-to-hidden weights were modifiable. The sensitivity-vectors of the hidden units in this simulation varied in magnitude and direction just as in Fig. 3 (but of course their actual distribution was different). In another 40-hidden-unit simulation, both the input-to-hidden and hidden-to-output weights were modifiable, and again the hidden unit vectors varied in magnitude and direction. These results indicate that, although the fixed hidden-to-output weights in our vertical VOR models (Table 1A) restricted convergence to the hidden unit layer, this constraint did not obscure the point that the

transformation from canal to muscle coordinates was distributed.

In another simulation, a 40-hidden-unit network was allowed to continue learning for over 25,000 cycles. Upon termination of learning, the average error for all outputs and all patterns was 0.0003, more than an order of magnitude lower than that for networks trained to the required tolerance of 0.01. The sensitivity-vectors in this simulation were again dispersed widely, indicating that hidden unit variability is independent of tolerance level. This result is expected, since all of the infinitely many neural-network solutions for the sensorimotor transformation are equally precise.

Although no claim is made that the back-propagation algorithm is physiological, the algorithm as used in our special case may correspond to a real learning mechanism. Because the hidden-to-output connections were fixed, the error signal that was "back-propagated" to the hidden units was proportional to retinal slip (Anastasio and Robinson 1989). The algorithm specifies that the change in input-to-hidden weights is mainly proportional to the error signal (retinal slip) and current input-unit (canal afferent) firing rate (*ibid.*). The retinal slip signal can be found in the accessory optic system and cerebellum – the latter is known to send a heavy, direct projection to the vestibular nuclei (Wilson and Melvill Jones 1979) and recent anatomical studies in the cat suggest that an indirect projection from the former also exists (Magnin et al. 1983). Also, the ability of retinal slip to produce plastic changes in the operation of the VOR is well documented (Wilson and Melvill Jones 1979). Thus, a learning mechanism that modifies vestibular afferent synapses to VN neurons according to the levels of vestibular input and retinal slip error is not improbable. Such a mechanism (or mechanisms) may operate both during development, as an ontogenetic process, and throughout life, to maintain proper VOR functioning under changing circumstances.

Implications of Distributed Processing

There are two theoretical implications about coordinate transformations by learning networks such as we used. Once the network has learned its task, what it does can be described in terms of covariant and contravariant vectors and metric tensors (Robinson 1982; Pellionisz and Graf 1987). A rigorous critique of the use of tensor theory for general sensorimotor transformations is beyond the scope of this study (but see Arbib and Amari 1985). For the special case of linear (Euclidean) spaces, a given vector in any coordinate frame has both a covariant and a contravariant representation (Pellionisz and Graf 1987). The projections of the vector onto the coordinate axes constitute a covariant representation while the components of the coordinate axes composing the vector (by the parallelogram rule) constitute a contravariant representation. In orthogonal coordinate frames covariant and contravariant representations are identical. Covariant and contravariant vector represen-

tations are related through a matrix known as the metric tensor – for orthogonal coordinate frames the metric tensor is the identity matrix.

For the VOR sensorimotor transformation, one can derive a canal matrix that converts a head-velocity vector from x - y to the canal coordinates, and also finds the projections of that vector onto the canal axes to produce a covariant canal vector. This canal matrix is represented mathematically by multiplying the coordinate transformation matrix by the metric tensor of the canal space. A motor matrix describes the coordinate transformation from the muscle axes back to the x - y coordinates. Since the motoneuron vector is the sum of vector components, it is a contravariant vector. Consequently, the hidden units are not only producing a coordinate transformation between the canal and muscle systems, but they also contain the inverse of the metric tensor to change the covariant canal vector into the contravariant muscle vector.

These concepts were introduced into neurophysiology by Pellionisz and Llinas (1980) who felt that the nervous system might be organized according to the rules of matrix algebra, performing computations such as the extraction of metric tensors and the determination of matrix inverses. They further postulated that the tensorial transformation is broken down into separate stages, with each station of the VOR subserving a separate and well defined portion of the overall computation. The present study does not dispute the utility of these man-made mathematical concepts in describing the functioning VOR, but it does question the notion that such concepts are specifically recognized by the learning nervous system. Our VOR neural-network models were charged with one goal – to eliminate the error – and they did this with a single layer of hidden units. Examining the hidden units illustrated in Fig. 3 would never lead one to interpret their behavior in terms of extracting a contravariant from a covariant vector or performing a matrix inversion. In our models, the synapses on the hidden units are modified by influences for which these concepts are irrelevant. Our results, supported by neurophysiological data, suggest that neural networks subserving sensorimotor transformations are shaped according to the exigencies of error-driven learning rather than by mathematical constructs such as tensor theory.

The second implication concerns overcompleteness. Tensor theory treats each of the six extraocular muscles independently (Pellionisz and Graf 1987). Since there are six muscles but only three rotational degrees of freedom for the eye, there are three equations with six unknowns relating eye muscle activity to eye rotation – an overcomplete system. Obviously there is no unique solution – there are an infinite number. Pellionisz has pointed out this problem and has suggested mechanisms (such as the Moore–Penrose generalized inverse) to make a unique solution possible (*ibid.*). We avoided this problem at the output stage of our models by rigidly coupling the eye muscles into agonist/antagonist pairs through reciprocal innervation. This reduced the number of degrees of freedom of the motor command

to three (Robinson 1982). But this reduction in dimensionality at the output stage does not necessarily reduce the redundancy at the hidden-unit level. Strictly speaking, any representation of the VOR sensorimotor transformation will be overcomplete if more than one hidden unit per dimension is used. It is clear that the actual VOR, subserved by an abundance of VN neurons, is extremely overcomplete.

Overcompleteness, however, is not a problem for learning networks. In our coordinate transformation model with 40 hidden units and two pairs of output units, each pair representing one degree of freedom, there is a 20:1 collapse in dimensions; the hidden-to-output relations would be described by two independent equations with 40 unknowns. Overcompleteness is simply irrelevant in such neural networks. The basic reason is that the network is not seeking a unique solution. It just wants one that works and there are an infinite number of them. The problem of overcompleteness is largely one of our own making because we are so used to thinking that a problem in physics must have a unique solution – the nervous system is under no such constraint.

Acknowledgements. This work was supported by the National Eye Institute of the National Institutes of Health, USA. We thank A. L. McCracken for preparing the manuscript and C. V. Bridges for illustrations.

References

- Anastasio TJ, Robinson DA (1989) The distributed representation of vestibulo-oculomotor signals by brain-stem neurons. *Biol Cybern* 61:79–88
- Arbib MA, Amari S-I (1985) Sensori-motor transformations in the brain (with a critique of the tensor theory of cerebellum). *J Theor Biol* 112:123–155
- Baker J, Goldberg J, Hermann G, Peterson B (1984) Optimal response planes and canal convergence in secondary neurons in vestibular nuclei of alert cats. *Brain Res* 294:133–137
- Curthoys IS, Markham CH (1971) Convergence of labyrinthine influences on units in the vestibular nuclei of the cat. I. Natural stimulation. *Brain Res* 35:469–490
- Ezure K, Graf W (1984) A quantitative analysis of the spatial organization of the vestibulo-ocular reflexes in lateral and frontal-eyed animals. I. Orientation of semicircular canals and extraocular muscles. *Neuroscience* 12:85–93
- Graf W, Ezure K (1986) Morphology of vertical canal related second order vestibular neurons in the cat. *Exp Brain Res* 63:35–48
- Hinton GE, McClelland JL, Rumelhart DE (1986) Distributed representations. 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 77–109
- Magnin M, Courjon JM, Flandrin JM (1983) Possible visual pathways to the cat vestibular nuclei involving the nucleus prepositus hypoglossi. *Exp Brain Res* 51:298–303
- Markham CH, Curthoys IS (1972) Convergence of labyrinthine influences on units in the vestibular nuclei of the cat. II. Electrical stimulation. *Brain Res* 43:383–397
- Miles FA, Lisberger SG (1981) Plasticity in the vestibulo-ocular reflex: A new hypothesis. *Ann Rev Neurosci* 4:273–299
- Pellionisz A, Graf W (1987) Tensor network model of the “three-neuron vestibulo-ocular reflex-arc” in cat. *J Theor Neurobiol* 5:127–151
- Pellionisz A, Llinas R (1980) Tensorial approach to the geometry of brain function: Cerebellar coordination via a metric tensor. *Neuroscience* 5:1125–1136
- Robinson DA (1982) The use of matrices in analyzing the three-dimensional behavior of the vestibulo-ocular reflex. *Biol Cybern* 46:53–66
- 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
- Wilson VJ, Melvill Jones G (1979) *Mammalian vestibular physiology*. Plenum Press, New York
- Thomas J. Anastasio, Ph.D.
USC Vestibular Laboratory
PMB C-103
1420 San Pablo Street
Los Angeles, CA 90033
USA