Institute for Neural Computation

Computational Analysis Suggests a New Hypothesis for Motor Learning in the Vestibulo-Ocular Reflex

Stephen G. Lisberger and Terrence J. Sejnowski

February 24, 1992

Technical Report INC-9201

Technical Report Series

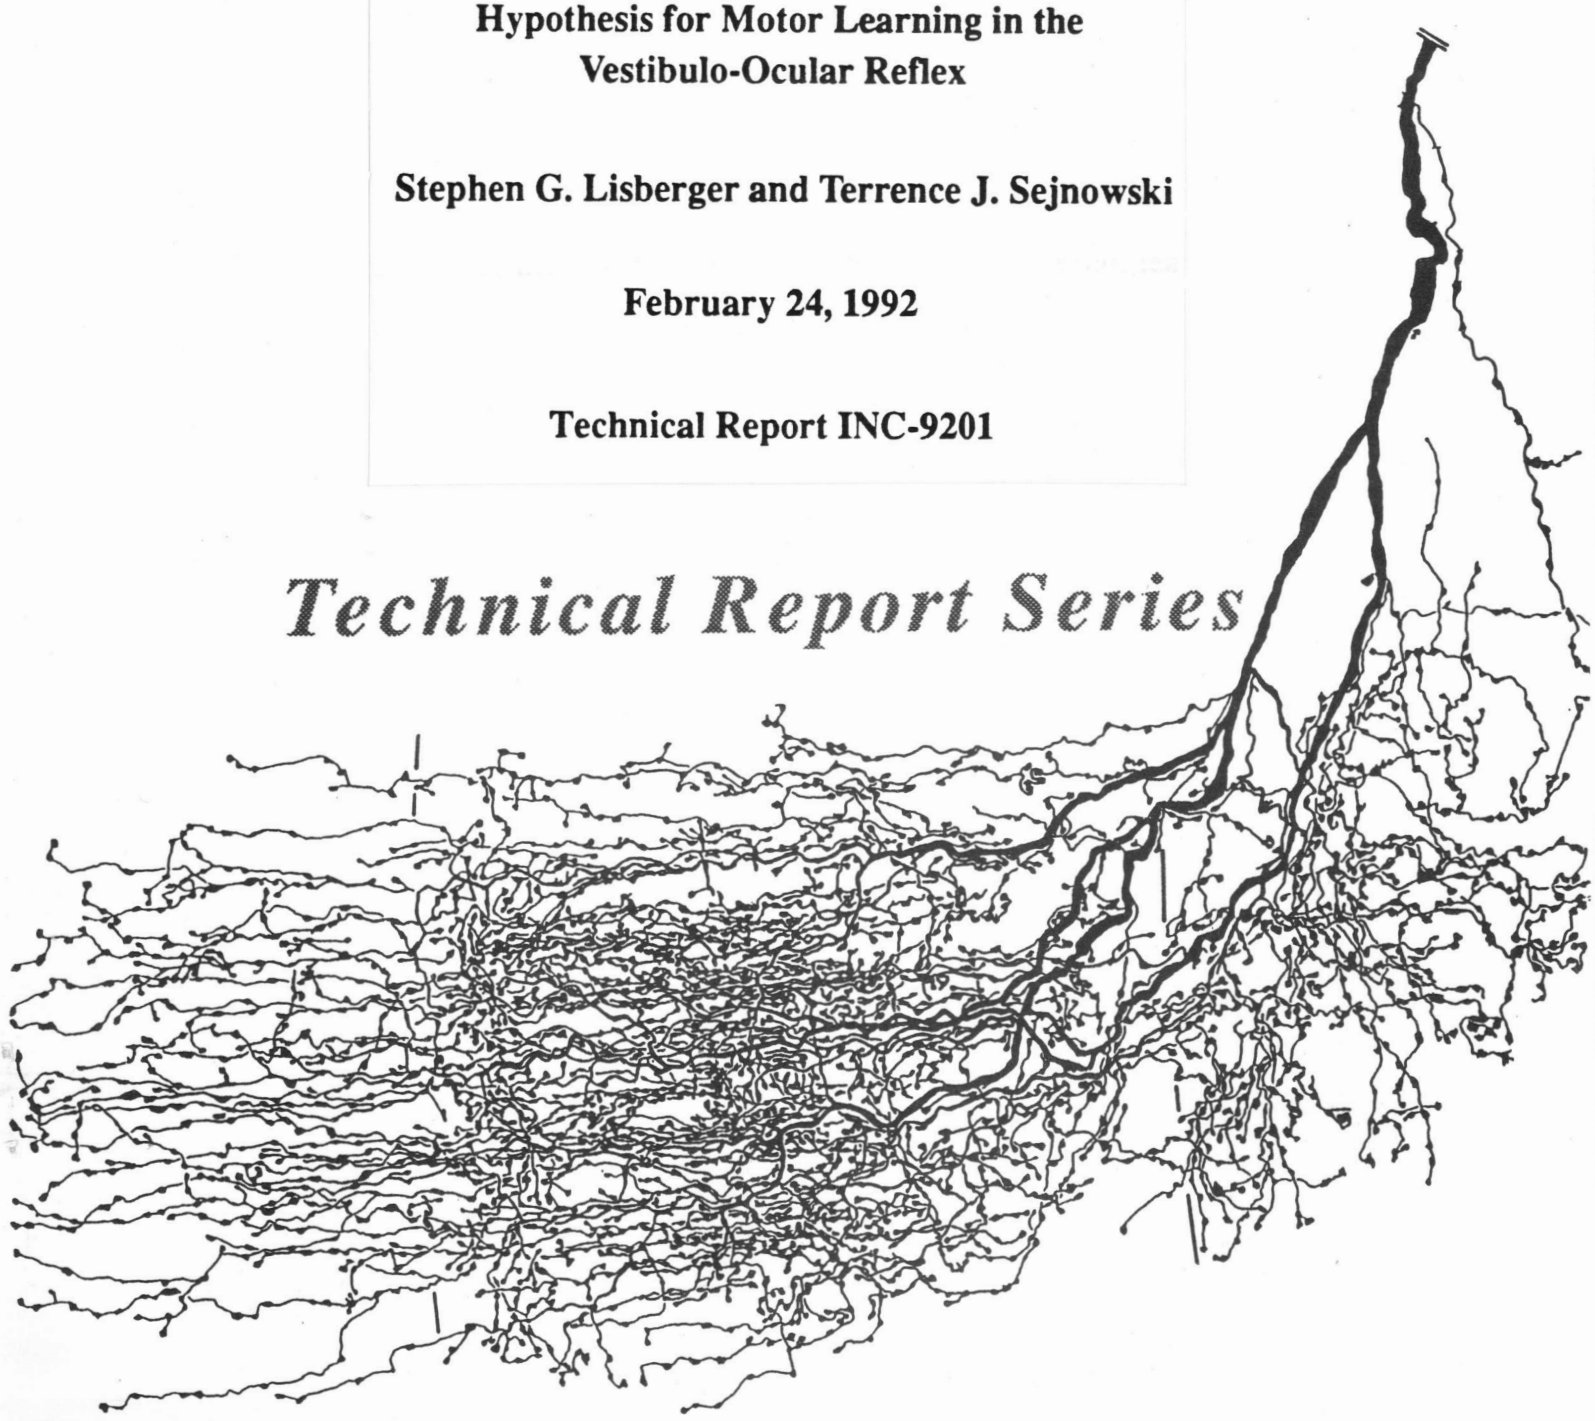

University of California, San Diego La Jolla, California 92093

Computational Analysis Suggests a New Hypothesis for Motor Learning in the Vestibulo-Ocular Reflex

Stephen G. Lisberger and Terrence J. Sejnowski

February 24, 1992

Technical Report INC-9201

Institute for Neural Computation
University of California, San Diego
9500 Gilman Drive
La Jolla, California 92093-0523

Stephen Lisberger is with the Department of Physiology, W.M. Keck Foundation Center for Integrative Neuroscience, and the Neuroscience Graduate Program at the University of California, San Francisco. Terrence Sejnowski is with the Salk Institute for Biological Science, Howard Hughes Medical Institute, and the Institute for Neural Computation, the Department of Biology, and McDonnell-Pew Center for Cognitive Neuroscience at the University of California, San Diego.
This research was supported by a grant for the Defense Advanced Research Project Agency, awarded through the Office of Naval Research. Correspondence concerning this article may be addressed to: Stephen Lisberger at the Department of Physiology, Box 0444, University of California San Francisco, San Francisco, CA 94143-0444 (email: sgl@phy.ucsf.edu).

Computational analysis suggests a new hypothesis for motor learning in the vestibulo-ocular reflex

S.G. Lisberger* and T.J. Sejnowski+

*Department of Physiology, W.M. Keck Foundation Center for Integrative Neuroscience, and Neuroscience Graduate Program, University of California, San Francisco

+Salk Institute for Biological Science, Howard Hughes Medical Institute, and Department of Biology and McDonnell-Pew Center for Cognitive Neuroscience, University of California, San Diego

Correspondence to:
Stephen G. Lisberger
Department of Physiology, Box 0444
UCSF
San Francisco, CA 94143-0444
phone: (415)-476-1062
FAX: (415)-476-4929
Email: sgl@phy.ucsf.edu

Studies of motor learning in the vestibulo-ocular reflex (VOR) of monkeys have identified brain cells that change their firing in association with learning. However, the site of learning has remained controversial[1,2]. A neural network model now shows how the sites of learning are constrained by physiological and behavioral data about the operation of the VOR. The model reveals that learning must be mediated by cellular changes in both the cerebellum and the brain stem and it suggests a new hypothesis in which changes in the dynamics of the signals in the brain cause modification of the amplitude of the VOR.

Good vision requires that images remain stationary on the retina in spite of movements of the subject or the target. The oculomotor system provides several mechanisms that use visual and vestibular sensory inputs to help maintain image stability. The *vestibulo-ocular reflex* (VOR) generates compensatory eye movements that maintain the stability of the eyes with respect to the stationary surroundings during head turns[3]. *Visual* tracking generates smooth eye movements that maintain eye speed equal to target speed, whether the target is a small object or a large moving field[4]. Experiments in awake monkeys have suggested that the brain stem and cerebellar pathways for these two kinds of movements are organized as shown in Fig. 1A. The direct VOR pathways include vestibular sensory inputs that encode angular head velocity, interneurons in the vestibular nucleus, and motor outputs through extraocular motoneurons[5]. A second VOR pathway goes through the flocculus and ventral paraflocculus of the cerebellum (hence called the flocculus), and acts as an inhibitory side loop of the direct VOR pathway[6]. The current command for smooth eye velocity is returned over a positive feedback pathway to the flocculus, where it forms part of the command for future smooth eye movements[4]. Visual inputs access the tracking system at least in the flocculus[7] and probably also in the brain stem.

If the VOR is functioning properly, then head turns will be associated with little motion of retinal images from the stationary surroundings. Whenever image motion occurs

consistently during head turns, the VOR undergoes motor learning to restore the proper compensatory eye movements[8]. In the laboratory, learning is induced by fitting monkeys with spectacles that magnify or miniaturize vision so that the VOR must become either larger or smaller than normal[9]. The performance of the reflex is estimated before and after learning by calculating the *gain of the VOR*, defined as eye speed divided by head speed during passive angular head rotation in darkness.

The organization of the VOR pathways has suggested two hypotheses for the site and mechanism of motor learning. Ito has postulated that the gain of the VOR is reduced by *increasing* the transmission of vestibular signals through the flocculus[10] (site **a**), while Miles and Lisberger have proposed that the gain of the VOR is reduced by *decreasing* the transmission in parallel through both the brain stem and the flocculus[11] (sites **a** and **d**). Miles *et al*[12] developed the model in Fig. 1B as a quantitative way to represent the organization illustrated in Fig. 1A. Analytical solution of their model reveals that the gain of the VOR in the steady-state is (d-e*a)/(1-e*b) and demonstrates that both hypotheses are feasible. In support of Ito's hypothesis, increases in the value of **a** would cause decreases in the gain of the VOR in the model. In support of Miles and Lisberger's hypothesis, appropriate decreases in the values of **a** and **d** would also decrease the gain of the VOR in the model. Miles *et al*[12] favored the latter alternative because their data showed that the value of **a** decreases when the gain of the VOR is decreased in monkeys.

The model in Fig. 1B, although instructive, is inadequate for evaluating how motor learning occurs in the brain. The brain subjects its inputs to dynamic transformations and processes inputs and outputs that vary as a function of time. The model, in contrast, lacks dynamics and predicts only the steady-state eye velocity. Most previous experiments on motor learning in the VOR have employed low-frequency sinusoidal head rotation and have not required dynamic models. In addition, previous models have assumed that the cellular mechanisms of learning in the VOR could be represented by multiplication factors like the elements in Fig. 1B.

Recent experiments have provided a new class of data that requires new models. The use of transient head motions has provided information about the performance of the VOR and about the responses of brain cells on a millisecond time scale. Interpretation of these data requires dynamic models that process time-varying inputs and generate time-varying outputs. We have now analyzed the performance of a dynamic model to determine how learning can occur in the VOR and to reevaluate the hypotheses of Ito[2] and of Miles and Lisberger[11] in the light of these new data. We will show that both hypotheses are incomplete. We will also suggest a novel mechanism of learning that can cause appropriate changes in the gain of the VOR without contradicting available data.

Structure of the network model

We have designed a network model (Fig. 1C) with the following constraints: the output of the model had to be correct when the sensory inputs and the eye movement outputs varied as a function of time; the structure the model had to account for the known flow of signals in the VOR pathways; and the model had to perform realistic visual as well as vestibular tracking. The model consisted of artificial processing units, each with an adjustable time constant of leaky integration and with an output that was a sigmoid function of its summed inputs. The units were grouped into clusters and the clusters were interconnected in a pattern that is consistent with Fig. 1A. Cluster F represents the flocculus, cluster V represents brain stem cells that receive monosynaptic inhibition from the flocculus[13], cluster M represents extraocular motoneurons, and cluster E consists of an eye velocity output unit. The performance of the model did not depend on the number of units in each cluster; Fig. 1C shows the number actually used. Each arrow in Fig. 1C indicates that all the units of one cluster were connected to all the units of another cluster. Each connection between two units was defined by an adjustable weight and a fixed time delay.

The model network employes sensory input to guide tracking continuously. Summing junction "R" compares target and eye motion with respect to the head and yields

a dynamically-varying sensory signal that is related to the velocity of the visual images. To mimic the latency of tracking, we placed a fixed delay of 100 ms in the visual inputs to cluster 10. Vestibular inputs related to head velocity and head acceleration allow the model to fixate stationary targets in spite of head motion; to mimic the latency of the VOR, we placed a fixed delay of 14 ms in the outputs from cluster 1. The simulation environment provided a gradient descent optimization algorithm that we used to find combinations of parameters that minimized output error[14].

The structure of the model in Fig. 1C was selected to account for the known flow of signals in the brain stem and cerebellar pathways that guide pursuit and the VOR. To show the homology between our model and the static model of Miles *et al*[12], the relevant connections in Fig. 1C have been labeled with lower case letters that correspond to the letters used as multiplication factors in Fig. 1B. We will refer to the connections by these letters. The model in Fig. 1C was not intended to be a literal representation of the anatomy of the VOR pathways in the brain, which is not entirely known. Some of the units, for example those in clusters 2 and 3, were required because of technical considerations related to the simulations. We do not think that these deviations from anatomical realism pose a problem because the issues addressed by the present paper deal with the flow of information and not with the exact synaptic organization of the VOR pathways.

Organization of the normal tracking model

We first fixed the weights in connections **a** and **b** at zero and allowed the optimization algorithm to adjust other parameters until the network generated an accurate VOR during rapid changes in head velocity in the dark (Fig. 2). Cluster 1 included four vestibular input units that were preprogrammed to have reciprocal responses; two were excited and two were inhibited by each head turn. Of each pair of input units, one provided a tonic input related only to head velocity while the other provided a phasic-tonic input related to head acceleration and head velocity[15]. The gradient descent optimization algorithm then modified the weights in the brain stem VOR pathways (connections 1->2,

d, 3->V, and V->M) until the actual eye velocity was nearly equal to and opposite head velocity and the gain of the VOR was close to 1. This resulted in the eye velocity and unit responses shown in Fig. 2. The model used both the tonic and the phasic-tonic vestibular inputs, so that the phasic behavior in the responses of the units in cluster V was intermediate between that of these two kinds of inputs (Fig. 2).

We next trained the model to track a step change in target velocity from 0 to 30^{o}/s. Through trial and error, we established the minimum number of parameters that needed to be adjusted by the optimization algorithm to achieve the excellent tracking shown in Fig. 3A. Good performance, defined as a rapid increase in output followed by stable tracking, was achieved only when the eye velocity pathway through cluster 7 was configured in positive feedback. If positive feedback of eye velocity was prevented by fixing the weights in connection **b** at zero, then the optimization algorithm could not find a combination of parameters that allowed the model to achieve a rapid rising phase in simulated eye velocity without developing severe oscillations like those shown in Fig. 3B. Stable tracking also required adjustment of the time constants in cluster 7: in four different simulations with fixed visual feedback delays of 40, 70, 100, and 130 ms, the optimization algorithm adjusted the time constants to be 44, 63, 78, and 89 ms, respectively.

The final step in the development of the normal tracking model was to free enough additional parameters for adjustment so that a single configuration could perform the VOR (Fig. 3C), visual tracking with the head fixed (Fig. 3A), and an additional tracking task called cancellation of the VOR in which the head and target moved together (not shown)[16]. To accomplish these tasks, the model required the vestibular pathway through clusters 9 and F, which became configured as an inhibitory side-loop of the VOR, mimicking the known anatomy of the vestibular pathways through the flocculus[6]. It also required that the time constants in cluster 9 were similar to those in cluster 7 and that the vestibular inputs to cluster F arose from the tonic units in cluster 1. Even if phasic-tonic inputs were available, the weights in their connections to cluster F were set to zero by the

optimization algorithm. In agreement with the physiology of the flocculus in the primate, the units in cluster F showed approximately equal sensitivity to head velocity and eye velocity inputs[17,18] and their output was almost unmodulated during the normal VOR[18,19].

The positive feedback configuration of the eye velocity input to cluster F is required in our model to produce realistic tracking of smooth target motion with the head stationary (compare Fig. 3A and B). Positive feedback would not be required in a static model like that in Fig. 1B, where it could be replaced with an amplifier. In a dynamic model, positive feedback acts as an integrator and affects the time course as well as the amplitude of the output from the model. In our model, positive feedback integrates the transient visual motion input caused by target motion and allows the eye velocity of the model to be sustained at a value close to target velocity during visual tracking. However, the same integration would cause a ramp increase in eye velocity for the sustained head velocity input that drives the VOR. In the model, the head velocity input to cluster F counteracts the eye velocity input so that the positive feedback loop is prevented from acting as an integrator during the VOR. The next section will show that the need to preserve balanced head and eye velocity inputs to cluster F during the VOR is a critical constraint on possible sites of motor learning in a dynamic model.

Sites of learning in the model

We conducted a number of simulations, each of which began after the parameters of the model had been optimized to produce realistic visual tracking and a normal VOR. The parameters of the model were readjusted by the gradient descent optimization algorithm under conditions that required motor learning in the VOR. To model the effect of wearing miniaturizing spectacles, eye velocity was required to be opposite in direction and 25% of the amplitude of head velocity. We began by allowing the optimization algorithm to readjust the weights at connections **a**, **b**, **d**, and **e** and using a training set that included both visual tracking and the VOR. After learning (Fig. 3D), simulated eye velocity during the

VOR (solid line) was close to the desired output (dashed line) and the model's performance during the other tracking tasks (not shown) was the same as in the original tracking model (Fig. 3A). Comparison of Fig. 3C and D shows that the change in the gain of the VOR was driven by changes in the output of units in cluster V with little change in the output from cluster F.

We next subjected the normal tracking model to the same training conditions used above to reduce the gain of the VOR to 0.25, but now with the adjustment of weights restricted to different combinations of sites. Table I shows that the model did not achieve stable performance when weights could be adjusted only at connection **a** (Model V, Fig. 3E) or only at connection **d** (Model III): constant head velocity did not elicit constant eye velocity and the asymptotic value of output error was large. Only if the optimization algorithm was allowed to adjust the weights at connections **a** and **d** was the model able to achieve a small asymptotic value of output error. Eye velocity positive feedback necessitates multiple sites of learning. Without parallel modification of the weights at connections **a** and **d**, any changes in these weights will upset the balance between vestibular and eye velocity inputs to cluster F during the VOR; the integrating action of eye velocity positive feedback will convert the imbalance into a ramp of eye velocity.

The presence or absence of training tasks that required visual tracking dictated how the model learned a VOR with a reduced gain. When the model was required to perform both the VOR and visual tracking but was allowed to change only the values of the weights in connections **a** and **b**, it was not able to learn a stable reduced VOR (Table I, Model VI). When the same weights were adjustable but visual tracking was not required, the optimization algorithm was able to achieve a small asymptotic output error in a VOR with a gain of 0.25. Motor learning in the VOR was then accomplished by parallel reductions in the weights in connections **a** and **b**, which compromised the performance of the model on visual tracking tasks. In monkeys, by contrast, motor learning in the VOR causes only tiny

changes in smooth visual tracking and no change in the positive feedback of eye velocity through the flocculus[20].

We next asked whether the magnitude of the changes in weights observed in the model were compatible with existing data from monkeys. In experiments on monkeys, Miles et al[12] have made measurements that should be comparable to the values of the weights in connection **a**. They found that the sensitivity to vestibular inputs of Purkinje cells in the flocculus was reduced to 67% of normal when the gain of the VOR was reduced to 0.18. Table I shows that the weights in connection **a** of the model were reduced to 39-44% of their original values when the gain of the VOR was reduced to 0.25. Further analysis of model II revealed that the weights in connection **a** had to be reduced to 31% of normal to reduce the gain of the VOR to 0.18. Reducing the weights in connection **a** to the 67% of normal found by Miles *et al* allowed the gain of the VOR in the model to be reduced only to 0.57. Thus, our model produces a stable VOR with a reduced gain when allowed to alter the weights in connections **a** and **d**, but it cannot make a given change in the gain of the VOR without changes in weight at connection **a** that exceed those recorded in monkeys.

A new way to implement learning in the VOR

Until now, models of the motor learning in the VOR have represented the cellular mechanism of learning as an multiplication factor. The optimization algorithm used for our simulations discovered that it is also possible to alter the gain of the VOR by changing the dynamics of the signals in the model. Fig. 4 shows how this mechanism worked when the optimization algorithm was allowed to adjust only the weights at connection **a**, a constraint that prevented learning of a stable VOR when the dynamics of the signals in the model were fixed. During the normal VOR, before learning, eye velocity was almost equal in amplitude to head velocity (Fig. 4A). Fig. 4C shows that the output from cluster F underwent very little modulation during the VOR, because the inputs to cluster F through connection **a** (labeled "Head velocity input" in Fig. 4C) and connection **b** ("Eye velocity

input") cancelled each other: they were equal in amplitude but opposite in direction. In this version of the model, the optimization algorithm was allowed to adjust the weights in the inputs from both the tonic and the phasic-tonic vestibular afferents to cluster F. Fig. 4E shows that the optimization algorithm adjusted these weights so that the vestibular inputs to cluster F arose only from the tonic afferents (Fig. 4E) when the gain of the VOR was 1.0.

Fig. 4B shows that this configuration of the model was able to learn a stable VOR with a reduced gain of 0.25. The dashed line labeled "H" in Fig. 4D shows that the head velocity input to cluster F became more transient when the gain of the VOR was 0.25 but retained the steady-state amplitude it had when the gain of the VOR was normal. Because the eye velocity of the VOR was reduced, the eye velocity input to cluster F ("E" in Fig. 4D) was smaller than normal and the output from cluster F now underwent strong modulation during the VOR. Finally, Fig. 4F shows that the change in the transient behavior of the head velocity input to cluster F reflected an increase in the contribution of the phasic-tonic vestibular inputs and a decrease in that of the tonic inputs.

The integrating action of the eye velocity positive feedback pathway makes it possible for the mechanism outlined in Fig. 4 to achieve a change in the gain of the VOR with little or no change in the steady-state amplitude of the vestibular inputs to the positive feedback pathway. The addition of a phasic component in the input to the positive feedback loop injects a transient signal into the loop that is integrated and remembered. If, as illustrated in Fig. 4, the transient is injected at cluster F, it causes a reduction in the gain of the VOR. Changes in the dynamics of the head velocity input to cluster V would also cause changes in the gain of the VOR, but would contradict other data[21].

Implications for the site of motor learning in the brain

Our simulations demonstrate that both extant hypotheses for motor learning in the VOR are incomplete. Ito has proposed that the site of learning is in the cerebellar cortex and that increases in the cellular equivalent of synaptic weights would decrease the gain of the VOR[10]. We have discovered that his hypothesis, when tested in a dynamic simulation,

does not provide a stable VOR *and* excellent visual tracking after motor learning. Miles *et al* [12] and Lisberger[1] proposed that parallel decreases in the cellular equivalent of weights in the brain stem and the cerebellum would work together to decrease the amplitude of the VOR. We have found that this hypothesis provides a stable VOR after learning, but that it fails to agree quantitatively with the following data. 1) In the model, the steady-state vestibular sensitivity in cluster F is reduced to 31% of normal when the gain of the VOR is reduced to 0.18. In the monkey, the same reduction in the gain of the VOR is associated with a reduction in the vestibular sensitivity of floccular Purkinje cells only to 67% of normal[12]. 2) In the model, learning occurs without changes in the output of cluster F (compare Fig. 3D and E). In the monkey, learning is associated with large changes in the responses of floccular Purkinje cells during the VOR[12,13,22].

We now propose a new hypothesis for motor learning in the VOR. The hypothesis allows stable performance of the VOR on a millisecond time scale when its inputs and outputs vary as a function of time. It also permits the model to maintain excellent visual tracking after motor learning in the VOR and it allows changes in weight at connection a that are consistent with the data of Miles et al[12]. The new hypothesis incorporates elements of those proposed by Ito[10] and by Miles *et al*[12] and takes advantage of our discovery that the gain of the VOR can be altered by changing the dynamics of the signals in our model. We present the hypothesis by showing how the sequential application of two separate mechanisms can mimic the data of Miles et al[12], but we assume that the mechanisms operate simultaneously in the brain. First, parallel decreases in the vestibular transmission through the brain stem and the flocculus would cause decreases in the gain of the VOR and would allow continued stable performance. The data of Miles *et al* require that the weights at connection a be reduced to 67% of normal, which means that this mechanism can reduce the gain of the VOR only to 0.57. Second, simultaneous increases in the weights for the phasic-tonic vestibular inputs to the flocculus and decreases in the weights for the tonic vestibular inputs would reduce the gain of the VOR to 0.18. After this

change in the dynamics of the vestibular inputs to the flocculus, steady-state vestibular sensitivity in the flocculus would still be 67% of normal and the output of the flocculus would become modulated during the VOR, in agreement with available data[12,13,22].

Fig. 5 outlines the constraints that determine possible sites of motor learning in the VOR and shows how those constraints would be affected by other, hypothetical VOR pathways. We have found that positive feedback of eye velocity is essential for realistic pursuit, but that it causes the model to become unstable unless there is balance in the strength of vestibular transmission through the brain stem (**d**) and the cerebellum (**a1** and **a2**). The strong evidence that such a positive feedback pathway exists in the flocculus[4], along with the fact that decreases in the gain of the VOR are associated with decreases in the steady-state transmission of vestibular inputs through the flocculus, argue that the hypotheses of Miles et al[12] and Lisberger[1] must form part of the mechanism of motor learning. Decreases in the gain of the VOR must be accomplished in part by parallel decreases in the value of **d** and in the steady-state component of **a1** and **a2**. This conclusion would not be changed by including the well-known neural integrator in the brain stem and a more realistic description of the orbit and eyeball[23], or by adding a parallel unmodified VOR pathway in the brain stem[24].

The dashed lines in Fig. 5 illustrate how the constraints on the site of learning might be different in other configurations of the model. The pathway through **a3** represents any other vestibular pathway that *both* drives eye movement *and* contributes to the positive feedback of eye velocity. This pathway, which could include floccular Purkinje cells that do not respond to eye velocity, operates under the same constraints as the pathways through **a1**, **a2**, and **d**. In fact, the inclusion of **a3** would not alter the operation of the model at all because changes in the value of **a3** would have the same effects as changes in **d**. The pathway through element **c** in Fig. 5 could supplement the mechanisms of motor learning proposed here. This pathway contains a site of learning and contributes to the signals that drive eye movements, but it does not provide inputs to the positive feedback through the

flocculus. If such a pathway existed in the brain[25], changes in its vestibular transmission (represented by the value of **c** in Fig. 5) would alter the gain of the VOR without affecting the stability of the VOR, the performance of visual tracking, or the discharge of the brain cells represented by clusters F and V in our model. The addition of other connections and pathways in Fig. 5 would not alter our conclusions because any other VOR pathway would be mathematically identical to one of the pathways that is already shown in this diagram.

Conclusions

Our analysis of a network model has pointed out the importance of considering motor learning in the VOR as a dynamical property of a system of neurons. We have found that earlier hypotheses based on individual sites of learning in the brain stem or cerebellum fail to account for the known behavior of the VOR. Even qualitative agreement between the performance of the model system and that of the biological system could be attained only by postulating highly-specific changes in vestibular transmission at a few carefully-chosen sites.

We have introduced a new hypothesis that uses a shift in the balance of phasic-tonic and tonic inputs to the cerebellum to contribute to motor learning. The new mechanism is a specific implementation of a more general mechanism in which changes in dynamic processing inside a model can cause changes in the amplitude of its output. The power of this mechanism in our model raises the possibility that subtle changes in the cellular properties of neurons could have major effects on the properties of a system of neurons. Thus, brain mechanisms as subtle as a change in adaptation properties of neuronal firing rate or a modification of the dynamics of individual ionic conductances could mediate behavioral learning by causing large changes in the operation of interconnected networks of neurons.

References

1. S.G. Lisberger, *Science* **242**: 728 (1988).
2. M. Ito, *Ann. Rev. Neurosci.* **5**: 275 (1982).
3. B. Cohen, In: *Handbook of Sensory Physiology, Pt. 1: Vestibular system.* ed: H.H. Kornhuber, **6**: 477 (1974).
4. S.G. Lisberger, E.J. Morris, and L. Tychsen, *Ann. Rev. Neurosci.* **10**: 97 (1987). Because the basic function of the brain stem and cerebellar pathways seems to be the same for pursuit of a small target and the fast component of the smooth eye movements evoked by the motion of a large field, our simulations did not differentiate between these two visual tracking mechanisms.
5. W. Precht and R. Baker, *Exp. Brain Res.* **14**: 158 (1972).
6. R. Baker, W. Precht, and R. Llinas, *Exp. Brain Res.* **14**: 364 (1972). The region we have been calling the flocculus in the primate actually includes both the ventral paraflocculus and the flocculus (N.M. Gerrits and J. Voogd, *Exp. Brain Res. Suppl.* **17**: 26 (1989)). For consistency with the existing literature on the monkey and ease of expression, we will refer to the whole structure as the flocculus, pending definitive evidence that its two embryological parts have different functions.
7. L.S. Stone and S.G. Lisberger, *J. Neurophysiol.* **63**: 1241 (1990).
8. A. Gonshor and G. Melvill Jones, *J. Physiol. (London)* **256**: 381 (1976).
9. F.A. Miles and B.B. Eighmy, *J. Neurophysiol.* **43**: 1406 (1980).
10. M. Ito, *Brain Research* **40**: 81 (1972).
11. F.A. Miles and S.G. Lisberger, *Ann. Rev. Neurosci.* **4**: 273 (1981).
12. F.A. Miles, D.J. Braitman, and B.M. Dow, *J. Neurophysiol.* **43**: 1477 (1980).
13. S.G. Lisberger and T.A. Pavelko, *Science* **242**: 771 (1988).
14. The units in the model were leaky integrators whose output was a sigmoid function [$y(x) = 1/(1+e^{-x})$] of their summed inputs. The model was run iteratively with two phases in each iteration. First, the model was provided with inputs of duration 1 s and

Euler integration was used to calculate the output of each unit and the output of the model in 500 time steps at 2 ms intervals. The desired output was provided for each input so that output error could be calculated as the sum of the sum of the squares of the difference between desired and actual output over the interval from the end of the delay in the input pathways to the end of the simulation. In this paper, summed error is reported without units; because the error was always calculated over the same interval for each type of behavioral trial, the numbers we report here are internally comparable. Second, each parameter in the model was adjusted in proportion to a gradient that reflected the partial derivative of output error with respect to that parameter (D.E. Rumelhart, G.E. Hinton, G.E., and R. Williams, *Nature* **323**: 533 (1986)), summed over the duration of the input stimulus. The gradients for each time step were computed using a variation of Pearlmutter's procedure for dynamic, recurrent back propagation (B.A. Pearlmutter, *Neural Computation* **1**: 263 (1989); S.R. Lockery, Y. Fang, and T.J. Sejnowski, *Neural Computation* **2**: 274 (1990)). With this algorithm, both the weights of the connections and the time constants of integration in the units can be adjusted. The model was started with random weights that varied between -0.2 and 0.2 for most of the adjustable weights and between -0.4 and zero for the output weights from cluster F, which were constrained to be negative to be consistent with the inhibitory action of cerebellar Purkinje cells.

15. S.G. Lisberger and T.A. Pavelko, *J. Neurosci.* **6**: 346 (1986) showed that both tonic and phasic-tonic trajectories of firing rate are found in the responses of primary afferents to rapid changes in head velocity.

16. The simulations reported here included cancellation of the VOR as a tracking condition, but other simulations revealed that neither the structure of the normal tracking model nor the site of learning were affected if the tracking conditions included only the VOR in the dark and visual tracking with the head stationary.

17. S.G. Lisberger S.G. and A.F. Fuchs, *J. Neurophysiol.* **41**: 733 (1978).

18. F.A. Miles, J.H. Fuller, D.J. Braitman, and B.M. Dow, *J. Neurophysiol.* **43**: 1437 (1980).

19. S.G. Lisberger and A.F. Fuchs, *Brain Research*. **69**: 347 (1974).

20. S.G. Lisberger, F.A. Miles, L.M. Optican, and B.B. Eighmy, *J. Neurophysiol.* **45**: 869 (1981) showed that motor learning in the VOR does not affect the fast component of visual tracking of a large stimulus. One of us has demonstrated recently that motor learning in the VOR causes only very small changes in the eye acceleration in the first 100 ms of the initiation of pursuit for a small target (S.G. Lisberger and J. Schwartz, unpublished observations). Miles *et al.*[12] found no evidence that the amplitude of the eye velocity input to the flocculus depended on the gain of the VOR. Lisberger[1] showed that motor learning does not alter the eye movements evoked by electrical stimulation of the flocculus.

21. Simulations of a linear control-theory model of the VOR (S.G. Lisberger, unpublished observations) have demonstrated that this mechanism cannot operate at the vestibular inputs to cluster V, because it contradicts the fact that decreases in the gain of the VOR cause an increase in the phasic overshoot in eye velocity during rapid changes in head velocity[15].

22. E. Watanabe, *Brain Res.* **297**: 169 (1984).

23. Our model operates in the velocity domain. Including a neural integrator in the brain stem (A.A. Skavenski and D.A. Robinson, *J. Neurophysiol.* **36**: 724 (1973)) and a realistic model of the orbit and eyeball would allow the output of the model to be eye position rather than eye velocity. These modifications would not affect the performance of the system because the elaboration of the brain stem is designed to cancel the more rigorous description of the orbit. Motor learning could not occur in the integrator without contradicting existing data by linking changes in the gain of the VOR to parallel changes in all other kinds of eye movements.

24. The presence of an unmodified brain stem pathway (S.G. Lisberger, *Science* **225**: 74, (1984)) would require larger changes in the modified pathways to achieve a given reduction in the amplitude of the VOR, but would not alter the conditions needed to achieve stability and would not improve the performance of either of the extant hypotheses for motor learning in the VOR.
25. Recordings from the ventral paraflocculus have implied that the eye velocity feedback signal represents the full command for smooth eye movement of visual and vestibular origin, making it unlikely that such a separate, modified VOR pathway exists. However, recent anatomical studies (N.M. Gerrits and J. Voogd, *Exp. Brain Res. Suppl.* **17**: 26 (1989)) have identified the caudal portion of the flocculus, from which few recordings have been made in monkeys, as a region that may be involved in motor learning in a way that is not included in our hypothesis. If the output from the caudal flocculus ultimately provides part of the eye velocity feedback to the rostral flocculus and ventral paraflocculus, then this structure would have to function under the same constraints described here for the units in cluster F. If the caudal flocculus possesses a separate and private output pathway to motoneurons that does not access the eye velocity positive feedback, then these caudal Purkinje cells could still contribute to motor learning in the simpler way envisaged in early models of the VOR[10,11].
26. Research supported by a grant from the Defense Advanced Research Project Agency, awarded through the Office of Naval Research. We thank numerous colleagues for their helpful comments on an earlier version of the manuscript.

Model	I	II	III	IV	V	VI
Connection						
e	99%	--	--	--	--	--
b	103%	--	--	96%	--	92%
a	39%	44%	--	--	107%	102%
d	28%	35%	81%	84%	--	--
error	**0.008**	**0.011**	**0.353**	**0.348**	**0.529**	**0.444**

Table I. Changes in connection weight during learning of a reduced amplitude VOR with adjustment allowed in different combinations of connections. For each modifiable set of connections between clusters, we summed the absolute values of the weights before and after learning. Each entry in the table shows the summed final connection weight as a percentage of the summed initial weight. Entries of -- indicate that the model was configured to prevent modifications in that connection. The bottom row shows the asymptotic value of output error in a single VOR trial. The asymptotic output error did not become smaller if models III-VI were altered so that the optimization algorithm was allowed to adjust the weights in connections **e** or **b**. We did not allow the optimization algorithm to adjust the weights in connection V->M because of evidence that this connection is not modified in the brain[1]. Each configuration of the model was run with the same simulation parameters[14]: decay, 10^{-7}; momentum, 0.5; learning rate for weights, 2.0; learning rate for time constants, 0.1. The performance of the model was similar in attempts to increase the amplitude of the VOR, although the non-linearities in the input-output characteristics of the processing units and the method for calculating output error yielded larger values for asymptotic output error.

FIGURE LEGENDS

Figure 1. A network model of the vestibulo-ocular reflex (VOR) that is based on the current understanding of the anatomy and physiology of the VOR pathways in the brain stem and cerebellum. **A**: Block diagram showing the basic organization of the brain pathways. **B**: A static model of the brain stem and cerebellar pathways that subserve visual tracking and the VOR in monkeys, after Miles *et al*[18]. The arrows indicate the flow of signals in the model, from the head velocity input on the left to the eye velocity output on the right. Circles represent summing junctions and the boxes containing small letters represent gain elements that multiply their inputs by the value of the variable inside the box. **C**: A network model based on the biological network. Each group of units represents a separate cluster and the arrows indicate the direction of specific interconnections between clusters. Cluster F, representing the flocculus, inhibits cluster V, representing the cells in the brain stem that receive inhibition from the flocculus. Cluster 1 contains the input units, which signal head velocity and head acceleration, cluster M represents two extraocular motoneurons, and cluster E consists of an eye velocity output unit, needed to compare the actual output of the model with the desired output. The letters next to some of the arrows indicate the analogy among the connections where modifications were allowed when testing the site of learning in the model, the multipliers defined in the static configuration of B, and the postulated sites of learning in A. The circle with an "R" in it is a summing junction that represents the retina and provides an output proportional to the different between target velocity and eye velocity, both of which were represented relative to the potentially-moving head. The model in panel C lacks an explicit three-neuron arc, but includes one implicitly in the chain of connections through clusters 1, 2, 3, and V to cluster M. Cluster 2 was included so that the balance of tonic and phasic-tonic vestibular inputs to the brain stem VOR pathways could be fixed by preventing the optimization algorithm from adjusting connection 1->2. Cluster 3 was included as a buffer to ensure

that the sites for learning in the two head velocity pathways to cluster V required back propagation from the output unit over the same number of connections. The number of connections in the pathway does not have a significant effect on the operation of the kind of artificial network studied here and the number of connections in a model of this kind bears no relation to the number of synaptic relays used by an equivalent circuit in the brain.

Figure 2. Combined use of tonic and phasic-tonic vestibular inputs to generate eye velocity during the VOR. The input to the model was a 50-ms ramp of head velocity from 0 to -30^{o}/s followed by a constant head velocity of 30^{o}/s. The tonic and phasic-tonic vestibular input units in cluster 1 provide inputs that are combined in cluster V to generate commands for eye velocity. The dashed line shows the desired eye velocity during the VOR.

Figure 3. Performance of the normal tracking model for visual and vestibular stimuli. **A:** Tracking of target motion at 30^{o}/s with the head stationary. There was a fixed delay of 100 ms in the visual inputs and the optimization algorithm had adjusted the weights in connection **b** as well as the time constants of integration in cluster 7. **B:** Optimal performance for the same stimulus shown in a, but with the weights in connection **b** constrained to be zero. **C-E:** Performance of the model and output from selected units during the VOR before (**C**) and after (**D,E**) motor learning to reduce the gain of the VOR to 25% of normal. In ^{D}, connections **a, b, d,** and **e** were modifiable. In **E**, the model had been optimized to have a reduced VOR with modifications allowed only in connection **a**. Each trace shows amplitude as a function of time in 2 ms steps. The dashed line in the bottom row of traces shows the desired eye velocity with respect to the orbit. The value at zero time in each trace shows zero eye or head velocity, or a unit output with a value of 0.5. The calibration bar holds for the eye and head velocity traces. The traces showing the outputs from cluster F and V have the same amplitude scale. Within each cluster, the

outputs from all the units were qualitatively similar; we selected the unit that had the largest amplitude response to illustrate the output from each cluster. .

Figure 4. Illustration of how changes in connection **a** can cause motor learning when the optimization algorithm is allowed to adjust the balance of phasic-tonic and tonic vestibular inputs to cluster F. The left and right columns show the performance before and after learning, respectively. **A,B**: Head velocity stimulus and eye velocity response during the VOR in the dark. The dashed lines show desired eye velocity. **C,D**: The solid line shows the output from a representative unit in cluster F. Because the vestibular stimulus was presented in the dark, the output from cluster F was determined entirely by the vestibular inputs from cluster 9 and the eye velocity inputs from cluster 7. The short dashed lines show the contribution of the vestibular inputs, derived by zeroing the connection weights between cluster 7 and this unit. The long dashed line shows the contribution of eye velocity, derived by subtracting the head velocity input from the output of the unit. **E,F**: The short and long dashed lines show the contributions of the phasic-tonic and tonic vestibular inputs to the output of cluster F.

Figure 5. Schematic diagram summarizing the conclusions of our simulations. The solid arrows show the flow of signals in the new hypothesis suggested here. The dashed arrows show the flow of signals in two other possible VOR pathways. Each box acts as a multiplication factor that represents a possible site of learning and the circles labeled F, V, and R represent the flocculus, the flocculus target neurons in the vestibular nucleus, and the retina, respectively. The separate head velocity pathways (**a1** and **a2**) to the circle labeled F provide a way to adjust separately the strength of the inputs from the tonic and the phasic-tonic vestibular afferents.

Figure 1

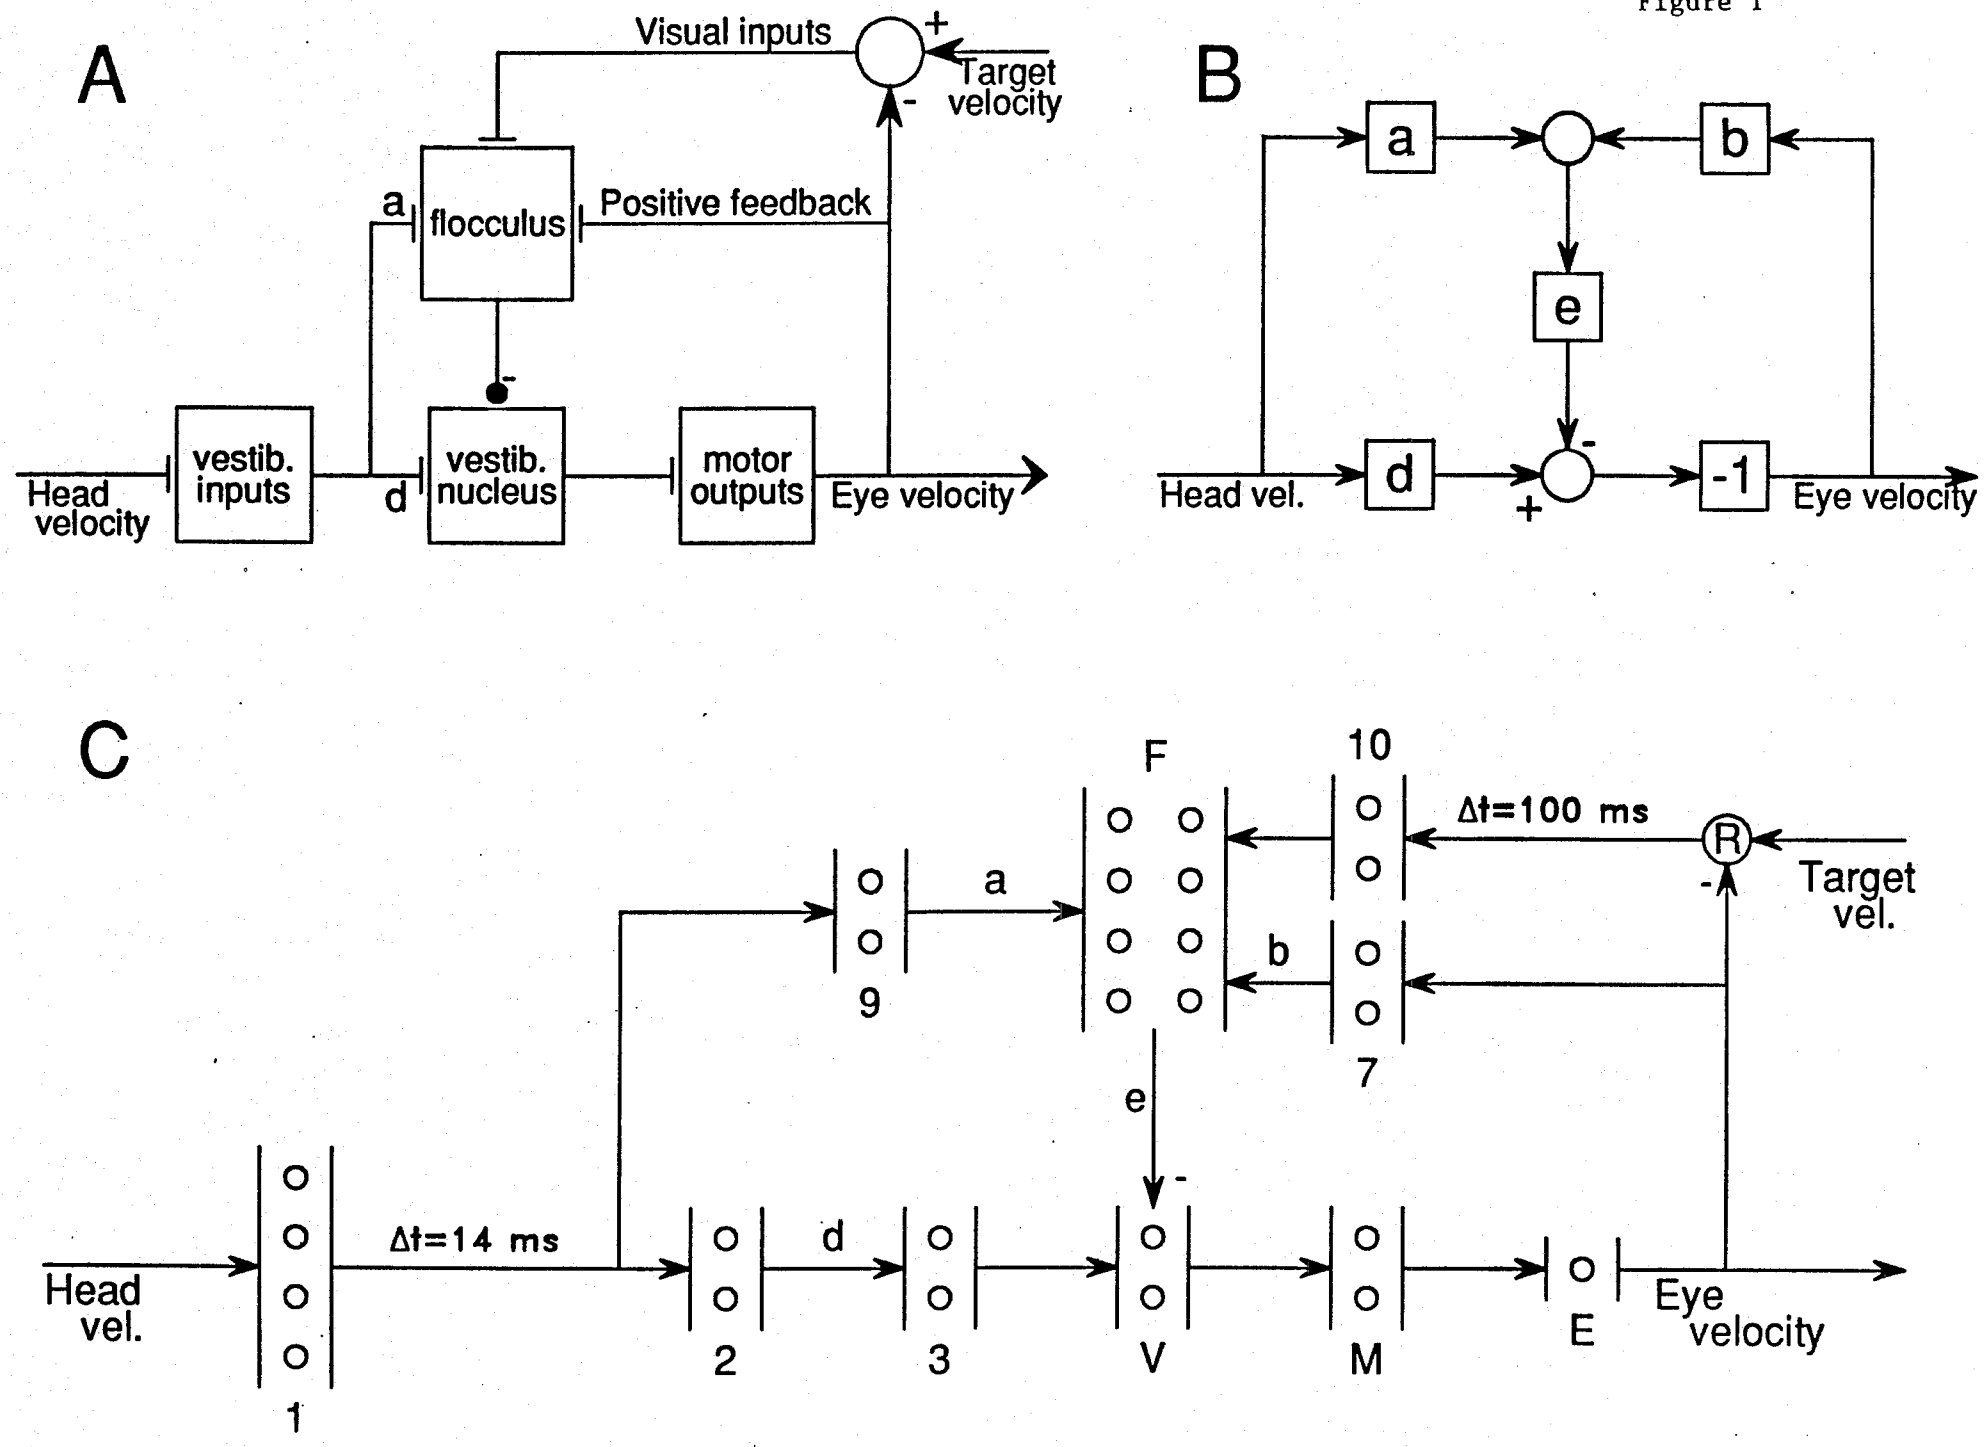
A
Visual inputs
+
Target velocity
-
a
flocculus
Positive feedback
-
Head velocity
vestib. inputs
d
vestib. nucleus
motor outputs
Eye velocity
B
a
b
e
-
Head vel.
d
+
-1
Eye velocity
C
F
10
Δt=100 ms
R
-
Target vel.
a
9
b
7
e
-
Δt=14 ms
Head vel.
1
d
2
3
V
M
E
Eye velocity

Figure 2

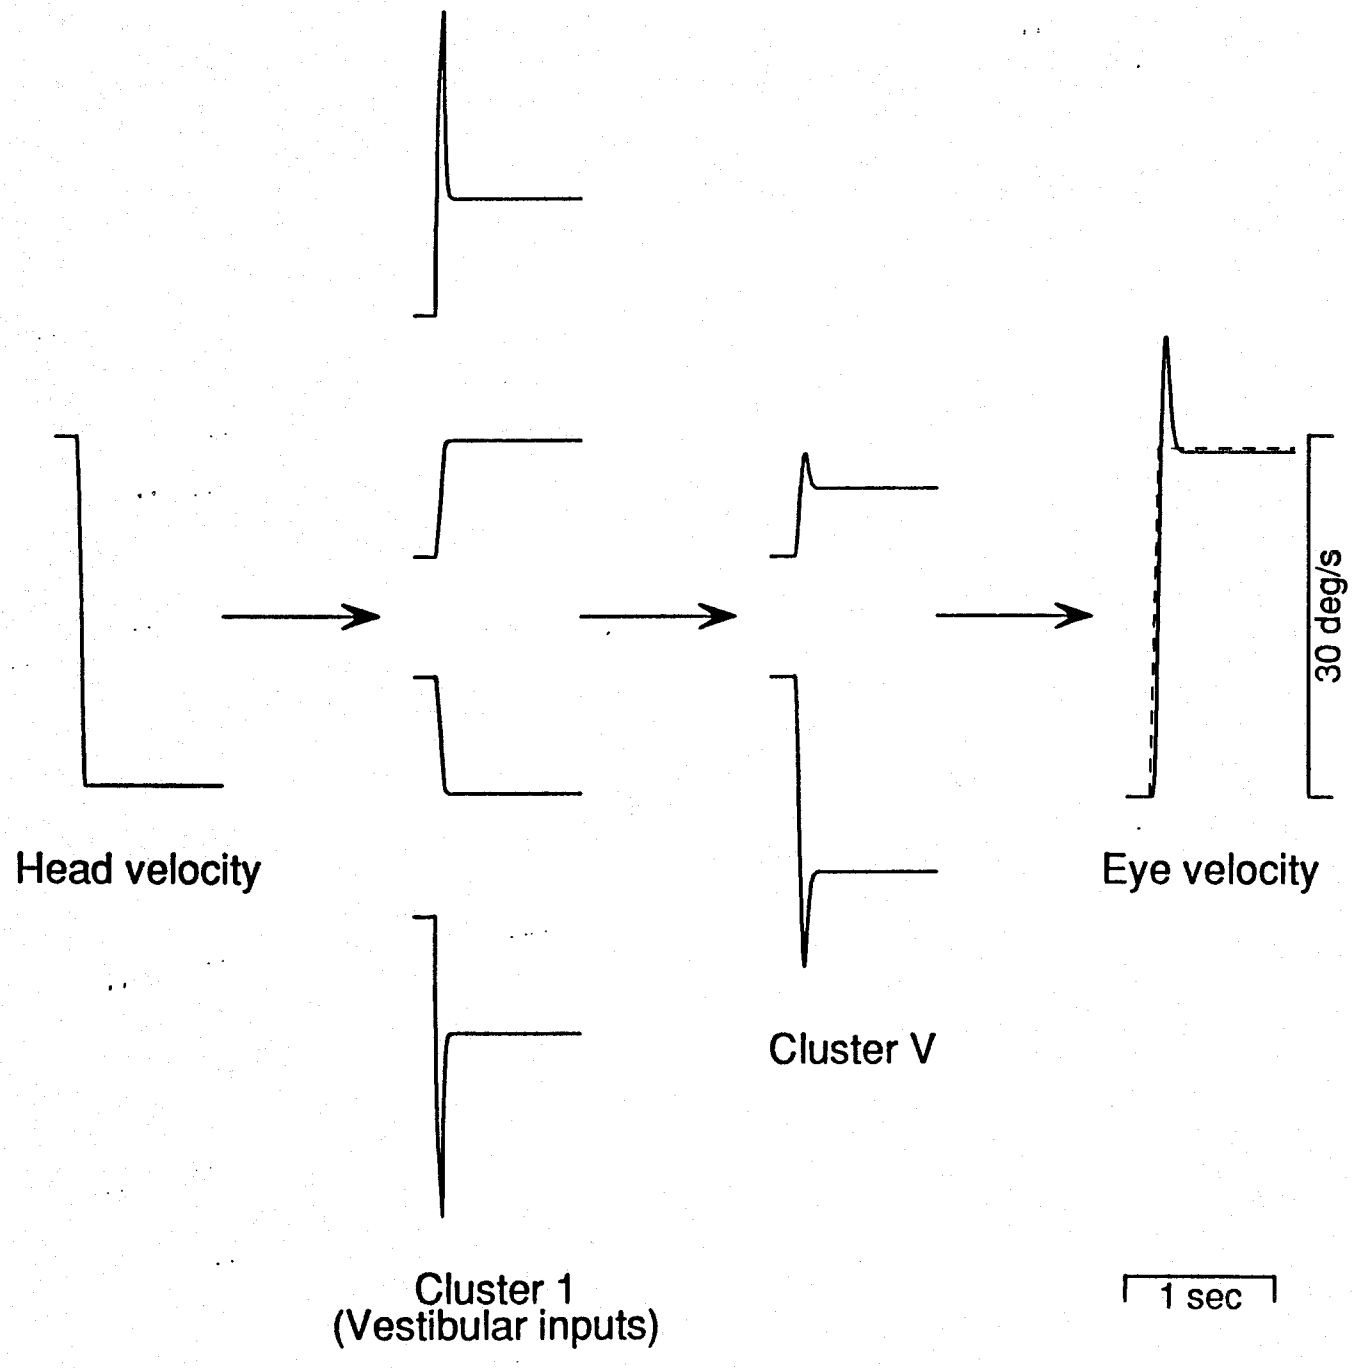
Head velocity
Cluster 1
(Vestibular inputs)
Cluster V
Eye velocity
30 deg/s
1 sec

Figure 3

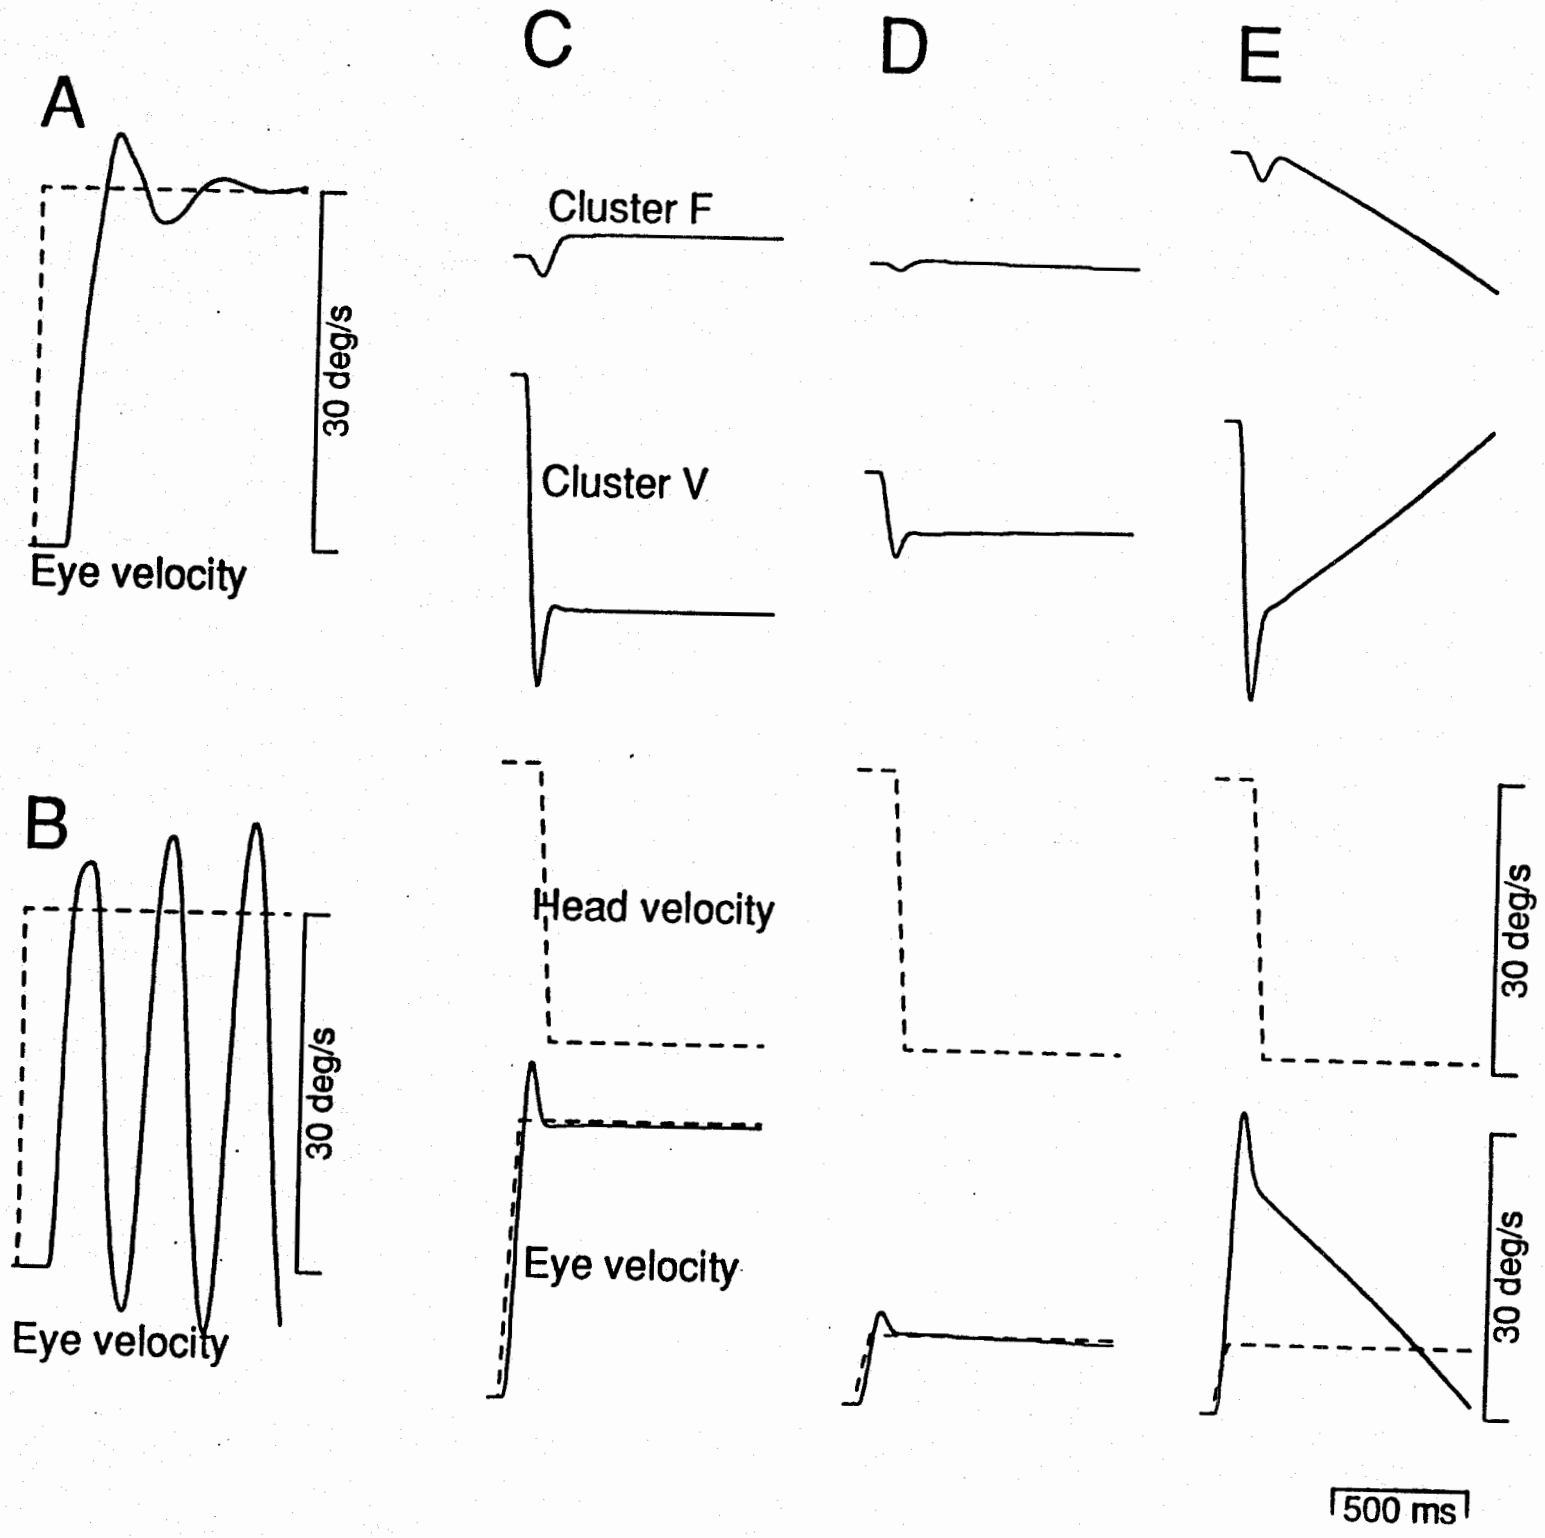
A
Eye velocity
30 deg/s
B
Eye velocity
30 deg/s
C
Cluster F
Cluster V
Head velocity
Eye velocity
D
E
30 deg/s
30 deg/s
500 ms

Figure 4

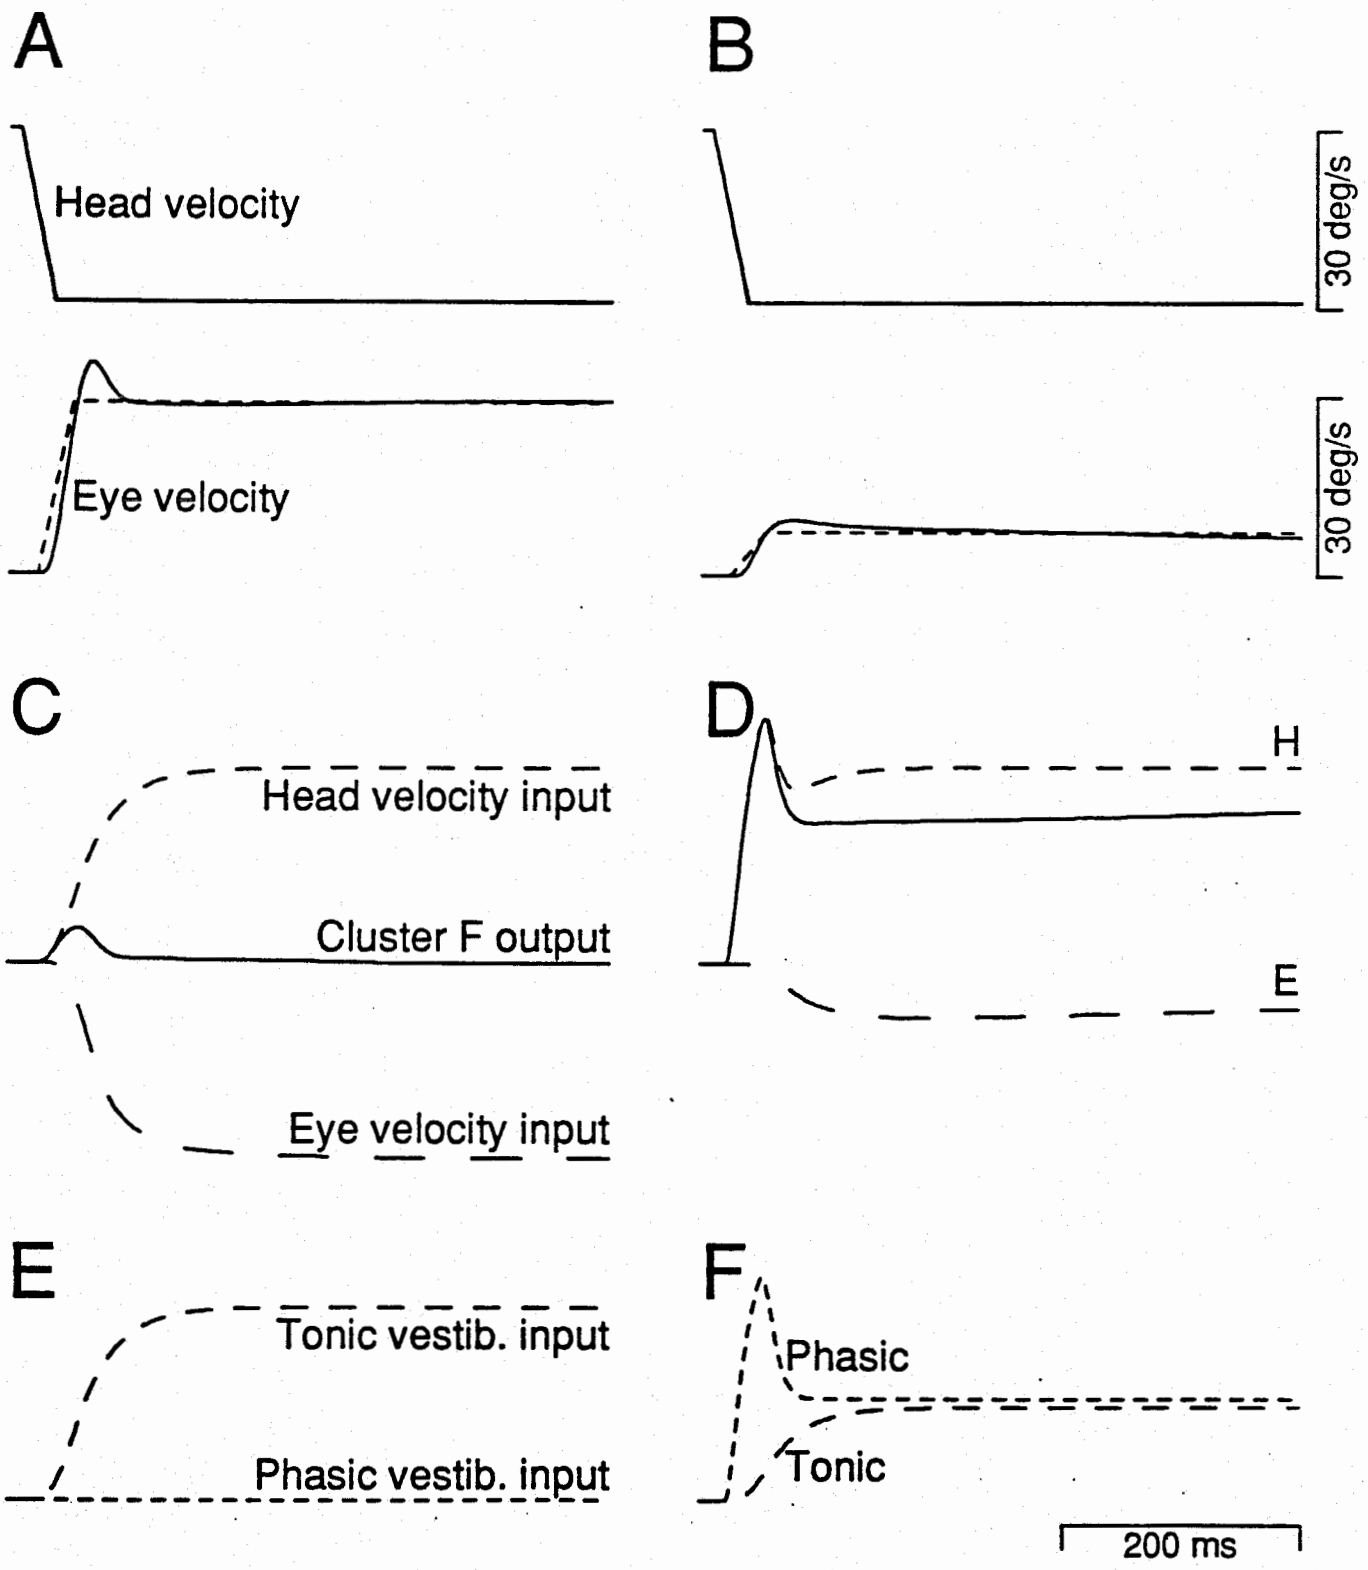
A
Head velocity
Eye velocity
B
30 deg/s
30 deg/s
C
Head velocity input
Cluster F output
Eye velocity input
D
H
E
E
Tonic vestib. input
Phasic vestib. input
F
Phasic
Tonic
200 ms

Figure 5

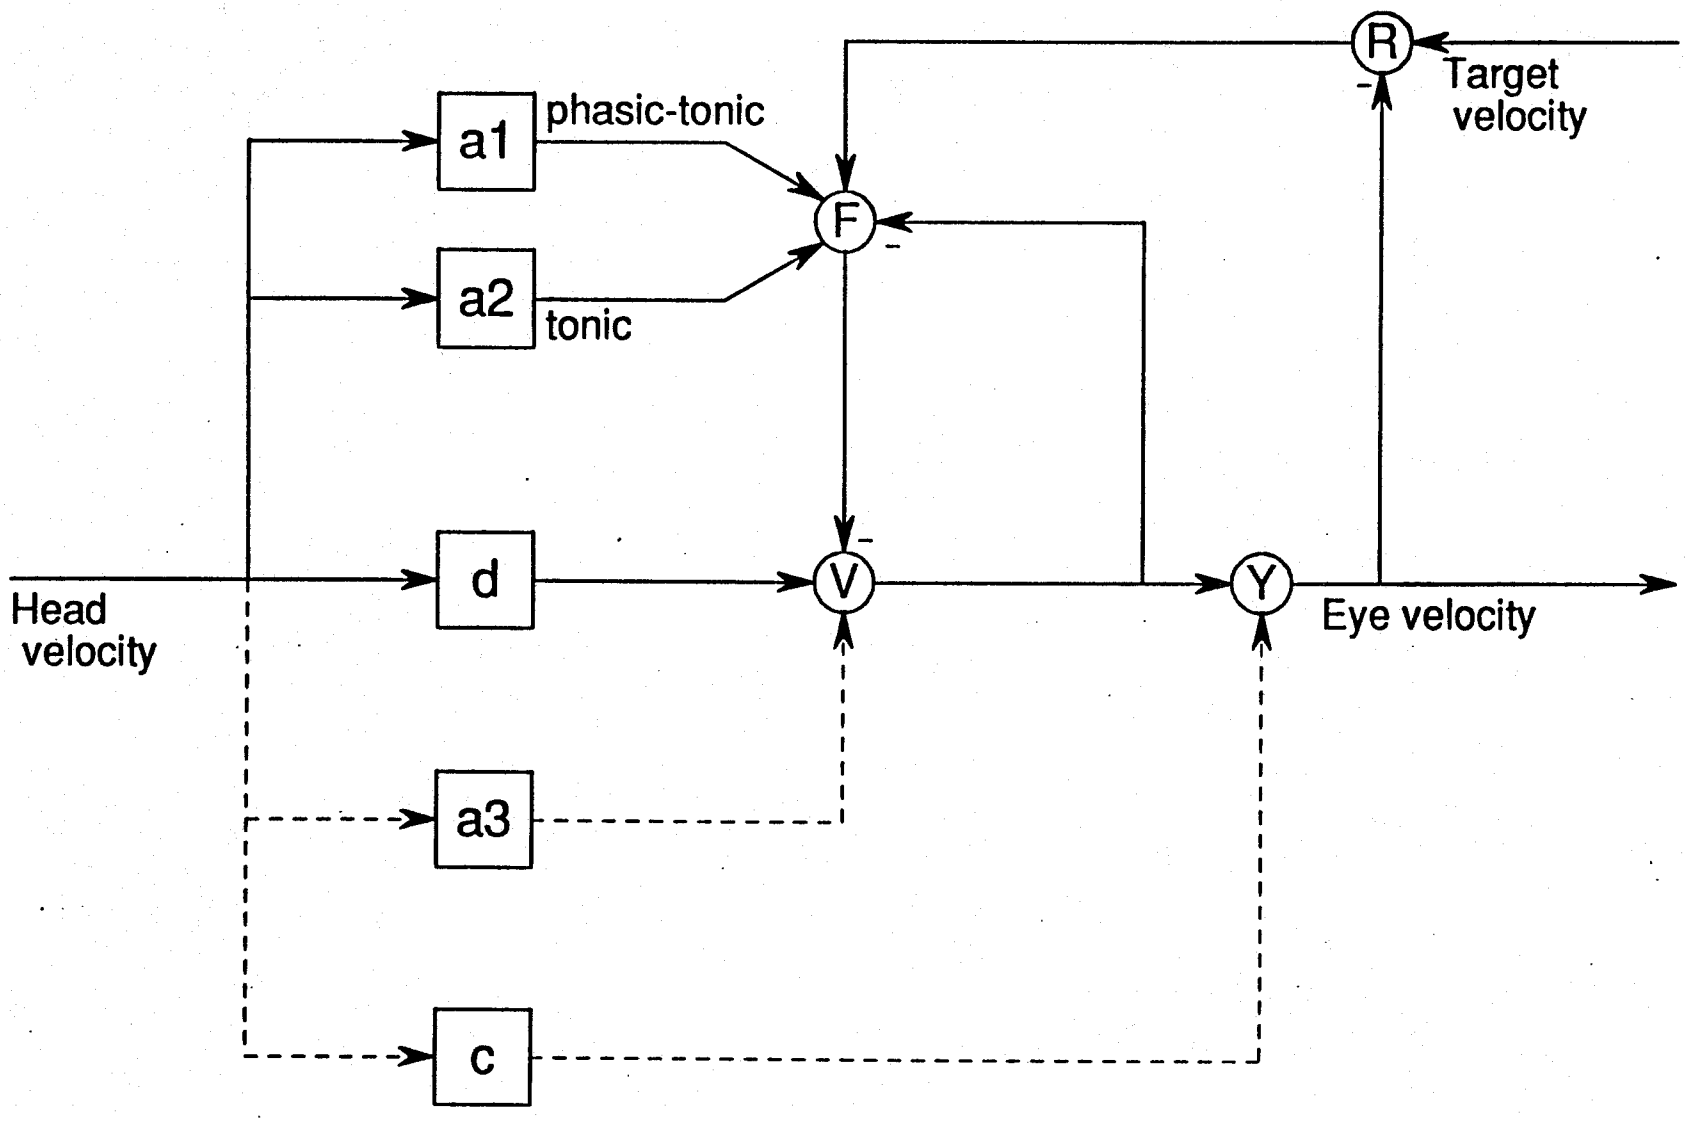
R
Target
velocity
a1
phasic-tonic
F
a2
tonic
d
V
Y
Head
velocity
Eye velocity
a3
c