PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

A Neural Network Model of the Vestibulo-Ocular Reflex Using a Local Synaptic Learning Rule [and Discussion]

Donald B. Arnold, David A. Robinson, J. G. Taylor and P. J. Simpson

Phil. Trans. R. Soc. Lond. B 1992 **337**, 327-330
doi: 10.1098/rstb.1992.0110

A neural network model of the vestibulo-ocular reflex using a local synaptic learning rule

DONALD B. ARNOLD AND DAVID A. ROBINSON

Departments of Biomedical Engineering and Ophthalmology, The Johns Hopkins University, School of Medicine, Baltimore, Maryland 21287-9131, U.S.A.

SUMMARY

Vertebrates use the vestibulo-ocular reflex to maintain clear vision during head movements. This reflex requires eye-velocity commands from the semicircular canals to be integrated (mathematically) to produce eye-position commands for the extraocular muscles. This is accomplished by a neural network in the caudal pons. A model of this network is proposed using positive feedback via lateral inhibition. The model has been adapted to a learning network. We have developed a synaptic learning rule using only local information to make the model more physiological.

1. INTRODUCTION

When flies and other winged insects zoom about with amazing angular velocities and accelerations, one must wonder what they can see during such maneuvers. Some, such as the hover fly, can hold quite still to examine a specific object, often a fly of the opposite sex, but in between, visual acuity must be seriously degraded during high-speed twists and turns unless they have visual capacities that we little suspect. Some invertebrates, unable to move their eyes, wiggle their retinas instead, but generally to scan their retinas across images of objects rather than following them. Invertebrates have invented an impressive variety of often ingenious ways of dealing with their visual worlds.

Vertebrates, on the other hand, have decided on a fixed way to solve the problem of vision during head movements. We keep our visual axes stationary in space when the head moves by moving our eyes in the opposite direction at the same speed. This reflex, the vestibulo-ocular reflex (VOR), evolved a long time ago. It allowed animals clear vision even though their heads were in motion and has been so successful that it has changed almost not at all over the eons. This ability is desirable whether you are a predator with forward-pointing eyes and a fovea, or a prey that is afoveate and lateral-eyed, so that the reflex appears in almost identical form in such widely disparate vertebrates as fish, birds and mammals.

A VOR might also be useful for robots. When we move in a visual environment, images flow across our retinas. This optic flow is enormously rich in information. Images of near objects flow more rapidly, giving us depth and form information. It is remarkable that most visual physiologists fail to appreciate that this primordial system must be a basic substrate for the mammalian visual cortex. Optic flow results from both translation and rotation within a visual environment. Visual robotocists worry about how to separate them. The vertebrate oculomotor system separates them very simply: it prevents rotational optic flow by stabilizing the angle of the visual axes in space. This solution suggests that robotic visual systems might also benefit from a VOR.

Our concern in this problem is to describe, by modelling, the operation of the VOR down to the single-unit level. When the head turns, the semicircular canals report it by modulating the discharge rates of primary vestibular afferents in proportion to head velocity. These push–pull signals enter the vestibular nuclei and pass directly to the motoneurons, also in push–pull, to generate a compensatory eye movement. Even a cursory attempt to deduce the signal processing in this reflex leads one quickly to the detection of the neural integrator. The signal from the canals is proportional to instantaneous head velocity. When passed on toward the appropriate motoneurons it becomes an eye velocity command. All very well, but muscles are mainly position actuators and need a command signal proportional to desired eye position. Attempts to explain how this transformation might be achieved by, for example, proprioceptive rate feedback, have been excluded, leaving us with the conclusion that there exists a neural network in the caudal pons that performs integration in the sense of Newtonian calculus to convert a neurally encoded velocity command to a position command (Robinson 1989). This network has been localized in the region of the nuclei prepositus hypoglossi and medial vestibu-

Phil. Trans. R. Soc. Lond. B (1992) **337,** 327–330
Printed in Great Britain

lar nuclei (NPH-VN) where lesions devastate the integrator (Cannon & Robinson 1987).

2. INTEGRATING WITH LATERAL INHIBITION

How does a neural network integrate? The popular theory is by positive feedback. The basic idea is that if cells can excite themselves via signals relayed through neighboring cells, an idea called reverberating collaterals, then activity, once started in a neural pool, will perseverate itself in the absence of new input: another way of describing an integrator. This idea, however, must be applied cautiously here because all the signals we are concerned with are modulations of a firing rate around a background rate of roughly 100 spikes per second. We want to integrate the modulations but definitely not the background rates. Fortunately, a positive feedback scheme in which cells inhibit their neighbours and are indebted by them does the trick. Cells now excite themselves by disinhibition via their neighbours, and an intrinsic property of this arrangement, called lateral inhibition, is that only the modulations on an input, push–pull, signal pair is integrated (Cannon *et al.* 1983).

3. A NETWORK THAT LEARNS TO INTEGRATE

There is evidence that, not surprisingly, the real network is adaptive and can repair itself after a lesion. Moreover, the network integrates only poorly in a newborn infant and improves its performance, once vision is available, until it reaches normal behaviour in about two months (Weissman *et al.* 1989). This suggests that part of the network's connectivity is genetic but that its synaptic weights are trimmed up after birth to achieve the goal of holding the eyes steady so that the images of the world do not slide about on the retinas. Signals from direction-selective cells, in the retinas of rabbits and cats and the accessory optic systems of primates, transduce retinal image velocity, the error signal in the training of this neural network, and make it readily available in the brain stem.

Consequently, we thought it wise to carry on with the model of the vestibulo-ocular reflex by allowing it to learn through modifiable synapses (Arnold & Robinson 1991). In this model, a group of interconnected units (model neurons with membrane time constants of 5 ms) received a push–pull input from a pair of canals and projected to a push–pull pair of motoneurons to rotate the eye (similar to the scheme in figure 1 which is a later version of this model). The mechanics of the eye consists of viscosity and elasticity so, to create an eye movement $E(t)$, the motoneuron must receive the command $E(t)$ to overcome the elasticity and the command $\dot{E}(t)$, or eye velocity, to overcome the viscosity. Thus, the neural network must receive head velocity, $\dot{H}(t)$, and produce two signals, $E(t)$ and $\dot{E}(t)$ in the correct ratio.

Initially, the synaptic weights between the units are randomized. The sum of $\dot{H}$ and $\dot{E}$ is the velocity of the eye in space, so it is equal and opposite to retinal slip velocity which should be zero. This signal is integrated over a trial run and provides the error which is used to calculate how to change the synaptic weights so that the error can be reduced. We chose a learning algorithm that we call the screwdriver method: each weight in turn is changed by a small amount and its effect on the error is measured. The ratio is an estimate of the partial derivative of the error with respect to that weight, and the weights are then changed in proportion to that ratio. This process is repeated again and again until the error is sufficiently close to zero. This worked well and the network learned to perform integration and move the eye in an accurate compensatory fashion.

The network could recover from lesions and even from a hemilabyrinthectomy. Its units carried a variety of combinations of the eye position and

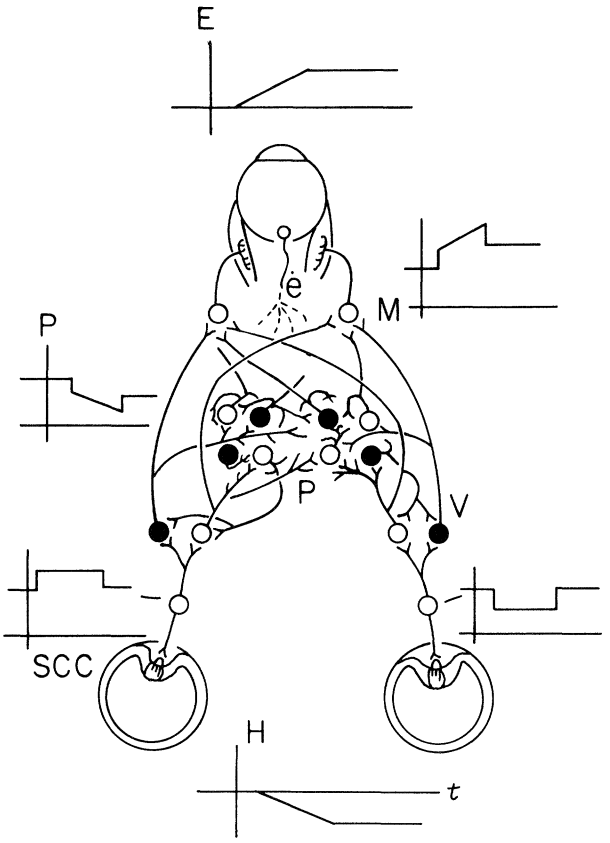

Figure 1. A test trial is a brief movement of the head, *H*. This creates a pair of push–pull signals from the semicircular canals, SCC, proportional to head velocity that are sent to excitatory (open cells) and inhibitory (filled) cells, V, in the vestibular nuclei. These project to motoneurons, M, with fixed polarities, inhibition ipsilaterally, excitation contralaterally. Neurons P are in the NPH-VN complex. Half are inhibitory, half excitatory. V cells project to all P cells. P cells project to all other P cells and all V cells. P cells project to M cells with fixed polarity, excitation ipsilaterally, inhibition contralaterally. As shown by the inset waveforms, many P cells and all M cells carry a combination of the eye-position and eye-velocity signals needed to make eye movement, *E*, equal and opposite to *H*. Retinal image slip, $\dot{e}$, is detected in the retina and is available to all modifiable synapses.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

velocity signals just as do real neurons recorded in the NPH-VN complex. The system still used positive feedback as its method of integration but now the feedback was hidden in a multitude of excitatory and inhibitory recurrent pathways. Nevertheless, an inhibitory commissural system did emerge; cells on one side of the brain stem tended to inhibit those on the other side so that lateral inhibition appeared in the organization.

4. LOCAL LEARNING RULE

We were still not satisfied because the learning algorithm was hardly physiological. It also ignored any genetic prewiring certain to occur in such an ancient reflex and did not distinguish between second and third order cells. Thus, we wanted to put more anatomy and genetics into the network and find a Hebbian-like learning algorithm that could change weights based only on local synaptic information. The scheme, shown in figure 1, differs from the earlier model by distinguishing between 2nd order vestibular neurons, V, and 3rd order cells, P, in the NPH-VN complex. Also, Dale's law is obeyed: cells are either excitatory (open circles) or inhibitory (filled circles) but not both. More important, 2nd order V cells excite motoneurons, M, contralaterally and inhibit them ipsilaterally and 3rd order P cells do just the opposite. The synaptic strengths may vary but these enforced polarities reflect what we believe to be genetically determined. Thus, each unit 'knows' which way its activity will rotate the eye. Otherwise, V cells contact all P cells, and P cells contact each other, V cells and M cells.

The main change was to alter the learning algorithm. The basic idea is that if a cell knew which way its activity turned the eye, then it could tell from the direction of the retinal slip, $\dot{e}$, whether the eye was going too slowly or too fast in 'its' direction. If it was going too slowly, one would want to increase synaptic weights from any excitatory input with a firing rate that was faster than normal, and vice versa if it was an inhibitory input. If the firing rate from the jth cell was R_j and its background rate was R_{jo}, so that $R_j - R_{jo}$ is the presynaptic modulating signal, we made the change in synaptic weight, W_{ij}, from the jth cell to the ith cell according to

$$\Delta W_{ij} = (+, -)(\dot{e})(R_j - R_{jo}).$$

If, for example, the ith cell causes the eye to go to the left and retinal slip, $\dot{e}$, is also to the left, the eye should go faster. If $(R_j - R_{jo})$ is positive and the synapse is excitatory, then one chooses the + or − sign in the first term so that the right hand side is positive and the synaptic weight, W_{ij}, is increased. It is increased both in proportion to the magnitude of retinal slip and the amount of presynaptic activity.

Recall that the network must produce both $E(t)$ and $\dot{E}(t)$. We found that if all the cells have the same learning rule they cannot produce both signals. Since the job of the reflex is twofold – to hold the eye still when the head is not moving (low frequencies) and move the eye briskly at the correct speed when the head does move (high frequencies) – we divided the cells into two groups, one responding to high frequency (phasic) signals in $\dot{e}$ and R_j, the other to low frequency (tonic) signals. This strategy allowed the network to converge to a solution rapidly and robustly.

In the resultant network, the commissural inhibitory network still emerged, indicating that positive feedback via lateral inhibition was still the basic means of achieving integration. The V cells encoded both E and $\dot{E}$, as is well known experimentally, and the P cells tended to divide into tonic cells (the E signal) and phasic-tonic cells (both the E and $\dot{E}$ signals, also known as burst-tonic cells). These results are not dissimilar from those recorded in the monkey NPH-VN complex by McFarland & Fuchs (1992).

5. DISCUSSION

In conclusion, we have modified our model to use a Hebbian-like learning algorithm that uses only local synaptic information that, at the least, is not inherently unphysiological. Obviously, we cannot tell if this is the way the synapses do learn, but we felt it necessary to divorce our models from clearly unphysiological algorithms such as backpropagation and our screwdriver method. Our results do illustrate that local Hebbian-like rules can obtain the same results when they are combined with modest genetic assumptions.

A major problem now is how one tests this hypothesis. One way to do this is to inject neural transmitter agonists and antagonists into the real network and interpret the results in terms of simulations in the model network. Preliminary results suggest that the computer network can simulate our experimental results, not only for injections but for stimulation and electrolytic lesions as well, but these experiments are in progress and the results are, as yet, only preliminary.

We thank A. McCracken for assistance with the manuscript and C. Bridges for assistance with the figure. This work was supported by a National Eye Institute Research Grant EY00598 of the National Institutes of Health (D.A.R.), and a Howard Hughes Medical Institute predoctoral fellowship (D.B.A.).

REFERENCES

Arnold, D.B. & Robinson, D.A. 1991 A learning network model of the neural integrator of the oculomotor system. *Biol. Cybern.* **64**, 447–454.

Cannon, S.C. & Robinson, D.A. 1987 Loss of the neural integrator of the oculomotor system from brain stem lesions in monkey. *J. Neurophysiol.* **57**, 1383–1409.

Cannon, S.C., Robinson, D.A. & Shamma, S. 1983 A proposed neural network for the integrator of the oculomotor system. *Biol. Cybern.* **49**, 127–136.

McFarland, J.L. & Fuchs, A.F. 1992 Role of the nucleus prepositus hypoglossi and adjacent medial vestibular nucleus in the control of horizontal eye movements in behaving macaques. *J. Neurophysiol.* **68,** 319–322.

Robinson, D.A. 1989 Integrating with neurons. *A. Rev. Neurosci.* **12**, 33–45.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

Weissman, B.M., Di Scenna, A.O. & Leigh, R.J. 1989 Maturation of the vestibulo-ocular reflex in normal infants during the first 2 months of life. *Neurology* **39**, 534–538.

Discussion

J. G. Taylor (*Kings College Centre for Neural Networks, London, U.K.*). Has there been any attempt to discover the LTD (reduction of synaptic weights) aspect of Dr Robinson's learning law; this is particularly of interest with respect to the on-going and very active debate about LTD versus LTP (increase in synaptic weights) among neurophysiologists studying synaptic change and memory?

D. A. Robinson. One thing that is fairly obvious in thinking about synaptic modifications is that one cannot go on endlessly turning synaptic strengths down (LTD) or up (LTP) without running into the stops, and it is becoming apparent that both mechanisms must have a way to change synaptic strengths in both directions although this is an idea that is still largely ignored by experimentalists. At any rate, in the model I presented, we assumed synapses could be strengthened and weakened by presynaptic conditions without invoking any specific mechanisms such as LTD or LTP. Although Ito has been proposing LTD between parallel T-fibers and Purkinje cells, his findings are not exactly robust or easily reproduced by others. To my knowledge, nothing is known about synaptic learning on cells of the deep cerebellar nuclei, such as those in the vestibular nuclei, or on other cell types in the latter nucleus. There is also a problem that although the cerebellum may play an important role in motor learning in the mature nervous system, motor learning may be very widespread throughout the developing nervous system. I'm afraid that Professor Taylor's phrase 'active debate' just about sums it up. This should be an area of vigorous research and productivity in the next decade or two.

P. J. Simpson (*Department of Psychology, University of Surrey, U.K.*). What are the implications of Dr Robinson's analysis and model of oculomotor control mechanisms for the clinical management of the Ménière's condition? Could the episodic nature of the condition reflect a change on vestibular function followed by adaptive 'recalibration' which is then made invalid by a further change in vestibular function? If so, should the clinical management include a programme of active relearning or adjustment to exploit the plasticity (as well as the treatment of the affective component)?

D. A. Robinson. The difficulty for the nervous system in Ménière's condition is that motor learning, which, fortunately, stays with us throughout life, tries to rebalance the background discharge rates from the left and right vestibular periphery but fails because of the great temporal variability of the disorder. If a fixed imbalance occurs, the adaptive rebalancing mechanism can succeed, but the episodic nature of this problem prevents anything but a temporary mitigation at best. For this reason, vestibular exercises, which attempt to encourage recalibration, are not recommended for Ménière's disease.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES