



A pattern correlation model of vestibulo-ocular reflex habituation

T.J. Anastasio*

Beckman Institute and Department of Molecular and Integrative Physiology, The University of Illinois at Urbana/Champaign, Urbana, USA

Received 15 September 2000; revised 22 September 2000; accepted 22 September 2000

Abstract

Through the process of habituation, the response of the vestibulo-ocular reflex (VOR) is decreased by prolonged, sinusoidal stimulation at lower frequencies (≤ 0.1 Hz). Research on goldfish has uncovered frequency-specific and nonlinear behaviors associated with habituation of the goldfish VOR, which include phenomena that cannot be explained using dynamic linear and static nonlinear models. The unexplained phenomena are abrupt decreases at peak response, gain decreases far in excess of linear predictions based on phase, and violation of superposition. Their existence was attributed to a hypothetical switch that closed in the appropriate context. The pattern correlation model provides a new perspective on the process of VOR habituation. Rather than treat the stimulus as a continuous sinusoid, the pattern correlation model breaks it up into a number of discontinuous patterns. The pattern most closely correlated with the current stimulus then decreases the VOR response by the amount of that correlation times a pre-assigned weight. The pattern correlation model explains how the frequency-specific and the nonlinear behaviors may be related, and how the apparent switching phenomena may occur. © 2001 Elsevier Science Ltd. All rights reserved.

Keywords: Vestibulo-ocular reflex; Learning; Plasticity; Habituation; Cerebellum; Linear systems; Nonlinear systems; Pattern correlation

1. Introduction

The vestibulo-ocular reflex (VOR) is an important model system for the study of motor learning. VOR behavior can be modified through experience in various ways. This paper describes a new model of VOR habituation, in which a sinusoidal stimulus is habituated not as a continuous signal but as a set of separate fragments or patterns. The discontinuous operation of the model offers an alternative view of VOR habituation.

The function of the VOR is to stabilize the retinal image by producing eye rotations that were roughly equal and opposite to head rotations (Wilson & Melvill Jones, 1979). The VOR can be studied in isolation from the visual system by testing it in the dark. Exposure in the dark to sinusoidal head rotation for a prolonged duration (1 h or longer), at frequencies lower than 0.1 Hertz (Hz, where 1 Hz = 1 cycle/s), produces a decrease in the response of the VOR (Schmidt & Jeannerod, 1985). This process is known as VOR habituation.

On time scales shorter than those required for experience-dependent modification (hours), the VOR has been considered to behave as a linear system (e.g. Robinson, 1981). The

gain and phase of the sinusoidal response of a linear system can change as a function of the frequency of the sinusoidal stimulus (Oppenheim & Willsky, 1983). Recent research in goldfish has shown that the habituated VOR produces frequency-specific effects that can be described using a linear system model (Dow & Anastasio, 1999a). Yet, for a system to be linear, its response to a sinusoidal stimulus should also be sinusoidal. Moreover, it should obey the superposition principle, which stipulates that its response to a stimulus at a particular frequency should be the same whether that stimulus is presented alone or in combination with a stimulus at another frequency. During habituation to low frequency (0.01 Hz) sinusoidal head rotation, the goldfish VOR eye rotation response exhibits waveform distortions that include magnitude-dependency, asymmetry, and abrupt decreases (Dow & Anastasio, 1998). Also, the goldfish VOR response habituated at 0.01 Hz increases over ten times when the 0.01 Hz rotation is combined with rotation at 0.3 Hz (Dow & Anastasio, 1996). These results show that the VOR response during habituation to a sinusoidal stimulus is not sinusoidal, and that the habituated VOR violates the superposition principle of linear systems. These nonlinear effects cannot be simulated using a linear system model.

To understand VOR habituation more fully, it is necessary to reconcile both the frequency-specific and the

* Tel.: +1-217-244-2895; fax: +1-214-244-5180.

E-mail address: tja@uiuc.edu (T.J. Anastasio).

nonlinear behaviors associated with it. This paper describes a new model of habituation, based on pattern correlation rather than linear system theory, which can simulate both aspects of the data. The new model suggests that the process of VOR habituation treats a prolonged sinusoidal stimulus not as a continuous signal, but as a set of discontinuous fragments or patterns that habituate separately. The pattern correlation model offers a new view of the process of VOR habituation.

The Introduction continues with a brief overview of the VOR, followed by a detailed description of the goldfish VOR habituation data that is to be modeled. The pattern correlation model is described in Methods, and its ability to simulate the goldfish VOR habituation data is presented in Results. Related issues are considered in Discussion.

1.1. Overview of the VOR

Classical studies have established that the VOR is mediated by a relatively simple three-neuron pathway through the brainstem (Wilson & Melvill Jones, 1979). Vestibular afferents (VAs) from the semicircular canals transmit a head rotational velocity signal to vestibular nuclei neurons (VNs) in the brainstem. The VNs convert that signal into an eye rotational velocity command and send it on to the motoneurons of the eye muscles. The conversion of the VA head rotational velocity signal to an eye rotational velocity command involves dynamic signal processing, which can be described using linear system models.

Linear system models of the VOR are based on the transfer function of the high-pass filter (Robinson, 1981). The high-pass transfer function $H(s)$ is (Oppenheim & Willsky, 1983):

$$H(s) = Y(s)/X(s) = g[s\tau/(s\tau + 1)] \quad (1)$$

where X and Y are the Laplace transforms of the input and output, respectively, g is the gain constant, τ is the time constant in seconds, and s is complex frequency ($s = j\omega$ where $j^2 = -1$ and ω is frequency in radians/s). Eq. (1) encodes the gain (output amplitude/input amplitude) and phase difference (output phase angle – input phase angle) of the high-pass filter as the modulus and argument, respectively, of its complex value at any frequency. Positive and negative phase values are called phase leads and lags, respectively. As input frequency increases, the gain of a high-pass filter increases in proportion to frequency and then levels off at a constant value (equal to g). Phase lead transitions to zero, and both of these transitions are centered on the break frequency f_b [$f_b = 1/\tau$ rad/s or $1/(2\pi\tau)$ Hz].

The eye rotational velocities produced by the VOR, and the neural signals carried by the VA and VN neurons in the VOR pathway, all have high-pass dynamics. Those dynamics reflect the high-pass dynamics of the semicircular canal receptor, from which the head rotational velocity signal originates (Mayne, 1950; Wilson & Melvill Jones, 1979). However, the time constant of the VAs is shorter than that

of the VNs and the VOR. The time constant of the VAs in goldfish has been estimated to be about 2 s (Hartmann & Klinke, 1975, 1980), which is much shorter than the goldfish VOR time constant (Weissenstein, Ratnam, & Anastasio, 1996).

The VA time constant is lengthened through a process called velocity storage (Raphan, Matsuo, & Cohen, 1979). Velocity storage gives the VOR a higher gain at lower frequencies than that of the VA head rotational velocity signal that drives it. To illustrate this, the frequency responses of two high-pass filters (Eq. (1)) with the same gain constant of one but different time constants of 2 s or 80 s are shown in Fig. 1A,B (solid and dashed lines, respectively). The frequency responses are exhibited as Bode plots in which log gain and phase are plotted separately against log frequency (Oppenheim & Willsky, 1983). The high-pass filter serves as the basis for linear system models of the habituated VOR.

1.2. Habituation of the VOR in the goldfish

Habituation of the horizontal VOR to prolonged sinusoidal stimulation at lower frequencies was studied in the goldfish. The experimental data have been published previously (Dow & Anastasio, 1996, 1998, 1999a). They are reproduced in this paper (with permission) so that model behavior can be compared easily to that of the real goldfish VOR.

The goldfish received whole-body rotation about a head-centered, earth vertical axis. Rotation occurred at seven discrete frequencies (0.01, 0.03, 0.05, 0.1, 0.3, 0.5 and 1.0 Hz), with peak rotational velocities of 60 deg/s, which were delivered alone or summed together in pairs (e.g. 0.01 + 0.3 Hz). The stimulus was delivered in the dark to exclude any influence on eye movements from the goldfish optokinetic system (Keng & Anastasio, 1997). The response measure taken was horizontal eye rotational velocity. Goldfish that had never experienced experimental rotation were considered naïve. Habituation was produced by rotating a naïve goldfish in the dark for one hour at the habituating frequency. The relationship between head and eye rotational velocity during and after habituation of the VOR is summarized below.

As in other vertebrates (Schmidt & Jeannerod, 1985), the response of the goldfish VOR decreases with time as the duration of lower frequency stimulation is prolonged. An example is shown in Fig. 2. The stimulus is a sinusoidal rotation delivered at a frequency of 0.01 Hz (and peak velocity of 60 deg/s). Head rotational velocity is shown in Fig. 2A. To simplify presentation and analysis of the data, head rotational velocity is inverted, removing the 180 deg phase difference between eye and head rotation that would occur if the VOR worked perfectly at each frequency. The stimulus endures for almost 20 min and encompasses more than eleven cycles. VOR eye rotational velocity in response to this stimulus is shown in Fig. 2B. Clearly, the eye rotational velocity produced by the VOR in response to this

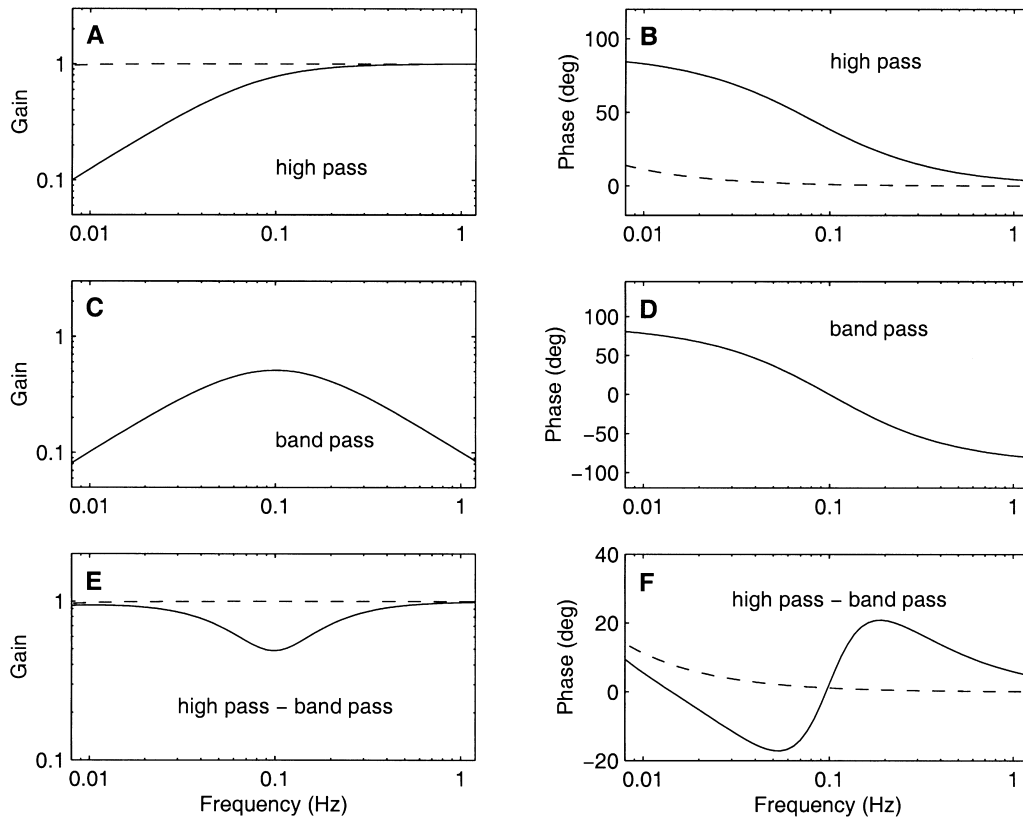


Fig. 1. Bode plots illustrating the frequency responses of various transfer functions. Gain plots (A, C, E) are log-log and phase plots (B, D, F) are log-linear. A and B, gain and phase of a high-pass filter (Eq. (1)) with time constant equal to 2 s (solid lines) or 80 s (dashed lines). C and D, gain and phase of a band-pass filter (Eq. (4)) with central frequency 0.1 Hz and corner frequency spacing 0.01 log units (Eq. (5)). E and F, gain and phase of a high-pass filter with time constant equal to 80 s (dashed lines), and of the same high-pass filter after subtracting a band-pass filter from it (Eq. (6)) with central frequency 0.1 Hz and corner frequency spacing 0.01 log units (solid lines). Gain constant is one for all transfer functions.

lower frequency stimulus decreases in time as the stimulus endures. This experience-dependent decrease, brought about through prolonged exposure to a lower frequency head rotation, is an example of VOR habituation.

The most striking feature of the experimental data in Fig. 2 is that the response is heavily distorted. Although the stimulus is sinusoidal, the response is nonsinusoidal. Of course, a habituating response to a sinusoidal stimulus is not expected to be perfectly sinusoidal, since the response decreases with time. However, the decrease might be expected to occur gradually and symmetrically over the duration of the stimulus. The habituating response shown in Fig. 2B shows abrupt decreases in amplitude, such as that appearing in the cycle beginning at the asterisk. The response is also asymmetrical. It appears to habituate more rapidly for rotations to one side than to the other (positive head or eye velocities represent rightward or leftward rotations, respectively). Distortions such as abrupt decreases and asymmetries can be characterized as nonlinear behaviors.

1.2.1. Nonlinear behavior associated with VOR habituation

Habituation was thought to be a long-term process that would not produce nonlinear distortions over the short-

term. The focus here is on distortions that occur within one cycle of the response. These can be characterized using the xy plot, in which VOR eye rotational velocity is plotted against head velocity for one cycle. The plot is made after any phase difference between stimulus and response is removed. This method has been used to analyze the nonlinear aspects of the VA response (Segal & Outerbridge, 1982). For a linear system the response to a sinusoidal stimulus is also sinusoidal, and the xy plot is a straight line. Response nonlinearity can be characterized by the departures of the xy plot from a straight line. VOR distortion was observed at all frequencies that produced habituation. However, the amount of habituation, and the amount of distortion, was greatest at the lowest frequency studied, which was 0.01 Hz. Some of the nonlinear behaviors of the VOR during and after habituation at 0.01 Hz were characterized using xy plots and are shown in Fig. 3.

The goldfish VOR, recorded in the dark, was almost never linear at 0.01 Hz. Effects due to habituation could be observed even within the first few cycles (Dow & Anastasio, 1998). An example of early distortion is shown in Fig. 3A. The response shows a dead-zone type of magnitude dependency, in which the response is weaker midrange than near the peaks of the stimulus. Another phenomenon that sometimes occurred after several cycles of stimulation




Fig. 2. Time series illustrating habituation of the goldfish VOR. A, sinusoidal head rotation at a frequency of 0.01 Hz and peak rotational velocity of 60 deg/s, delivered for almost 20 min in the dark. B, the VOR eye rotational velocity response to the prolonged stimulus decreases over time. Positive velocities indicate rightward head rotations or leftward eye rotations. The head rotational velocity record has been inverted to remove the 180 deg phase difference that would be observed if VOR eye rotation perfectly opposed head rotation. (Redrawn with permission from Dow and Anastasio, 1998).

at 0.01 Hz was the abrupt response decrease. An abrupt decrease occurs in the time series of Fig. 2B within the cycle marked with the asterisk. Abrupt decreases were observed only at the peak of the response (Dow & Anastasio, 1998). When the abrupt decrease occurred at the peak of one cycle it often recurred at the peak of one or more subsequent cycles (Fig. 2B). The abrupt decrease produced a figure-eight xy plot, as illustrated in Fig. 3B. The asterisks in Fig. 2B and 3B mark the same time point.

As habituation at 0.01 Hz progressed, the VOR response usually became asymmetrical, where the response was decreased more to one side than to the other. Asymmetrical responses produced crescent xy plots. Fig. 3C shows an xy plot of a crescent due to left-side habituation, as in the time series of Fig. 2B. Following 1 h of rotation at 0.01 Hz, VOR gain decreased by over 20 times (Dow and Anastasio, 1998). Both the time series and the xy plot appeared as flat lines (not shown). Habituation at 0.01 Hz both decreased and distorted the VOR response at nearby frequencies. Fig. 3D shows that a pronounced dead-zone xy plot characterized the response at 0.03 Hz following habituation at 0.01 Hz.

The dead-zone and crescent xy plots characterize static nonlinearities that could be explained on the basis of the essentially sigmoidal nonlinearity of the neurons that mediate the VOR (Dow & Anastasio, 1998). However, the figure-eight characterizing the abrupt decrease could only

be explained by invoking a hypothetical switch that decreases the VOR response at its peak (ibid.). The pattern correlation model offers a consistent explanation for all of the nonlinear distortions illustrated in Fig. 3 (see below).

The dead-zone, crescent, and figure-eight profiles represent idealizations of the nonlinear distortions that were observed during VOR habituation in goldfish. The three types of distortion could be present in varying degrees and usually occurred in combination. For example, the abrupt decrease illustrated in Fig. 2B occurs in combination with a partial, asymmetrical habituation on the same side. As habituation progresses the VOR response expresses different combinations of distortions. The pattern correlation model applies equally to sequences of combinations as to individual distortions, but simulations focus on the latter for simplicity.

1.2.2. Frequency-specific behavior associated with VOR habituation

Despite nonlinear distortions, the frequency response of the habituated VOR was studied as though it were linear by fitting least-squares sinusoids to the responses. Responses to sinusoidal stimulation at a single frequency were fit with an equation of the form:

$$y_e(t) = a \sin(2\pi ft + \theta) + o \quad (2)$$

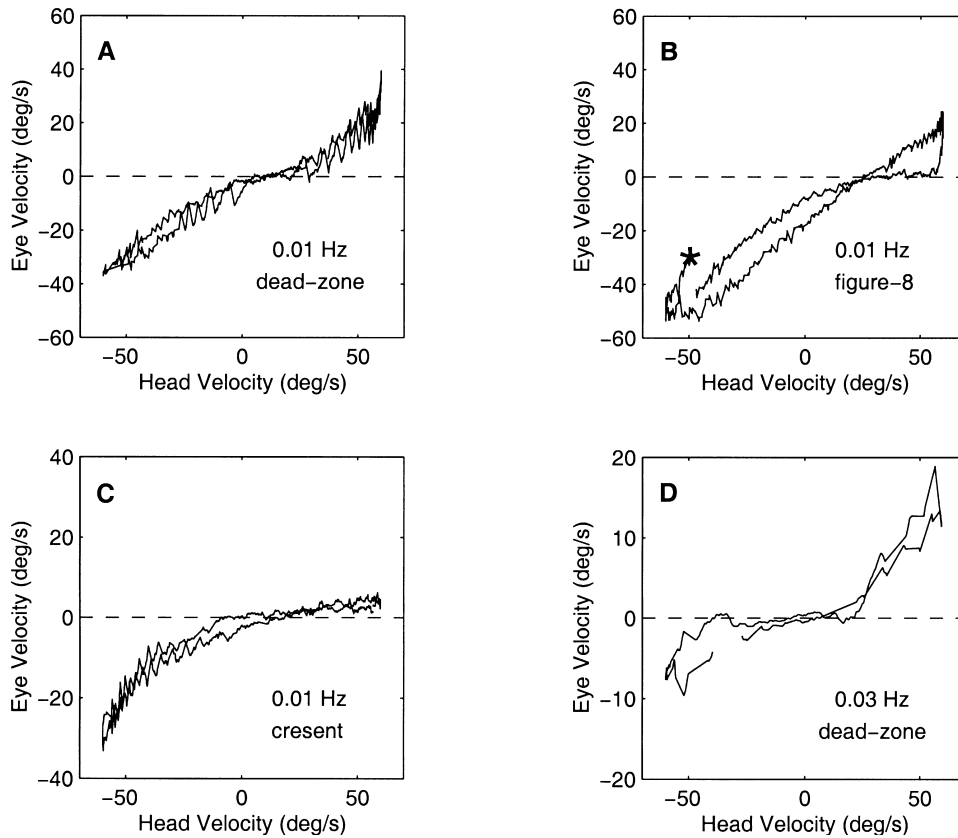


Fig. 3. xy plots illustrating nonlinear responses during and after habituation of the goldfish VOR to sinusoidal rotation at 0.01 Hz. A, dead-zone at 0.01 Hz observed at an early stage of 0.01 Hz habituation. B, figure-eight at 0.01 Hz observed at an intermediate stage of 0.01 Hz habituation. C, crescent at 0.01 Hz observed at an intermediate stage of 0.01 Hz habituation. D, pronounced dead-zone observed at 0.03 Hz following habituation at 0.01 Hz. (Redrawn with permission from Dow and Anastasio, 1998).

where $y_e(t)$ is the estimated value of the sinusoid at each discrete time step t , f is the frequency of stimulation, and a and θ are the amplitude and phase angle (in deg) of the response at that frequency. The parameter o is the offset, which was practically zero and can be ignored. Fitting was done using an optimization routine that minimized the least-squares error E between the experimental data $y(t)$ and the estimated sinusoid:

$$E = \sum_t [y_e(t) - y(t)]^2 \quad (3)$$

where the summation is over all time points t . Fitting least-squares sinusoids provides linear estimates of the amplitude and phase angle of the VOR responses at each frequency.

The frequency response was first estimated in naïve goldfish to establish a baseline (Weissenstein et al., 1996; Dow & Anastasio, 1999a). Linear estimates of gain (response amplitude/stimulus amplitude) and phase (response phase angle – stimulus phase angle) were made at each frequency in the spectrum (0.01, 0.03, 0.05, 0.1, 0.3, 0.5, and 1.0 Hz). A high-pass filter transfer function, similar to Eq. (1), was fit to the naïve frequency response data using an optimization routine. The gain (A, C, E, and G) and phase (B, D, F, and

H) of the high-pass filter fit to the naïve VOR frequency response data are reproduced in Fig. 4 (dashed lines). The Bode plots show that naïve gain is roughly constant and phase is near zero over a broad range of the frequency spectrum.

To study the VOR frequency response following habituation at single frequencies, linear estimates of gain and phase at each frequency in the spectrum were made following rotation for 1 h at 0.01, 0.03, 0.05, or 0.1 Hz to habituate the VOR at that frequency. The frequency responses are shown in Fig. 4 as Bode plots (gain, A, C, E, and G; and phase, B, D, F, and H). VOR frequency response data following habituation at 0.01 Hz (A and B), 0.03 Hz (C and D), 0.05 Hz (E and F), and 0.1 Hz (G and H) are shown as open circles in Fig. 4. Each circle in Fig. 4 represents the median value for several goldfish of the gain or phase estimated at each frequency. The habituating frequency is indicated on each plot, and gain and phase at the habituating frequency are marked with crosses.

The Bode plots in Fig. 4 show that VOR gain at the habituating frequency is decreased relative to naïve. Gain at 0.01 Hz decreased over 20 times following habituation at that frequency. Gain at the habituating frequency decreased

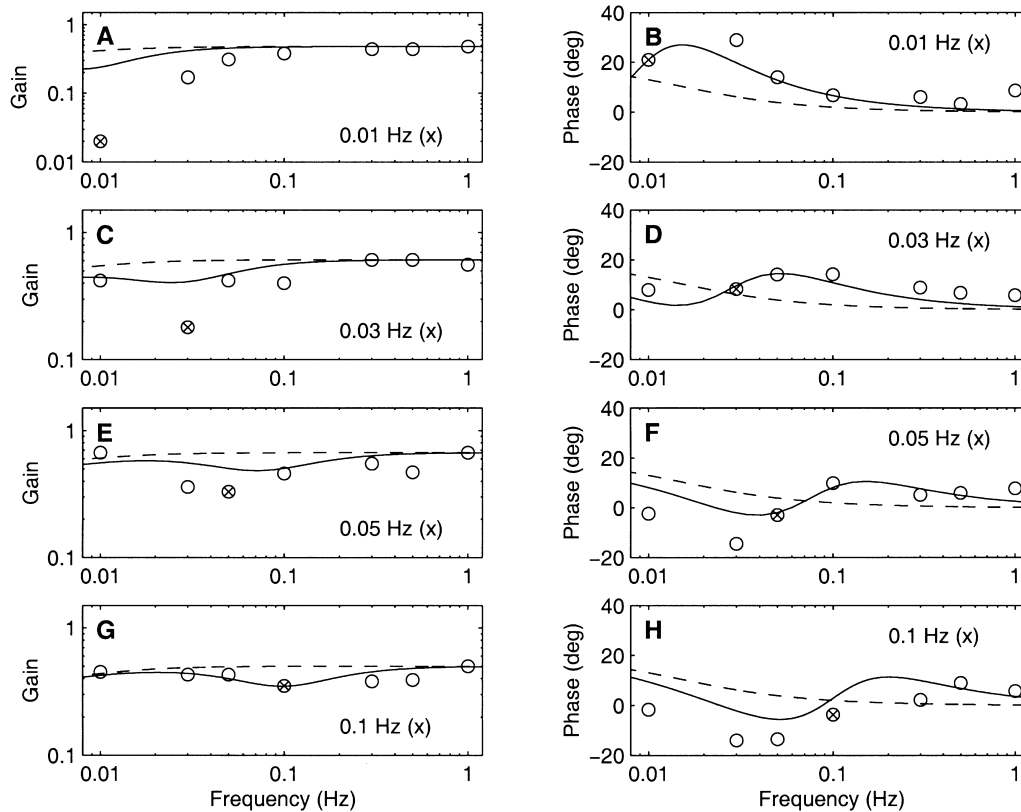


Fig. 4. Bode plots illustrating the frequency responses of the goldfish VOR following habituation at various single frequencies of head rotation. Gain plots (A, C, E, G) are log-log and phase plots (B, D, F, H) are log-linear. The habituating frequency is 0.01 Hz for A and B, 0.03 Hz for C and D, 0.05 Hz for E and F, and 0.1 Hz for G and H. Open circles represent median gain and phase values from several goldfish. Gain and phase at the habituating frequency are marked with crosses. The dashed line shows the frequency response of a high-pass filter transfer function that was fit (least-squares) to data from naïve goldfish. The solid line shows the frequency response of a transfer function in which a band-pass filter is subtracted from the high-pass filter describing the naïve VOR. The high-pass – band-pass transfer function is fit to habituated VOR phase data only. (Redrawn with permission from Dow and Anastasio, 1999a, from which additional details can be obtained.)

progressively less as the habituating frequency was increased up to 0.1 Hz. Habituation was not observed at frequencies of 0.3 Hz and higher (Dow & Anastasio, 1997). Habituation at 0.1 Hz and lower decreased VOR gain not only at the habituating frequency but also at nearby frequencies. This produced a dip in gain centered at the habituating frequency.

The phase of the VOR frequency response following habituation at single frequencies also had a characteristic profile. The difference in phase between the naïve and the habituated response was practically zero at the habituating frequency. In contrast, habituated VOR phase led naïve phase at nearby higher frequencies, and lagged naïve phase at nearby lower frequencies. This produced a phase crossover centered at the habituating frequency.

Similar gain-dip and phase-crossover effects have been described for down-adaptation of the VOR (Lisberger, Miles, & Optican, 1983; see Discussion), and have been simulated using a linear system model based on the band-pass filter (Miles, Optican, & Lisberger, 1985). The band-pass filter is a linear system element that, in its simplest form, is constructed from a high-pass filter (Eq. (1)) and a

low-pass filter in series (Oppenheim & Willsky, 1983). The transfer function for such a simple band-pass filter $B(s)$ is:

$$B(s) = Y(s)/X(s) = g[s\tau_h/(s\tau_h + 1)][1/(s\tau_l + 1)] \quad (4)$$

where τ_h and τ_l are the time constants of the high-pass and low-pass filters, respectively. It is convenient to express the time constants as functions of a common central frequency (f_c in Hz) and spacing (λ), where λ determines the separation in log units of the high-pass and low-pass break frequencies from f_c :

$$\tau_h = 1/(2\pi f_c 10^{-\lambda}) \text{ and } \tau_l = 1/(2\pi f_c 10^{\lambda}). \quad (5)$$

The frequency response of a band-pass filter for which $g = 1$, $f_c = 0.1$ Hz, and $\lambda = 0.01$ is illustrated in Fig. 1C,D. The gain dip and phase crossover observed for the goldfish VOR habituated at a single frequency have been modeled using a simple transfer function, labelled $D(s)$, in which a band-pass filter is subtracted from the high-pass filter model of the naïve (unhabituated) VOR (Dow & Anastasio, 1999a):

$$D(s) = H(s) - B(s) \quad (6)$$

Fig. 1E,F (solid lines) illustrate the frequency response of $D(s)$ (Eq. (6)), where the high-pass filter time constant $\tau = 80$ s, and the band-pass filter $f_c = 0.1$ Hz and $\lambda = 0.01$, and $g = 1$ for both filters. The high-pass filter frequency response by itself is also shown for comparison (Fig. 1E,F, dashed lines).

The high-pass – band-pass transfer function (Eq. (6)) could not completely describe the VOR frequency response following single-frequency habituation (Dow and Anastasio, 1999a). Therefore, this transfer function (Eq. (6)) was fit only to the phase characteristic of the frequency response using a least-squares routine to minimize the following error:

$$E = \sum_f [\phi_e(f) - \phi(f)]^2 \quad (7)$$

where $\phi_e(f)$ are the estimated and $\phi(f)$ are the experimental phase values, and the summation is over the seven frequencies f in the spectrum. Both the gain and the phase of the least-squares transfer function are reproduced in Fig. 4 (solid lines).

The high-pass – band-pass transfer function (Eq. (6)) does a reasonably good job of fitting the phase characteristic of the VOR frequency response following habituation at each single frequency. It also provides a good fit to gain for habituation at 0.1 Hz. However, it overestimates gain for habituation at lower frequencies. The amount of overestimation grows as the habituating frequency decreases. Gain is overestimated by an order of magnitude at 0.01 Hz for habituation at that frequency (Fig. 4A). Increasing the gain constant (g) of the band-pass filter increased the depth of the dip, but also produced deviations in phase that were much larger than those observed. Clearly, the linear, high-pass – band-pass model is inadequate to fully describe the VOR frequency response following single-frequency habituation. The decrease in gain below that predicted by the linear model was explained by invoking a hypothetical switch, which decreased VOR gain at and near the habituating frequency but not at other frequencies, and did not affect phase at any frequency (Dow & Anastasio, 1999a). Experimental support for this hypothetical switch came from studies of superposition.

1.2.3. Violation of superposition by the habituated VOR

The response of a linear system to a stimulus at a particular frequency should be the same whether that stimulus is presented alone or in combination with a stimulus at another frequency. This is the superposition principle of linear systems (Oppenheim and Willsky, 1983). Superposition was tested by first habituating the VOR at 0.01 Hz, and then measuring VOR gain at 0.01 Hz when that stimulus was presented alone or in combination with a stimulus at 0.3 Hz (Dow & Anastasio, 1996). Gain was measured by fitting sinusoids to the responses. Responses to stimulation at a single frequency were fit with Eq. (2), while responses to stimulation at two summed frequencies were fit with an

equation of the form:

$$y_e(t) = a_1 \sin(2\pi f_1 t + \theta_1) + a_2 \sin(2\pi f_2 t + \theta_2) + o \quad (8)$$

where $y_e(t)$ is the estimated value of the sinusoid at each discrete time step t , the f_i [$i = (1,2)$] are the two summed frequencies of stimulation, and the a_i and θ_i are the amplitudes and phase angles (in deg) of the response at the corresponding frequency. The parameter o is the offset, which was practically zero and can be ignored. The error measure minimized is given in Eq. (3).

The results of the superposition test are reproduced in Fig. 5. One hour of rotation in the dark at 0.01 Hz reduced the naïve gain by more than a factor of 20 (Fig. 5A,B). The habituated response in Fig. 5B is re-plotted in Fig. 5D on an expanded scale, to illustrate that it was very small but still measurable. When rotation at 0.01 Hz was combined with rotation at 0.3 Hz, gain at 0.01 Hz immediately increased to half of its naïve value (Fig. 5C). Gain decreased again to its habituated level immediately after the 0.3 Hz component was removed (not shown). These data show that the habituated VOR violates the superposition principle of linear systems. The data were explained by invoking a hypothetical switch that is closed when the habituating frequency is presented alone and further decreases VOR gain. The switch opens when the habituating frequency is presented in combination with some other frequency, presumably leaving only the linear component of VOR habituation.

In summary, previous research on the habituated VOR uncovered frequency-specific and nonlinear behaviors that were simulated using dynamic linear and static nonlinear models, respectively. These models had to rely on a hypothetical switch to fully explain the observed phenomena. The pattern correlation model can explain both the frequency-specific and nonlinear behaviors, and offers insight into the nature of the switching mechanism.

2. Methods

The main feature of the pattern correlation model is that the habituation process does not treat the sinusoidal input as a continuous function, but as a set of discontinuous fragments or patterns that habituate separately. The motivation for this alternative assumption comes directly from the raw data. A time series record of the VOR as it habituates to a 0.01 Hz sinusoidal head rotation is shown in Fig. 2B. Examination of the time series suggests that different portions of the response habituate independently of one another. For example, the record clearly shows that the positive response half cycle habituates more rapidly than the negative half cycle. Segments even smaller than a half cycle may also be treated independently by the habituation process. For example, the positive half cycle following the asterisk appears split into quarter cycles, where the second quarter cycle is habituated more than the first. A possible choice of patterns that emerges from this examination of the raw data

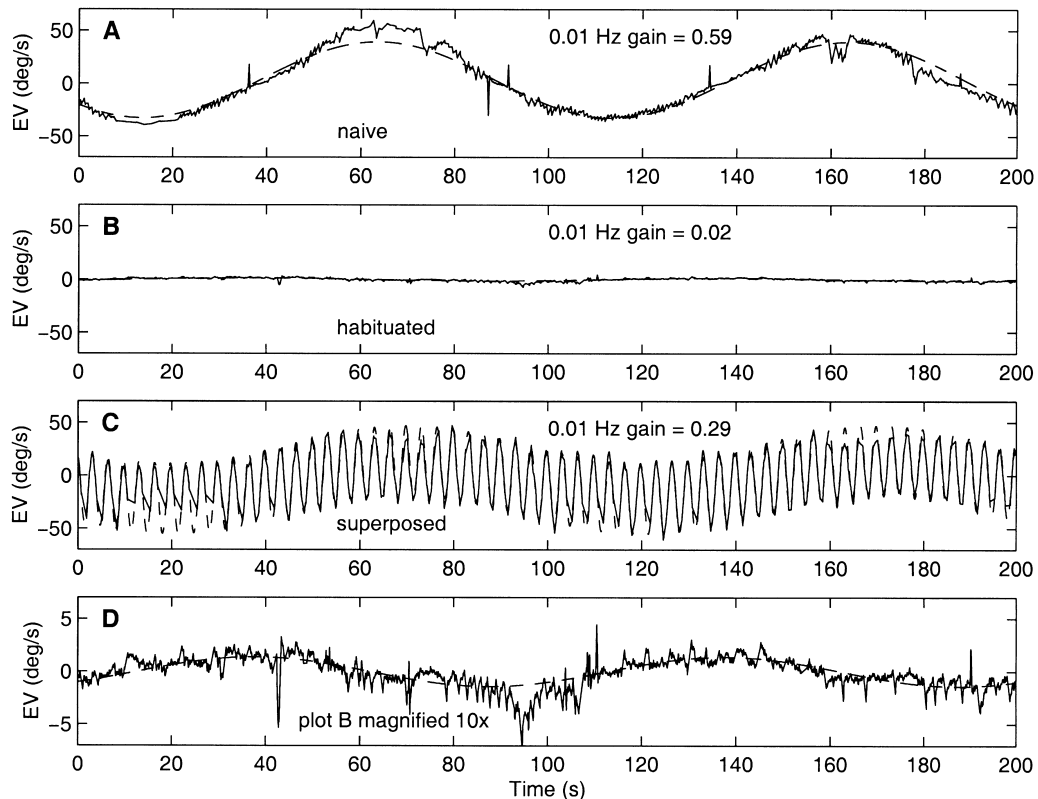


Fig. 5. Violation of the superposition principle by the VOR in habituated goldfish. A, the naïve response at 0.01 Hz (solid line) and best fit, least-squares sinusoid (dashed line). B, the response at 0.01 Hz following habituation at 0.01 Hz (solid line) and least-squares sinusoid (dashed line). C, the response to a combined head rotation at 0.01 Hz and 0.3 Hz following habituation at 0.01 Hz (solid line) and least-squares sum of sinusoids (dashed line). The gain of the 0.01 Hz component is greatly increased over its habituated level by superposition of the 0.3 Hz component. D, the plot in B magnified 10 times. EV, eye rotational velocity. (Redrawn with permission from Dow and Anastasio, 1996).

would be to divide the sinusoidal input into quarter cycles and have those habituate separately. A more well reasoned choice of patterns can be derived from an argument based on input uncertainty and the purpose of VOR habituation.

2.1. Setting the patterns in the pattern correlation model

Under the uncertainty hypothesis, the purpose of VOR habituation is to learn to reduce the response of the VOR to any repeated input pattern that is uncertain. Uncertain inputs are those that have a low signal to noise ratio (Cover & Thomas, 1991). If we assume a constant level of background noise, then the signal to noise ratio of the input will depend upon its amplitude. The amplitude of the input to the VOR is dependent on both frequency and phase.

The VOR receives input concerning head rotational velocity from the vestibular afferents (VAs). Although the VAs differ in their firing-rate variability, all are stochastic to some extent (e.g. Fernandez & Goldberg, 1971; Goldberg & Fernandez, 1971a,b; Anastasio, Correia, & Perachio, 1985). By the central limit theorem, the distribution of the combined input from the population of VAs could be approximated as Gaussian. An example of Gaussian background noise, at an amplitude chosen arbitrarily for purposes of illustration, is shown in Fig. 6A. This background noise

would be combined with any signal that is transmitted by the VAs. Signals that are large or small relative to the background noise will be certain or uncertain, respectively.

Due to the high-pass properties of the VA signal, input peak amplitude will decrease as frequency decreases below the break frequency. For goldfish VAs, the break frequency is near 0.1 Hz, and gain is almost ten times higher at 0.1 Hz than at 0.01 Hz (Fig. 1A, solid line). The effects of this ten-fold peak amplitude difference on signal to noise ratio are illustrated in Fig. 6B,C. Signal to noise ratio is lower, and uncertainty is higher, at 0.01 Hz than at 0.1 Hz. If the hypothesis is correct that the purpose of habituation is to reduce the VOR response to uncertain inputs, then the habituation process should reduce the VOR response more at lower frequencies. This is indeed borne out by the data (Fig. 4). The frequency dependence of VOR habituation is considered further in Discussion. Input uncertainty is also dependent upon the phase of the input, and this has implication for setting the patterns in the model.

The VA response will be modulated periodically by a sinusoidal head rotation, even if the response is not perfectly sinusoidal. The peak response will have a higher signal to noise ratio than the other response phases. This is illustrated for the 0.01 Hz input in Fig. 6C. Signal to noise ratio is poor overall, but it is lower for midrange




Fig. 6. Signal to noise ratio as a function of signal amplitude. A, background noise (dots) drawn from a standard Gaussian distribution (mean = 0, variance = 1) and scaled by a factor of 0.4. B, background noise by itself (dots) and after it has been combined with a sinusoidal signal at a frequency of 0.1 Hz and an amplitude of 10 (solid line). C, background noise by itself (dots) and after it has been combined with a sinusoidal signal at a frequency of 0.01 Hz and an amplitude of 1 (solid line). The overall signal to noise ratio is obviously higher in B than in C. In C, segments near the peak have a higher signal to noise ratio than midrange.

amplitudes than at the peaks. If the sinusoidal signal is broken up into fragments, then the signal to noise ratio of each fragment will depend upon its amplitude, or distance from the baseline. Patterns in the pattern correlation model are distinguished on this basis, as illustrated using a perfect sinusoid in Fig. 7.

Fig. 7 shows a sinusoid at 0.01 Hz with amplitude one. Each dot on the curve represents the value of the sinusoid at each second. Isoamplitude lines (dashed lines) divide the sinusoid into fragments that are each 10 s long. These fragments are labelled 1-5 for the positive half cycles and a-e for the negative half cycles. The amplitudes of the isoamplitude lines are arbitrary, and are chosen for convenience to make all the fragments the same length. The important point is that individual fragments are distinguished from each other on the basis of amplitude, or distance from the baseline. Those with the greatest distance have the largest signal to noise ratio and are the least uncertain. It is reasonable to suppose that the habituation process distinguishes patterns on this basis.

2.2. Model construction and operation

The pattern correlation model is schematised in Fig. 8. The structure of the model is consistent with vestibular

neuroanatomy, but it is constructed as simply as possible to facilitate the presentation, analysis, and interpretation of the pattern correlation mechanism. Because the VOR is a bilateral structure, the complete model is actually composed of two copies of the model shown in Fig. 8, one on either side. The sinusoidal input to the model is half-wave rectified, and only the positive or negative half-cycles are sent to the positive or negative sides of the model. The positive and negative inputs are shown in Fig. 9A,B, respectively (dot-dash lines). This separation of the two sides is motivated by the data, which clearly indicate that habituation occurs independently on each side (Dow & Anastasio, 1998). Each side of the model operates independently on its input, and the outputs are combined at the end.

The vestibulocerebellum, (VC) is necessary for VOR habituation to occur and be maintained (see Discussion). Both primary and secondary vestibular neurons send projections up to the VC, and the VC sends inhibitory projections back down to secondary vestibular neurons (Butler & Hodos, 1996). The model is subdivided into two parts, one representing the VC and the other representing the brainstem VOR pathway (BS). The vestibular nucleus (VN) computes the output of the model as the difference between its inputs from BS and VC. The structure and




Fig. 7. Discontinuous segments of a sinusoid constitute the patterns in the pattern correlation model. Each dot marks the value of a 0.01 Hz sinusoid at intervals of 1 s. Isoamplitude reference lines (dashed lines) are drawn at varying levels relative to the baseline, which cut the sinusoid into segments each 10 s long. This produces 10 distinct patterns labeled a-e and 1-5.

function of the vestibulocerebellar (VC) and brainstem (BS) parts of the model are described in detail below.

2.2.1. Brainstem model

VOR habituation is known to involve events occurring at the level of the VN, since the responses of the vestibular afferents (VAs) are not altered by habituation (Courjon, Precht, & Sirkin, 1987). To produce VOR habituation in the model, VC must remove from VN some amount of the vestibular signal to be habituated. For simplicity, the vestibular signal provided as input to the model is analogous to the output of the velocity storage mechanism. As such, the dynamic properties of the input to the model are more similar to those of the VOR than of the VAs.

The gain and phase of the goldfish VA signal changes a lot over the frequency spectrum (Fig. 1A,B, solid lines). In contrast, the naïve (unhabituated) goldfish VOR frequency response has nearly constant gain and zero phase over most of the spectrum (Fig. 4, dashed lines). This can be modeled as a direct (gain of one) connection to VN for practical purposes. Representing the brainstem VOR pathway (BS) as a direct connection implies that, in the absence of any influence from VC, the output of the model will be identical to the input. This facilitates the analysis of the model, because any difference between the output and the input must be due only to the influence of the pattern correlation

mechanism. The model could use either the more complicated VA signal or the less complicated VOR signal with equal validity. The simpler signal is used because it makes the operation of the model transparent.

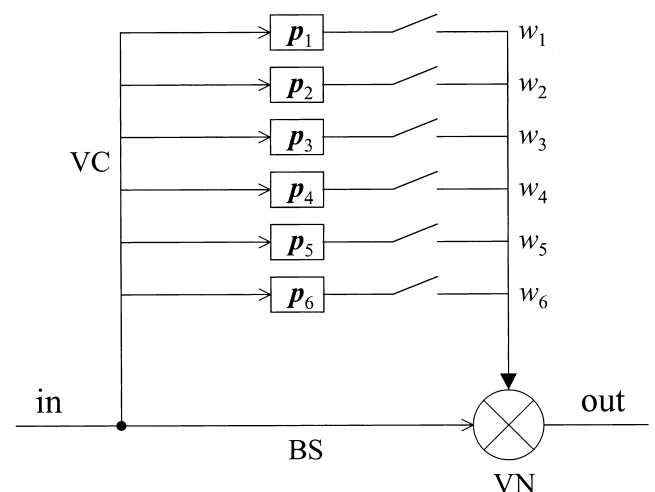


Fig. 8. Schematic diagram of the pattern correlation model. Open arrows, excitatory connections; closed arrow, inhibitory connection. BS, brainstem; VN, vestibular nucleus neuron; VC, vestibulocerebellum. The p_i and w_i ($i = 1, 2, \dots, 6$) are pattern vectors and weights, respectively.




Fig. 9. Operation of the pattern correlation model. The complete model consists of two copies of the model schematised in Fig. 8. Each copy is thought to represent one side of the real, bilateral VOR. A, half-wave rectified positive input (dot-dash line) and output (solid line) of the positive side of the model. B, half-wave rectified negative input (dot-dash line) and output (solid line) of the negative side of the model. C, combined output of the complete model (solid line), which is the sum of the outputs of the positive and negative sides following positive and negative rectification, respectively. The least-squares sinusoid fit to the combined output is also shown (dashed line). Numbers and letters mark the contribution to the output of each of the patterns shown in Fig. 7. D, the combined output is filtered and dithered (solid line). The least-squares sinusoid from C is reproduced for comparison (dashed line).

2.2.2. Vestibulocerebellar model

The VC mediates pattern correlation in the model. The VC part of the model stores six different patterns. Each pattern is associated with a switch and a weight. At any point in time one of the switches will be closed. The closed switch is the one whose pattern is best correlated with the input at the current time. That pattern then inhibits VN by the amount of its correlation with the input, multiplied by the assigned pattern weight. A more detailed description follows.

Each pattern represents a possible time history of the input to the VC. Every pattern is a vector with elements that are the values of the pattern at equally spaced time points. On the positive side of the model, the first five patterns (p_1 through p_5) are contiguous, non-overlapping segments taken from the positive half cycle of a sinusoid. The negative side of the model (not shown) stores five patterns taken from the negative half cycle of a sinusoid. The sixth pattern on either side (p_6) is a vector of positive ones. This pattern represents a constant input.

The model is considered habituated at a particular frequency when the patterns are taken from an input sinusoid at that frequency. The patterns for simulated habitua-

tion at 0.01 Hz are shown in Fig. 7. Patterns 1-5 constitute the first five patterns on the positive side, while patterns a-e constitute the first five patterns on the negative side of the model. Alternatively, patterns can be taken from an input sinusoid at 0.03, 0.05, or 0.1 Hz to simulate habituation at each of those frequencies. All input sinusoids have amplitude one. The same isoamplitude lines are used to separate each input sinusoid into ten patterns. Thus, all of the patterns at a given habituating frequency have the same duration, which is 1/10 of the period of the input sinusoid. For purposes of illustration, values for the sinusoid in Fig. 7 have a temporal spacing of 1 s, but the sinusoids used as the actual inputs to the model have spacing of 50 ms (0.01, 0.05, and 0.03 Hz) or 25 ms (0.1 Hz). All of the patterns at a given habituating frequency have the same number of elements n which are: 200 at 0.01 Hz, 60 at 0.03 Hz, 40 at 0.05 Hz, and 40 at 0.1 Hz.

At any time t , the input to the six patterns in the VC is a sliding window of n elements that advances along the input to the model, one value at a time. This vector $h(t)$ represents the recent history of the input. It contains the value of the input at the current time step t and at each of the previous $n - 1$ time steps. At each point in time, the correlation $r_i(t)$

between the history vector $\mathbf{h}(t)$ and each of the pattern vectors \mathbf{p}_i is computed as follows:

$$r_i(t) = \mathbf{h}(t)^T \mathbf{p}_i / \|\mathbf{h}(t)\| \|\mathbf{p}_i\| \quad (9)$$

where the numerator on the right-hand side is the inner product of the two vectors (superscript T denotes the transpose operation) and the denominator is the product of the norms of the two vectors. This method is commonly used in artificial neural systems for computing the correlation between two vectors (Haykin, 1999). Once the correlation between the history vector $\mathbf{h}(t)$ and each of the pattern vectors \mathbf{p}_i is computed, the index k of the pattern having the maximum correlation with the input history vector is found:

$$k = \arg \max_i (r_i(t)). \quad (10)$$

The output at time step t [$y(t)$] is computed by subtracting from the input [$x(t)$] the value of the maximum correlation, multiplied by the weight value w_k corresponding to the pattern k having the maximum correlation with the history vector at time step t :

$$y(t) = x(t) - w_k r_k(t). \quad (11)$$

The history vector continues to advance along the input one step at a time, and the output is computed at each step through application of Eqs. (9)–(11), until the end of the input is reached. The algorithm is run through one full cycle before results are taken to avoid artefacts due to history vector initialization. The process is carried out separately for the positive and negative half cycles of the input by the positive and negative sides of the model. This produces two separate outputs, which are combined after a step designed to model VN rectification.

The model focuses on signal processing occurring at the VN level. As in mammals (Baker, Evinger, & McCrea, 1981), VN neurons in goldfish are prone to rectification (Allum, Graf, Dichgans, & Schmidt, 1976; Allum & Graf, 1977). Thus, VNs respond with firing rate modulations to rotations in their excitatory directions, but have zero firing rate for rotations in their inhibitory directions. There is also evidence in mammals that VN rectification is associated with the responses to habituated stimuli (Kileny, Ryu, McCabe, & Abbas, 1980). VN rectification is simulated in the model by rectifying the outputs from the model. A detailed example of the operation of the model is given below.

2.2.3. Operation of the pattern correlation model

The outputs on the positive and negative sides for sinusoidal input at 0.01 Hz, from a model habituated at 0.01 Hz, are shown in Fig. 9A,B, respectively (solid lines). After processing is terminated, the output is rectified. All negative points are removed from the output on the positive side and vice versa for the negative side. The outputs from the two sides are then summed to produce the final output of the model. The final output of the model, habituated at 0.01 Hz, for sinusoidal input at 0.01 Hz is shown in Fig. 9C (solid line).

The pattern weights are set by hand, to produce roughly the amount of gain decrease observed following actual VOR habituation at the habituating frequency, with minimal phase difference (Fig. 4). Gain and phase of the model are determined as they are for the goldfish VOR, by fitting least-squares sinusoids to the data (see above). The least-squares sinusoid for the output at 0.01 Hz of the 0.01 Hz habituated model is shown in Fig. 9C (dashed line). The gain of the habituated model is 0.05. Since the gain of the naïve model is one, this represents a gain decrease due to habituation of 20 times. The phase of the model is near zero. A gain reduction of 20 times with a phase near zero is also observed for the goldfish VOR at 0.01 Hz following habituation at that frequency (see above).

The values of the weights for patterns 1–5 that produced habituation at 0.01 Hz are in order: 0.555, 0.915, 0.915, 0.640, and 0.100. The weight values for patterns a–e are the additive inverse of those for patterns 1–5. The weight for the sixth (constant) pattern is zero for both sides of the model. The weights for habituation at other frequencies (see below) are essentially scaled versions of the weights for habituation at 0.01 Hz. Differences in the pattern weight values largely reflect differences in the amplitudes of the input segments to which the patterns best correlate. They also reflect differences in the amounts by which the different input segments are eliminated from the habituated response. At the habituating frequency, each pattern in the sequence (Fig. 7) is best correlated with the input history vector until the input reaches the midpoint of the subsequent pattern (see below). As a consequence, input segments that best correlate with patterns 1, 4, and 5 have lower amplitudes than those that best correlate with patterns 2 and 3. The weights chosen eliminate three times as much of the lower as of the higher amplitude input segments. This is consistent with the hypothesis that lower amplitude input segments should habituate more because they are more uncertain (see Discussion).

When the input is at the frequency to which the model has been habituated, each pattern in turn will become maximally correlated with the history vector. Thus, the output will jump discontinuously from one level to another. This discontinuity is evident in the combined output shown in Fig. 9C (solid line). The various pieces of the response can be identified easily with the patterns that produced them. The number of the pattern that produced each piece is marked in Fig. 9C. On the positive side, pattern 1 produced the first, large nib. Patterns 2 and 3 produced the first and second halves of the hump, respectively. Patterns 4 and 5 produced the second, large, and third, small, nibs, respectively. Patterns a–e have the same relationship to the response pieces on the negative side. This easy ability to see the effects of individual patterns on the output is helpful in characterizing the nonlinear behavior of the model (see below). Superficially, however, the output of the model does not seem to resemble the response of the real goldfish VOR. It can easily be made to do so.

2.2.4. Filtering and dithering the model output

The pattern correlation model processes its input as a set of discontinuous fragments and applies them separately in forming the output. The output of the model is therefore a series of discontinuous fragments. Although discontinuities are also apparent in the data (e.g. Fig. 2B), those in the raw output of the model are sharper. The reasons for this difference probably involve physiological and experimental factors that affect the recorded response of the real VOR but are absent from the model.

The output of the model represents the VOR eye rotational velocity command that is transmitted by the vestibular nucleus neurons (VNs) to the motoneurons of the eye muscles, and is then used to produce rotations of the eye (see above). The real VN command is low-pass filtered by the neuromuscular mechanisms of the VOR and by the physiological properties of the eye muscles and eyeball before it is expressed as VOR eye rotation. The transduced eye rotation signal is deliberately low-pass filtered to prevent aliasing prior to digitisation (Dow & Anastasio, 1996, 1998, 1999a). Eye rotation is corrupted by noise introduced by the stochasticity inherent in neuromuscular mechanisms, and the eye rotation signal is corrupted by noise introduced in the experimental acquisition and processing of the data. By the time the VN eye rotation command is observed as a measured eye rotation, the filtering (low-pass) and dithering (added noise) would have smoothed out the sharp corners of the discontinuities that may have been present in it. The output of the model can be made superficially to resemble the VOR eye rotation data by low-pass filtering and dithering it.

For purposes of illustration, one model output time series is low-pass filtered and dithered for comparison with VOR eye rotational velocity data. The model output is digitally filtered using a sliding boxcar average that is 150 time points long:

$$y(t) = 1/m [x(t) + x(t - \Delta t) + x(t - 2\Delta t) + \dots + x(t - (m - 1)\Delta t)] \quad (12)$$

where $x(t)$ and $y(t)$ are the input to and output from the filter at time t , m is 150, and Δt is 50 ms. This filter is applied to the model output in both the forward and reverse directions to obviate phase distortions. Noise is drawn from a standard Gaussian distribution (mean = 0, variance = 1) and scaled by a factor of 0.1. The noise is first filtered using a sliding boxcar average that is 20 time points long (Eq. (12) with $m = 20$) before it is added to the filtered model output. Filtering and dithering are done for purposes of demonstration only, so the parameters are chosen arbitrarily. The result is cosmetically more similar to VOR eye rotational velocity.

Of the frequencies examined, habituation of the goldfish VOR is the most pronounced at 0.01 Hz, and simulated habituation at that frequency is the most discontinuous.

Thus, in terms of lack of superficial resemblance to the experimental data, the raw model output at 0.01 Hz represents the worst case. The raw output of the model at 0.01 Hz following simulated habituation at that frequency (Fig. 9C, solid line) is filtered and dithered and the result is shown in Fig. 9D (solid line). The same least-squares sinusoid fit to the raw model output in Fig. 9C (dashed line) is redrawn in Fig. 9D (dashed line) over the filtered and dithered output. The filtering and dithering process does not change the amplitude and phase angle estimate of the output, but it does provide it with a better cosmetic resemblance to the data. This can be appreciated by comparing the filtered, dithered output of the model habituated at 0.01 Hz (Fig. 9D, solid line) with the habituated VOR at 0.01 Hz (Fig. 5D, solid line). Filtering and dithering could be used in all cases to provide the model output with a better cosmetic resemblance to the goldfish VOR data. However, this superficial change would not affect the basic operation of the model, which is better appreciated by observing the raw model output.

3. Results

Generation, analysis, and interpretation of the output of the pattern correlation model provide the new results to be presented in this paper. The model is used to simulate the frequency-specific and nonlinear behaviors that are associated with habituation of the goldfish VOR.

3.1. Nonlinear behavior of the pattern correlation model

The nonlinear distortions observed during habituation of the goldfish VOR at 0.01 Hz can be simulated by changing the weights of the patterns in the model habituated at that frequency. This is consistent with the hypothesis that individual patterns habituate separately. To keep matters simple, pattern weights can take only one of two values. Either a weight takes its fully habituated value, as enumerated above for habituation at 0.01 Hz, or it takes the value of zero. An even closer fit between the model and the data could be obtained with arbitrary changes to the pattern weights, but this is not done for simplicity.

Setting all pattern weights to zero produces the naïve model output, which is identical to the input (Fig. 10, dot-dash lines in each plot). Nonlinear distortions observed during VOR habituation are modeled by setting a subset of the pattern weights to their fully habituated values. The effect on the model output time series of setting patterns 1, 5, a, and e to their full values is shown in Fig. 10A (solid line). The model output is decreased midrange but not at the peaks. An xy plot is constructed from this data (see above) and is shown in Fig. 11A. The xy plot has a dead-zone profile, which can be compared to the goldfish VOR in Fig. 3A. The dead-zone distortion was commonly observed during early stages of habituation at 0.01 Hz (Dow & Anastasio, 1998). The model suggests that patterns on the

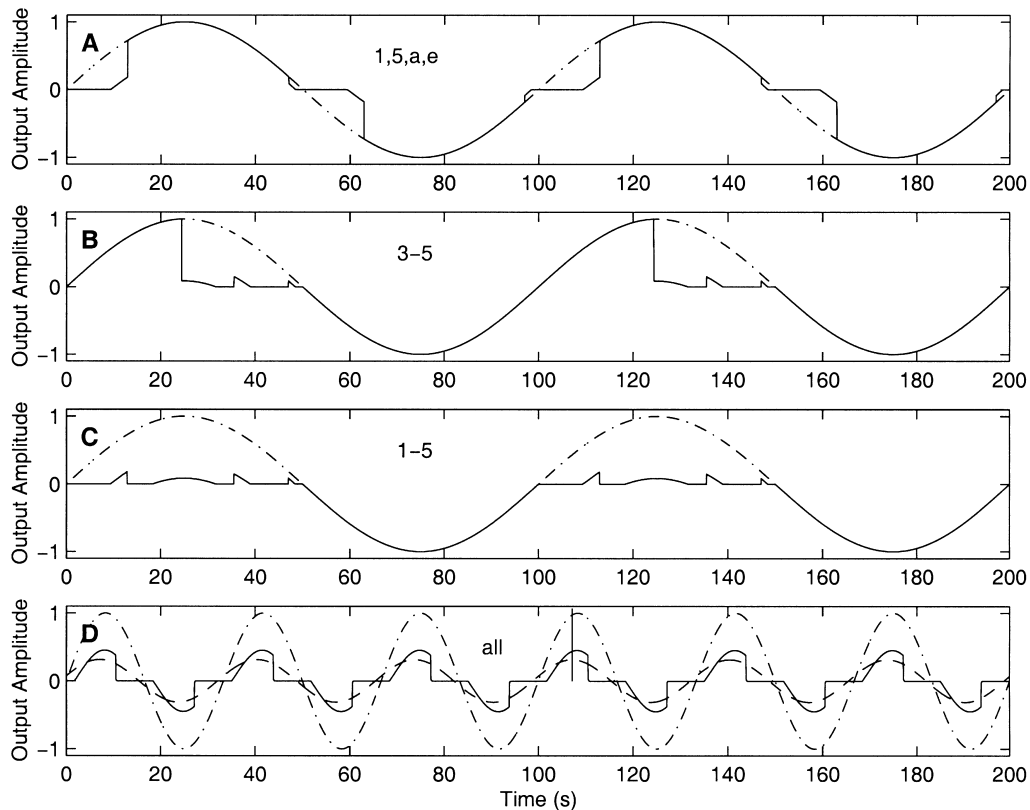


Fig. 10. Simulating VOR response distortion using the pattern correlation model. Pattern weights for the 0.01 Hz sinusoid are set either at zero or at their fully habituated values. A, fully weighting only patterns 1, 5, a, and e produces a distortion in which the shoulders of the output at 0.01 Hz are decreased. B, fully weighting only patterns 3 through 5 decreases the 0.01 Hz output for the second positive quarter cycle. C, fully weighting only patterns 1 through 5 decreases the 0.01 Hz output for the positive half cycle. D, fully weighting all 0.01 Hz patterns greatly decreases the shoulders of the 0.03 Hz output. The least-squares sinusoid fit to the 0.03 Hz output (dashed line) shows reduced amplitude relative to the input (dot-dash line in all plots). The phase of the least-squares sinusoid leads the input at 0.03 Hz (marked with a vertical rule).

shoulders of the sinusoid habituate first, perhaps because they are the closest to baseline and so are the most uncertain.

The effect on the model output time series of setting patterns 3, 4, and 5 to their full values is shown in Fig. 10B (solid line). Note that pattern 3 straddles the positive peak of the sinusoid, and patterns 4 and 5 are taken from the falling phase of the positive half cycle (Fig. 7). All pattern weights on the negative side are zero for this case, as are pattern weights 1 and 2 on the positive side. This weight configuration produces an abrupt decrease precisely at the peak of the positive half cycle. The model output resembles the abrupt decrease observed in the habituating VOR time series (Fig. 2B). An abrupt decrease at the peak is observed whenever pattern 3 is more habituated than pattern 2, regardless of the weights of the other patterns.

The reason for the abrupt decrease at the peak is that pattern 3, which straddles the peak, only becomes the most well correlated with the input after the input history vector has advanced enough to overlap pattern 3 more than the preceding pattern 2. In this case pattern 2 has zero weight, so it has no effect on model output. The input history vector overlaps pattern 3 one element more than pattern 2 right at the peak of the input sinusoid. At that point pattern 3

becomes the best correlated and decreases the output according to its nonzero weight. The xy plot for this case is shown in Fig. 11B, which can be compared to the goldfish VOR in Fig. 3B. Both the model output and the VOR show the figure-eight profile. The model simulates the abrupt decrease without invoking an unspecified switch.

The abrupt decrease follows as a consequence of pattern correlation and the hypotheses that patterns are defined on the basis of amplitude (distance from the baseline). Thus, the peak pattern would always straddle the peak of the input, because it would begin and end at the intersections of the input sinusoid and the uppermost isoamplitude reference line (Fig. 7). The rest of the VOR response half-cycle was always observed to be low following an abrupt decrease (Dow & Anastasio, 1998). The model suggests that when the peak pattern habituates then the falling phase patterns also habituate, perhaps because the amplitude of the sinusoidal input will only decrease and become more uncertain after the peak has been reached.

The effect on the model output time series of setting patterns 1 through 5 to their full values is shown in Fig. 10C (solid line). All pattern weights on the negative side are zero for this case. The model output is completely

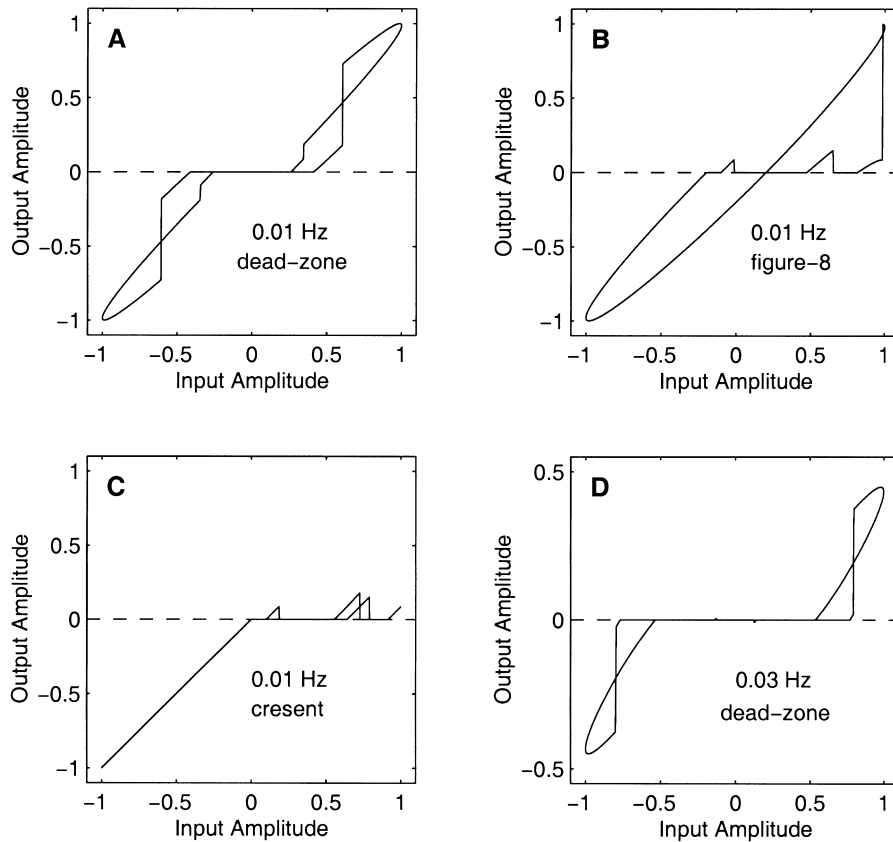


Fig. 11. Simulating the xy plots that characterize the nonlinear responses of the goldfish VOR during and after habituation at 0.01 Hz. A, dead-zone at 0.01 Hz during 0.01 Hz habituation. B, figure-eight at 0.01 Hz during 0.01 Hz habituation. C, crescent at 0.01 Hz during 0.01 Hz habituation. D, pronounced dead-zone at 0.03 Hz following 0.01 Hz habituation.

habituated on the positive side and naïve on the negative side. It appears as an extreme form of the asymmetrical habituation that is observed during habituation of the goldfish VOR at 0.01 Hz (Fig. 2B). A small difference between the model and the VOR is that asymmetrical VOR habituation is sometimes accompanied by a baseline shift (Dow & Anastasio, 1998). This is probably due to bilateral interactions in the real VOR that are excluded from the model for simplicity. The xy plot for this case is shown in Fig. 11C, which can be compared to the goldfish VOR in Fig. 3C. Both the model output and the VOR show the crescent profile. Setting all patterns to their full values produces full habituation at 0.01 Hz and an essentially flat xy plot (not shown).

Patterns perfectly correlate with their corresponding input segment when the input is at the frequency to which the model is habituated. Patterns in a model habituated at one frequency may correlate with segments of inputs at other frequencies, but to a more limited extent. The output of the model habituated at 0.01 Hz for input at 0.03 Hz is shown in Fig. 10D (solid line). The shoulders of the output are decreased but a relatively strong response remains at and around the peak. The xy plot for this case is shown in Fig. 11D, which can be compared to the goldfish VOR in Fig. 3D. Both the model output and the VOR show a pronounced

dead-zone. The model suggests that a pronounced dead-zone is observed in the VOR response at 0.03 Hz following habituation at 0.01 Hz because patterns stored for habituation at 0.01 Hz correlate strongly at the shoulders of the 0.03 Hz input but not at and around the peak.

Using the pattern correlation model to simulate the dead-zone and the crescent during habituation at 0.01 Hz is simply a matter of setting the weights of the appropriate patterns to full or zero. In contrast, the ability of the model to reproduce the abrupt decrease during 0.01 Hz habituation, and the pronounced dead-zone at 0.03 Hz following habituation at 0.01 Hz, are emergent properties of the model that are not set directly but follow from its assumptions and organization. Similar emergent properties allow the model to simulate the frequency-specific behavior associated with VOR habituation.

3.2. Frequency-specific behavior of the pattern correlation model

The model output, like the response of the real habituated VOR, is distorted and nonlinear. Despite this the model output, like the real habituated VOR, can be analysed as though it were linear by fitting least-squares sinusoids to it, and using the best-fit parameters to estimate gain and

phase (see above). This is illustrated for the output at 0.03 Hz of the model habituated at 0.01 Hz in Fig. 10D, where the least-squares sinusoid (dashed line) is plotted over the distorted output itself (solid line). An indication of how habituation at 0.01 Hz affects the gain and phase of the model at 0.03 Hz can be given by comparing the least-squares sinusoid fit to the output with the input (dot-dash line), and noting that the input is identical to the output of the naïve model. Relative to naïve, the gain (output amplitude/input amplitude) at 0.03 Hz is decreased. Also, a phase lead (positive phase difference) has been introduced. This is indicated by the vertical rule, which shows that the peak of the sinusoid best fit to the habituated response at 0.03 Hz precedes the peak of the naïve output at that frequency. It appears that the model habituated at 0.01 Hz has distorted the output at 0.03 Hz in such a way that a lead in the phase estimate has been introduced. A lead in the phase estimate relative to naïve is also observed at 0.03 Hz following habituation of the goldfish VOR at 0.01 Hz (Fig. 4B).

This analysis can be repeated at other frequencies, to observe the effect of habituating the model at 0.01 Hz on its output over the frequency spectrum. It can also be repeated over the spectrum after the model is habituated at other frequencies. As for the goldfish VOR, the model is habituated at 0.01, 0.03, 0.05, and 0.1 Hz, and tested over the spectrum at 0.01, 0.03, 0.05, 0.1, 0.3, 0.5, and 1.0 Hz. For simulated habituation at each frequency, the pattern weights are set to produce roughly the same amount of gain decrease at the habituating frequency, with minimal phase difference, as observed for the habituated goldfish VOR. The pattern weights for the 0.01 Hz habituated model are set to produce a gain decrease of 20 times at that frequency (see above). These weights are scaled to produce gain decreases of 5, 2, and 1.5 times at habituating frequencies of 0.03, 0.05, and 0.1 Hz, respectively. Some small adjustments in pattern weights are made to ensure minimal phase difference at the habituating frequency. Thus, the gain and phase of the model are set directly to match the observed values at the habituating frequency.

With the weights held fixed, the models habituated at each habituating frequency are then tested over the frequency spectrum. When the input is at a frequency other than the habituating frequency, the best-correlated patterns do not occur in order as they appear in Fig. 7. Instead, they can switch from one pattern to another and back again in no apparent order. Whether the input is at the habituating frequency or not, model outputs are discontinuous and nonlinear, but they can be analyzed as though they were linear by fitting least-squares sinusoids to them. Model gain and phase at each frequency can be computed using the parameters of the least-squares sinusoids fit to model output.

The simulated frequency responses are shown in Fig. 12 as Bode plots (gain, A, C, E, and G; and phase, B, D, F, and H). Simulated gain and phase for models habituated at

0.01 Hz (A and B), 0.03 Hz (C and D), 0.05 Hz (E and F), and 0.1 Hz (G and H) are shown as open circles in Fig. 12. The habituating frequency is indicated on each plot, and gain and phase at the habituating frequency are marked with crosses. As mentioned before the output of the naïve model is identical to the input for simplicity. The frequency response of the identity transfer function (the constant one) has gain one and phase zero over the frequency spectrum, and this is indicated in each plot with a dashed line.

The simulated frequency response data are fit with a transfer function in which a band-pass filter is subtracted from one:

$$C(s) = 1 - B(s). \quad (13)$$

This is analogous to the linear model fit to the habituated VOR (Eq. (6)), because the transfer function describing the dynamics of the brainstem VOR pathway has been simplified to the identity for the model. As for the experimental data, the transfer function of the simplified model is fit only to the phase characteristic of the simulated data (Eq. (7)). In fitting the transfer function, the optimisation routine is free to adjust the gain constant, central frequency, and spacing of the band-pass filter (g , f_c , and λ , Eqs. (4) and (5)). Parameter f_c is close to the frequency of simulated habituation in all fits. No consistent trends are observed in the g and λ parameters, because they both affect the gain of the band-pass filter and so the fitting problem is overdetermined. The point of the analysis is to show that the one – band-pass transfer function (Eq. (13)) does not fit the simulated frequency response data, so the precise values of the least-squares parameters are unimportant. Both the gain and the phase of the least-squares transfer function are reproduced in Fig. 12 (solid lines). The simulated and experimental (Fig. 4) habituated VOR frequency responses can be compared.

In both the simulated and experimental cases, gain is decreased most at the habituating frequency and is decreased progressively less for higher and lower frequencies. This produces a gain dip centered on the habituating frequency. The pattern correlation model reproduces the dip in gain observed in the VOR frequency response following single frequency VOR habituation. It suggests that the dip results because stored patterns that correlate well with segments from sinusoids at the habituating frequency correlate less well with segments from sinusoids at progressively higher and lower frequencies. The constant pattern correlates well with segments from sinusoids at frequencies much lower than the habituating frequency. When the constant pattern is the most well correlated it passes the input unaltered because its weight is zero.

The phase characteristics of the model also agree with the experimental data. In both the simulated and experimental cases, habituated phase leads naïve phase for frequencies higher than the habituating frequency, and habituated phase lags naïve phase for frequencies lower than the habituating frequency. This produces a crossover in phase centered on the habituating frequency. The model suggests

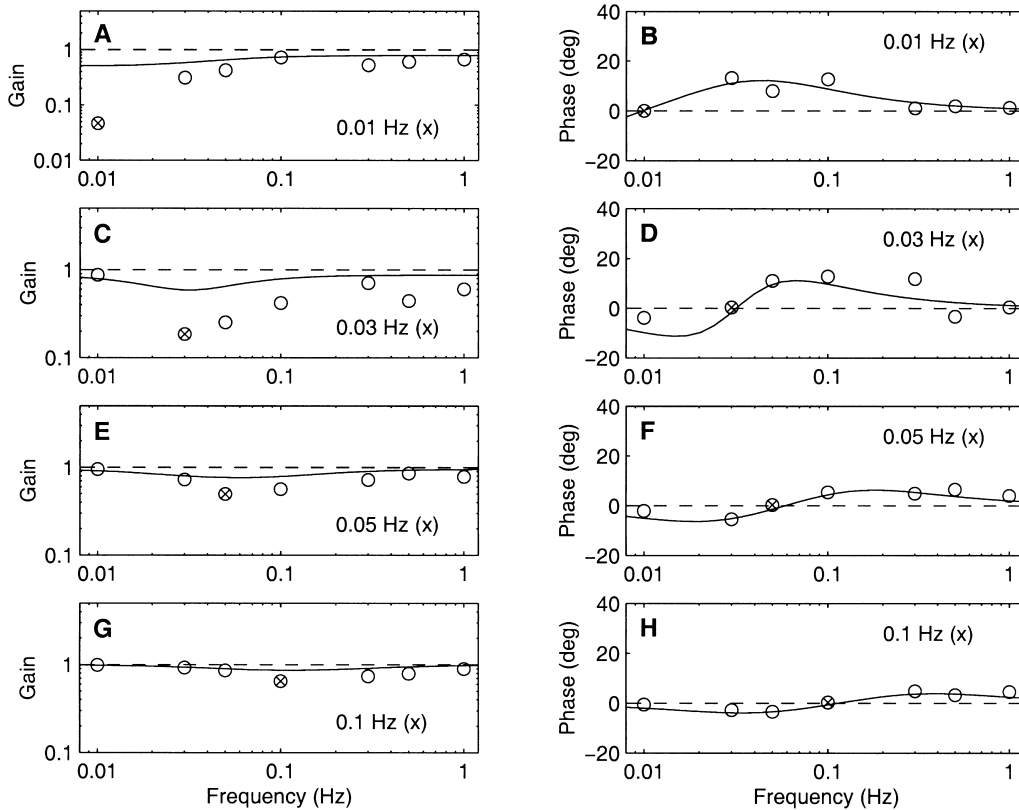


Fig. 12. Simulating the frequency responses of the goldfish VOR following habituation at various single frequencies. Gain plots (A, C, E, G) are log-log and phase plots (B, D, F, H) are log-linear. The habituating frequency is 0.01 Hz for A and B, 0.03 Hz for C and D, 0.05 Hz for E and F, and 0.1 Hz for G and H. Open circles represent gain and phase values estimated from the output of the model. The gain and phase estimates at the habituating frequency are marked with crosses. The dashed lines show the frequency response of the constant one (gain = 1 and phase = 0 at all frequencies) that is chosen for the naïve output of the model for simplicity. The solid lines show the frequency response of a transfer function in which a band-pass filter is subtracted from one. The one-band-pass transfer function (Eq. (13)) is fit (least squares) to model phase data only.

that phase crossover is due to a consistent pattern of response distortion that is produced by the pattern correlation mechanism, which causes the phase of the least-squares sinusoid to shift in characteristic ways. The pattern of distortion is robust, and produces the same form of phase crossover regardless of the habituating frequency or the amount of gain decrease at that frequency.

A noteworthy difference between the simulated and experimental frequency responses concerns the maximum deviation between habituated and naïve phase characteristics (maximum distance between open circles and dashed lines). For the real VOR, the maximum deviations between habituated and naïve phase are all of about the same size, regardless of the frequency of habituation or the amount of gain decrease observed at the habituating frequency (Fig. 4). In contrast, the maximum deviation between habituated and naïve phase in the model is proportional to the amount of gain decrease at the habituating frequency (Fig. 12). The reason for the difference may be that mechanisms of dynamic signal processing, which would affect real VOR phase (Dow & Anastasio, 1997), have been omitted from the model for simplicity.

The simulated and experimental, habituated VOR

frequency responses are similar not only in their gain and phase characteristics, but also in the way in which the gain characteristics fail to be described by the linear system model of frequency-specific habituation. The problem in using the linear system model (Eq. (6)) to fit goldfish VOR frequency responses following habituation at single frequencies is that it grossly overestimates gain at and near the habituating frequency (Fig. 4). In other words, the amount of gain decrease at the habituating frequency is much greater than expected on the basis of the relatively small shifts in phase, given the linear system model of frequency-specific habituation.

The pattern correlation model overcomes this problem. It is set to greatly decrease gain at the habituating frequency without affecting phase at that frequency. The gain dip occurs because patterns at the habituating frequency will correlate to some extent with segments at nearby frequencies. This produces distortions of the response at nearby frequencies that appear as comparatively small but consistent shifts in phase in the linear estimate. The corresponding linear system model for simulated habituation (Eq. (13)) provides a good fit to the simulated phase characteristic, but overestimates gain for all habituating frequencies

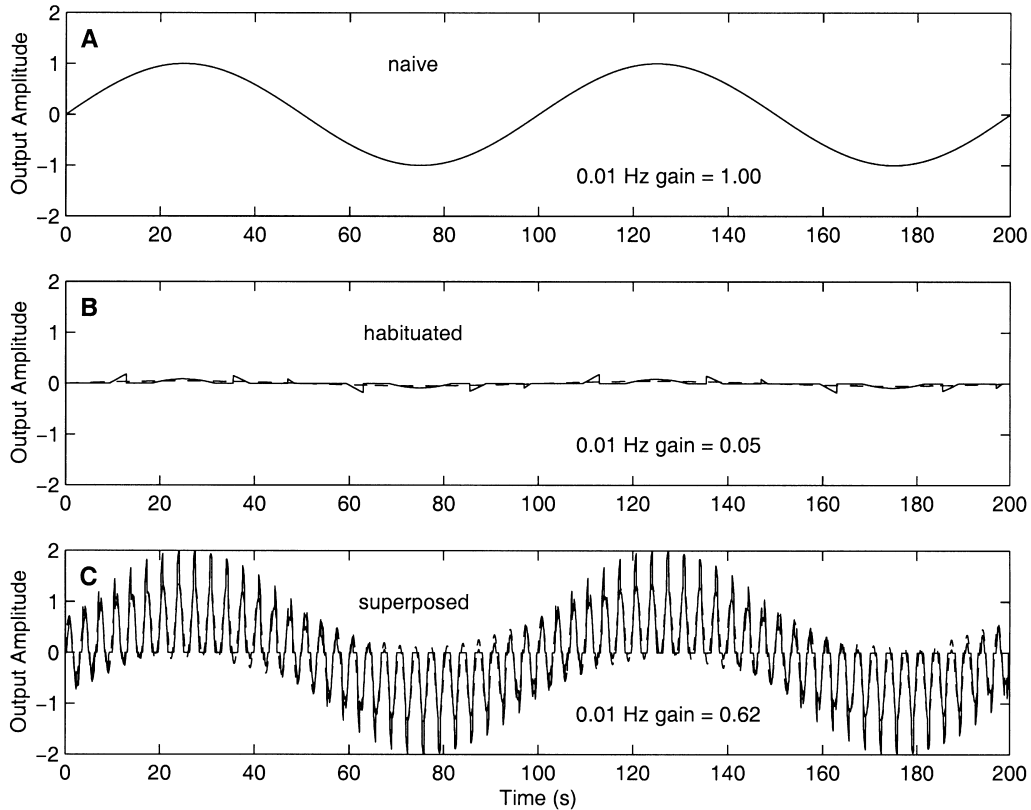


Fig. 13. Simulating the violation of superposition by the VOR in habituated goldfish. A, the naïve output of the model at 0.01 Hz (solid line), which is identical to the input. B, model output at 0.01 Hz following habituation at 0.01 Hz (solid line) and best fit, least-squares sinusoid (dashed line). C, model output for a combined input at 0.01 Hz and 0.3 Hz following habituation at 0.01 Hz (solid line) and least-squares sum of sinusoids (dashed line). The gain of the 0.01 Hz component is greatly increased over its habituated level by superposition of the 0.3 Hz component.

(Fig. 12). Also, the amount of gain overestimation is greater for lower habituating frequencies, for the simulated as for the experimental data. Thus, the pattern correlation model is capable of simulating the large decrease in gain, which can greatly exceed linear predictions based on relatively small shifts in phase, without invoking a hypothetical switch that further decreases the gain of the habituated VOR without affecting its phase.

3.3. Violation of superposition by the pattern correlation model

Experimental support for the hypothetical switch came from studies of superposition. These studies showed that the habituated goldfish VOR violates the superposition principle, by demonstrating that the gain of the VOR habituated at 0.01 Hz is immediately increased by an order of magnitude when the 0.01 Hz stimulus is combined with a stimulus at 0.3 Hz (Fig. 5). It seemed as though a switch was thrown that increased gain as soon as the stimulus deviated from the habituated stimulus. This hypothetical switch was also invoked to explain the nonlinear abrupt decrease, and the frequency-specific gain decrease exceeding linear predictions based on phase (see above). The pattern correlation model obviates the need for a hypothetical switch in those

two cases. It is also necessary to determine whether the pattern correlation model can reproduce the superposition results.

The superposition results are simulated by first setting the pattern correlation model for habituation at 0.01 Hz. Time series showing the naïve and habituated model responses are reproduced in Fig. 13A,B. Then the input was altered by combining the 0.01 Hz sinusoid with a sinusoid at 0.3 Hz. This is done simply by generating another sinusoid at 0.3 Hz and adding it to the 0.01 Hz sinusoid. The combined input is applied to the pattern correlation model and the output is shown in Fig. 13C. The output at 0.3 Hz is little affected by habituation of the model at 0.01 Hz, as expected from the frequency specificity of the model (Fig. 12A,B). However, the output at 0.3 Hz is modulated at 0.01 Hz by a substantially dishabituated response component at that frequency. A sum of two sinusoids (Eq. (8)), with frequencies at 0.01 Hz and 0.3 Hz, is fit to the combined response, and the linear gain estimate at 0.01 Hz is found to be 0.62. This represents an increase to about half the naïve level, and an increase over the habituated level of more than an order of magnitude, as for the goldfish VOR (Fig. 5).

Thus, the pattern correlation model reproduces the goldfish VOR superposition results. The model suggests that the habituated VOR fails to obey superposition because the

correlations with patterns taken from a single habituating sinusoid are degraded by combining the input sinusoid at the habituating frequency with a higher frequency sinusoid. This decreases the amount of inhibition due to habituation and results in a substantial dishabituation. Again, invocation of a hypothetical switch is not necessary.

4. Discussion

The frequency-specific and nonlinear behaviors associated with habituation of the goldfish VOR include phenomena that could not be explained using dynamic linear and static nonlinear models (Dow & Anastasio, 1996, 1998, 1999a). The unexplained phenomena are abrupt decreases at peak response, gain decreases far in excess of linear predictions based on phase, and violation of superposition. Their existence was attributed to a hypothetical switch that closed (or opened) at the appropriate time and in the appropriate context. The pattern correlation model explains how the frequency-specific and the nonlinear behaviors may be related, and how the apparent switching phenomena may occur.

The pattern correlation model is based on the hypothesis that the purpose of habituation is to decrease the VOR response to any repeated input pattern that is uncertain. Uncertain inputs are those with low signal to noise ratios. Since signal to noise ratio depends upon signal amplitude, the patterns correspond to segments of a sinusoid that differ from one another on the basis of amplitude. The pattern best correlated with the model input vector at any time then decreases model output by the value of its correlation multiplied by its weight. Thus, the amount of output reduction switches as one pattern and then another becomes the best correlated. Pattern correlation and selection processes in the model produce nonlinear and discontinuous model behavior that reproduces the experimental findings, including the phenomena that had been attributed to a hypothetical switch.

The model suggests that abrupt response decreases occur at the peak because the pattern that straddles the peak only becomes the best correlated after the input vector advances more than halfway into it, and that point is right at the peak. It suggests that the frequency-specific gain dip results because patterns from the habituating sinusoid, which produce potentially large gain decreases at the habituating frequency, may correlate to some extent with segments of inputs at nearby frequencies but not at more distant frequencies. It suggests that the phase crossover is produced by a pattern of distortion at nearby frequencies, which causes a relatively small but consistent pattern of phase shift as estimated by fitting least-squares sinusoids. The model suggests that violation of superposition results because the addition of a higher frequency input component degrades the correlation with patterns taken from a single habituating sinusoid.

Thus, the pattern correlation model reconciles the

frequency-specific and nonlinear behaviors associated with VOR habituation. It suggests that the hypothetical switch may be a consequence of pattern correlation and selection. The model offers an alternative to the continuous view of VOR habituation. It suggests that the habituation process operates discontinuously, recognizing and responding to fragments of signals such as sinusoids, in apparent disregard for their status in mathematics as elementary functions.

4.1. Vestibular input uncertainty and the purpose of VOR habituation

The pattern correlation model is based on the hypothesis that the purpose of habituation is to reduce the VOR response to any repeated input pattern that is uncertain. Presumably, it would be costly for the VOR to make a response to an uncertain input. The input to the VOR is the head rotational velocity signal that is transmitted by the vestibular afferents (VAs). The two most well established facts concerning the VAs are that their firing rates are stochastic, and that their dynamics are essentially high-pass (e.g. Fernandez & Goldberg, 1971; Goldberg & Fernandez, 1971a,b; Anastasio et al., 1985). This implies that VA discharge is noisy, in that it varies stochastically around some mean value that encodes the head rotational velocity signal being transmitted. Due to the high-pass dynamics of VAs, the peak amplitude of the signal decreases as frequency decreases. Consequently, as frequency decreases, the VA signal to noise ratio decreases, and the uncertainty of the input increases (Cover & Thomas, 1991). If the purpose of VOR habituation is to decrease the response of the VOR to uncertain input, then it might be expected that the amount of habituation would increase with input uncertainty. This appears to be the case for the goldfish, since both input uncertainty and the amount of VOR habituation increase as the frequency of habituation decreases (Fig. 4). Thus, the uncertainty hypothesis may explain why VOR habituation is frequency dependent.

4.2. Prediction of the pattern correlation model

Among the predictions that can be derived from the model, one nicely complements the results on violation of superposition. The patterns are taken from a sinusoid not only at a specific frequency but also at a specific amplitude. The model predicts that the gain of the habituated VOR should depend not only upon the frequency but also upon the amplitude of head rotation, even at the habituating frequency. Thus, the pattern correlation model predicts that the habituated VOR should also violate the homogeneity principle, which stipulates that the gain of a linear system should be independent of input amplitude (Oppenheim & Willsky, 1983). This prediction has been verified experimentally (Sudlow & Anastasio, 1999).

4.3. Habituation and adaptation of the vertebrate VOR

Habituation of the VOR has been observed in various vertebrates including rats, rabbits, cats, monkeys, and humans (Crampton, 1962; Kleinschmidt & Collewijn, 1975; Buettner, Henn, & Young, 1981; Clément, Courjon, Jeannerod, & Schmid, 1981; Jäger & Henn, 1981a,b; Tempia, Dieringer, & Strata, 1991). Other vertebrates also show nonlinear behavior associated with VOR habituation, especially response asymmetry (Crampton, 1962; Clément et al., 1981; Tempia et al., 1991). The pattern correlation model may be applicable to VOR habituation in other vertebrates. Further experimentation on the goldfish has shown that VOR habituation can be specific for two frequencies of head rotation presented simultaneously (Dow & Anastasio, 1999b). The pattern correlation model should be extendable to any repeated input including sums of sinusoids.

Experience-dependent modification of the VOR is usually studied under the adaptation paradigm, in which the normal relationship between visual and vestibular input is altered (Melvill Jones, 1985). Normally, the head rotates within a stationary visual surround and perfect VOR gain would be one. VOR gain can be adapted down or up by rotating the visual surround with or against the head, respectively. Adaptation and habituation of the VOR share certain features in common. For example, context-dependent switches have been identified in studies of VOR adaptation (Baker, Wickland, & Peterson, 1987; Shelhamer, Robinson, & Tan, 1991). The pattern correlation model may provide explanations for context-dependent switches observed in VOR adaptation studies.

Frequency-specific effects are also observed when VOR gain is adapted using sinusoidal visual/vestibular mismatch in cat, monkey, and goldfish (Godaux, Halleux, & Gobert, 1983; Lisberger et al., 1983; Schairer & Bennett, 1986; Raymond & Lisberger, 1996). When VOR gain is adapted down at a single frequency, a gain dip and phase crossover are observed in the VOR frequency response that are similar to those reported for goldfish habituated at single frequencies (Dow & Anastasio, 1999a). When VOR gain is adapted up at a single frequency, a gain bump is observed, and it is accompanied by a phase crossover that is flipped top-to-bottom relative to that due to down adaptation. The pattern correlation model is also capable of simulating the observed frequency-specific effects following either up or down adaptation (not shown). Applying the pattern correlation model to data on VOR adaptation is an area for future research.

4.4. Comparison with models of VOR adaptation

Many previous models of VOR adaptation concerned the site at which adaptive synaptic weight changes might be made (e.g. Lisberger, 1994). Other models of VOR adaptation are based on more abstract notions that include the linear adaptive filter (Fujita, 1982), the Kalman filter (Paulin, 1989), and other ideas from adaptive control and

neural networks (Gomi & Kawato, 1992). A recent study using an adaptive neural network model of the VOR directly addressed the phenomenon of frequency-specific adaptation (Arnold & Robinson, 1997). That model was trained using a modified delta-rule (Haykin, 1999) to increase the gain of the VOR network model at a single frequency. Following training, the roughly frequency-specific gain increase was similar to that which is observed behaviorally. Unfortunately, phase effects in the neural network model were not studied, so its ability to reproduce phase-crossover cannot be evaluated. In any case, examination of the trained neural network model provided no insight into how the frequency-specific effects associated with simulated VOR adaptation came about. In contrast, the pattern correlation model provides an explicit mechanism that may account for both the nonlinear and the frequency-specific effects associated with VOR habituation.

Current models of VOR adaptation make use of a feedback error signal known as retinal slip, which is produced during visual/vestibular mismatch and is thought to drive VOR adaptation. The retinal slip signal cannot be used in models of VOR habituation because habituation only occurs in the dark. Retinal slip drives VOR adaptation but prevents VOR habituation. The assumption of the pattern correlation model is that VOR habituation is driven by vestibular afferent input uncertainty. Retinal slip probably prevents habituation because it provides a signal from an independent sensory modality (i.e. vision) that also indicates head rotation. Availability of retinal slip could greatly reduce any uncertainty associated with the vestibular afferent signal, and thereby prevent VOR habituation.

4.5. The role of the cerebellum

Different regions of the mammalian vestibulocerebellum mediate VOR adaptation and habituation (Torte, Courjon, Flandrin, Magnin, & Magenes, 1994). Vestibular adaptation involves the flocculus (Robinson, 1976; Ito, 1982; Waespe, Cohen, & Raphan, 1983; Lisberger, Miles, & Zee, 1984), while vestibular habituation involves the nodulus and uvula (Waespe, et al., 1985, Cohen, Cohen, & Raphan, 1992). A vestibulocerebellum without distinct lobules has been identified in the goldfish (Pastor, de la Cruz, & Baker, 1994). Some amount of VOR adaptation can occur in vestibulocerebellectomized goldfish (*ibid.*). In contrast, VOR habituation can neither occur nor be maintained following vestibulocerebellectomy in the goldfish (Dow & Anastasio, 1996, 1998).

Experiments using brief pulses of head rotation show that the VOR begins responding within milliseconds, and before any influence from the vestibulocerebellum (VC) can be exerted on it (Lisberger et al., 1984; Pastor et al., 1994). The brainstem VOR pathway may respond to any VA signal that it receives, without regard for the certainty of the signal. Conversely, the VC, in mediating habituation, may evaluate the certainty of the VA signal. The VC could learn to

recognize patterns of activity that are associated with uncertain VA signals, and to decrease the response of the VOR by an amount appropriate for any particular pattern.

The pattern correlation mechanism is thought to reside in the VC. Most current models of the cerebellum are based on the Marr/Albus paradigm, in which the cerebellum can adaptively modify ongoing behavior on the basis of distinct input patterns that it can learn to recognize (Marr, 1969; Albus, 1971). The pattern correlation model of VOR habituation follows this basic paradigm. To simplify the model in its initial presentation, the pattern correlation mechanism has been stripped to its essentials. Many interesting issues are left for future work. One issue concerns the mechanism by which the VC could gauge the uncertainty of the vestibular input signal. Another concerns how the VC could learn temporal patterns of long duration.

Pattern durations in the model are arbitrary, but the patterns have to be long enough to ensure adequate specificity for fragments taken from a sinusoid at the frequency of simulated habituation. Patterns of 10 s and shorter are used in the model, and recent studies confirm that the goldfish vestibulo-optokinetic system can be adapted to patterns of 10 s and longer (Marsh & Baker, 1997). It has been suggested that the cerebellum can recognize temporal patterns of arbitrary length, for example, through waves of activity over parallel fibers (Braitenberg, Heck, & Sultan, 1997). Exactly how this happens in the VC remains to be determined.

Patterns may be stored in the VC on parallel fiber to Purkinje cell synapses. The input vector could be carried by a set of parallel fibers, and the correlations could be computed on Purkinje cells as the sums of the parallel fiber inputs weighted by the synapses encoding the patterns. Selection of the maximally correlated pattern could occur through lateral inhibition among Purkinje cells that produces selection via a winner-take-all mechanism (Llinas & Walton, 1990). Pattern synapses could be learned through a neurobiologically plausible strategy similar to that which has been used to stimulate learning in the smooth pursuit system (Kettner et al., 1997). In that study (ibid.), the cerebellum received multiple inputs that varied in their temporal response properties, and learned to recognize distributed input states that indicated the recent history of the pursuit stimulus. A similar strategy could be used to learn the recent history of head rotation inputs. Endowing the pattern correlation model with greater neurobiological detail is an important area for future research.

Acknowledgements

I thank Andrew Barto and Daniel Bullock for helpful discussions, and Ernst Dow, Aaron Seitz, and anonymous reviewers for comments on the manuscript. This research was conducted during a sabbatical at Boston University in the Center for Adaptive Systems. It was supported by the University of Illinois and NIH grant MH50577.

References

- Albus, J. S. (1971). A theory of cerebellar function. *Mathematical Biosciences*, *10*, 25–61.
- Allum, J. H. J., Graf, W., Dichgans, J., & Schmidt, C. L. (1976). Visual-vestibular interactions in the vestibular nuclei of the goldfish. *Experimental Brain Research*, *26*, 463–485.
- Allum, J. H. J., & Graf, W. (1977). Time constants of vestibular nuclei neurons in the goldfish: a model with ocular proprioception. *Biological Cybernetics*, *28*, 95–99.
- Anastasio, T. J., Correia, M. J., & Perachio, A. A. (1985). Spontaneous and driven responses of semicircular canal primary afferents in the unanesthetized pigeon. *Journal of Neurophysiology*, *54*, 335–347.
- Arnold, D. B., & Robinson, D. A. (1997). The oculomotor integrator: testing of a neural network model. *Experimental Brain Research*, *113*, 57–74.
- Baker, J., Wickland, C., & Peterson, B. (1987). Dependence of cat vestibulo-ocular reflex direction adaptation on animal orientation during adaptation and rotation in darkness. *Brain Research*, *408*, 339–343.
- Baker, R., Evinger, C., & McCrea, R. A. (1981). Some thoughts about the three neurons in the vestibular ocular reflex. *Annals of the New York Academy of Science*, *374*, 171–188.
- Braitenberg, V., Heck, D., & Sultan, F. (1997). The detection and generation of sequences as a key to cerebellar function: experiments and theory. *Behavioral and Brain Sciences*, *20*, 229–277.
- Buettner, U. W., Henn, V., & Young, L. R. (1981). Frequency response of the vestibulo-ocular reflex (VOR) in the monkey. *Aviation, Space and Environmental Medicine*, *52*, 73–77.
- Butler, A. B., & Hodos, W. (1996). *Comparative vertebrate neuroanatomy: evolution and adaptation*. New York: John Wiley and Sons.
- Clément, G., Courjon, J. H., Jeannerod, M., & Schmid, R. (1981). Unidirectional habituation of vestibulo-ocular responses by repeated rotational or optokinetic stimulations in the cat. *Experimental Brain Research*, *42*, 34–42.
- Cohen, H., Cohen, B., Raphan, T., & Waespe, W. (1992). Habituation and adaptation of the vestibuloocular reflex: a model of differential control by the vestibulocerebellum. *Experimental Brain Research*, *90*, 526–538.
- Courjon, J. H., Precht, W., & Sirkin, D. W. (1987). Vestibular nerve and nuclei unit responses and eye movement responses to repetitive galvanic stimulation of the labyrinth in the rat. *Experimental Brain Research*, *66*, 41–48.
- Cover, T., & Thomas, J. (1991). *Elements of information theory*. New York: Wiley.
- Crampton, G. H. (1962). Directional imbalance of vestibular nystagmus in cat following repeated unidirectional angular acceleration. *Acta Otolaryngologica*, *55*, 41–48.
- Dow, E. R., & Anastasio, T. J. (1996). Violation of superposition by the vestibulo-ocular reflex of the goldfish. *NeuroReport*, *7*, 1305–1309.
- Dow, E. R., & Anastasio, T. J. (1997). Induction of periodic alternating nystagmus in intact goldfish by sinusoidal rotation. *NeuroReport*, *8*, 2755–2759.
- Dow, E. R., & Anastasio, T. J. (1998). Analysis and neural network modeling of the nonlinear correlates of habituation in the vestibulo-ocular reflex. *Journal of Computational Neuroscience*, *5*, 171–190.
- Dow, E. R., & Anastasio, T. J. (1999a). Analysis and modelling of frequency specific habituation of the goldfish vestibulo-ocular reflex. *Journal of Computational Neuroscience*, *7*, 55–70.
- Dow, E. R., & Anastasio, T. J. (1999b). Dual-frequency habituation and dishabituation of the goldfish vestibulo-ocular reflex. *NeuroReport*, *10*, 1729–1734.
- Fernández, C., & Goldberg, J. M. (1971). Physiology of peripheral neurons innervating semicircular canals of the squirrel monkey. II. Response to sinusoidal stimulation and dynamics of peripheral vestibular system. *Journal of Neurophysiology*, *34*, 661–675.
- Fujita, M. (1982). Adaptive filter model of the cerebellum. *Biological Cybernetics*, *45*, 195–206.

- Godaux, E., Halleux, J., & Gobert, C. (1983). Adaptive change of the vestibulo-ocular reflex in the cat: the effects of a long-term frequency-selective procedure. *Experimental Brain Research*, 49, 28–34.
- Goldberg, J. M., & Fernández, C. (1971a). Physiology of peripheral neurons innervating semicircular canals of the squirrel monkey. I. Resting discharge and response to constant angular accelerations. *Journal of Neurophysiology*, 34, 635–660.
- Goldberg, J. M., & Fernández, C. (1971b). Physiology of peripheral neurons innervating semicircular canals of the squirrel monkey. III. Variations among units in their discharge properties. *Journal of Neurophysiology*, 34, 676–684.
- Gomi, H., & Kawato, M. (1992). Adaptive feedback control models of the vestibulocerebellum and spinocerebellum. *Biological Cybernetics*, 68, 105–114.
- Hartmann, R., & Klinke, R. (1975). System analysis of properties of primary vestibular fibers. *Experimental Brain Research Suppl.*, 23, 165.
- Hartmann, R., & Klinke, R. (1980). Discharge properties of afferent fibers of the goldfish semicircular canal with high frequency stimulation. *Pflügers Archiv*, 388, 111–121.
- Haykin, S. (1999). *Neural networks: a comprehensive foundation*, Upper Saddle River: Prentice Hall.
- Ito, M. (1982). Cerebellar control of the vestibulo-ocular reflex—around the flocculus hypothesis. *Annual Review of Neuroscience*, 5, 275–296.
- Jäger, J., & Henn, V. (1981a). Habituation of the vestibulo-ocular reflex (VOR) in the monkey during sinusoidal rotation in the dark. *Experimental Brain Research*, 41, 108–114.
- Jäger, J., & Henn, V. (1981b). Vestibular habituation in man and monkey during sinusoidal rotation. In B. Cohen, *Vestibular and oculomotor physiology: international meeting of the Barany Society* (pp. 330–339). New York: New York Academy of Sciences.
- Keng, M. J., & Anastasio, T. J. (1997). The horizontal optokinetic response of the goldfish. *Brain, Behavior and Evolution*, 49, 214–229.
- Kettner, R. E., Mahamud, S., Leung, H.-C., Sitkoff, N., Houk, J. C., Peterson, B. W., & Barto, A. G. (1997). Prediction of complex two-dimensional trajectories by a cerebellar model of smooth pursuit eye movement. *Journal of Neurophysiology*, 77, 2115–2130.
- Kileny, P., Ryu, J. H., McCabe, B. F., & Abbas, P. J. (1980). Neuronal habituation in the vestibular nuclei of the cat. *Acta Otolaryngologica*, 90, 175–183.
- Kleinschmidt, H. J., & Collewijn, H. (1975). A search for habituation of vestibulo-ocular reactions to rotary and linear sinusoidal accelerations in the rabbit. *Experimental Neurology*, 47, 257–267.
- Llinas, R. R., & Walton, K. D. (1990). Cerebellum. In G. M. Shepherd, *The synaptic organization of the brain* (pp. 214–245). New York: Oxford University Press.
- Lisberger, S. G., Miles, F. A., & Optican, L. M. (1983). Frequency-selective adaptation: evidence for channels in the vestibulo-ocular reflex? *Journal of Neuroscience*, 3, 1234–1244.
- Lisberger, S. G., Miles, F. A., & Zee, D. S. (1984). Signals used to compute errors in monkey vestibuloocular reflex: possible role of flocculus. *Journal of Neurophysiology*, 52, 1140–1153.
- Lisberger, S. G. (1994). Neural basis for motor learning in the vestibulo-ocular reflex of primates: III. Computational and behavioural analysis of the sites of learning. *Journal of Neurophysiology*, 72, 974–998.
- Marr, D. (1969). A theory of cerebellar cortex. *Journal of Physiology, London*, 202, 437–470.
- Marsh, E., & Baker, R. (1997). Normal and adapted visuomotor reflexes in goldfish. *Journal of Neurophysiology*, 77, 1099–1118.
- Mayne, R. (1950). The dynamic characteristics of the semicircular canals. *Journal of Comparative Physiology and Psychology*, 43, 304–319.
- Melville Jones, G. (1985). Adaptive modulation of VOR parameters by vision. In A. Berthoz & G. Melville Jones, *Adaptive mechanisms in gaze control* (pp. 21–50). New York: Elsevier.
- Miles, F. A., Optican, L. M., & Lisberger, S. G. (1985). An adaptive equalizer model of the primate vestibulo-ocular reflex. In A. Berthoz & G. Melville Jones, *Adaptive mechanism in gaze control* (pp. 313–326). New York: Elsevier.
- Oppenheim, A. V., & Willsky, A. S. (1983). *Signals and systems*, Englewood Cliffs: Prentice Hall.
- Pastor, A. M., de la Cruz, R. R., & Baker, R. (1994). Cerebellar role in adaptation of the goldfish vestibuloocular reflex. *Journal of Neurophysiology*, 72, 1383–1394.
- Paulin, M. G. (1989). A Kalman-filter theory of the cerebellum. In M. A. Arbib & S. Amari, *Dynamic interactions in neural networks: models and data* (pp. 241–259). Berlin: Springer-Verlag.
- Raphan, Th., Matsuo, V., & Cohen, V. (1979). Velocity storage in the vestibulo-ocular reflex arc (VOR). *Experimental Brain Research*, 35, 229–248.
- Raymond, J. L., & Lisberger, S. G. (1996). Behavioral analysis of signals that guide learned changes in the amplitude and dynamics of the vestibulo-ocular reflex. *Journal of Neuroscience*, 16, 7791–7802.
- Robinson, D. A. (1976). Adaptive gain control of vestibuloocular reflex by the cerebellum. *Journal of Neurophysiology*, 39, 954–969.
- Robinson, D. A. (1981). The use of control systems analysis in the neurophysiology of eye movements. *Annual Review of Neuroscience*, 4, 463–503.
- Schairer, J. O., & Bennett, M. V. L. (1986). Changes in gain of the vestibulo-ocular reflex induced by combined visual and vestibular stimulation in goldfish. *Brain Research*, 373, 164–176.
- Schmidt, R., & Jeannerod, M. (1985). Vestibular habituation: an adaptive process? In A. Berthoz & G. Melville Jones, *Adaptive mechanism in gaze control* (pp. 113–122). New York: Elsevier.
- Segal, B. N., & Outerbridge, J. S. (1982). Vestibular (semicircular canal) primary neurons in bullfrog: nonlinearity of individual and population response to rotation. *Journal of Neurophysiology*, 47, 545–562.
- Shelhamer, M., Robinson, D. A., & Tan, H. S. (1991). Context-specific gain switching in the human vestibulo-ocular reflex. *Annals of the New York Academy of Sciences*, 656, 889–891.
- Sudlow, L. C., & Anastasio, T. J. (1999). Violation of homogeneity by the vestibulo-ocular reflex of the goldfish. *NeuroReport*, 10, 3881–3885.
- Tempia, F., Dieringer, N., & Strata, P. (1991). Adaptation and habituation of the vestibulo-ocular reflex in intact and inferior olive-lesioned rats. *Experimental Brain Research*, 86, 568–578.
- Torte, M. P., Courjon, J. H., Flandrin, J. M., Magnin, M., & Mageses, G. (1994). Anatomical segregation of different adaptive processes within the vestibulocerebellum of the cat. *Experimental Brain Research*, 99, 441–454.
- Waespe, W., Cohen, B., & Raphan, T. (1983). Role of the flocculus and paraflocculus in optokinetic nystagmus and visual-vestibular interactions: effects of lesions. *Experimental Brain Research*, 50, 9–33.
- Waespe, W., Cohen, B., & Raphan, T. (1996). Dynamic modification of the vestibulo-ocular reflex by the nodulus and uvula. *Science*, 228, 199–202.
- Weissenstein, L., Ratnam, R., & Anastasio, T. J. (1996). Vestibular compensation in the horizontal vestibulo-ocular reflex of the goldfish. *Behavioral Brain Research*, 75, 127–137.
- Wilson, V., & Melville Jones, G. (1979). *Mammalian vestibular physiology*, New York: Plenum Press.