J Comp Physiol A (1995) 177: 723–736

ORIGINAL PAPER

M. E. Nelson · M. G. Paulin

Neural simulations of adaptive reafference suppression in the elasmobranch electrosensory system

Accepted: 18 June 1995

Abstract The electrosensory system of elasmobranchs is extremely sensitive to weak electric fields, with behavioral thresholds having been reported at voltage gradients as low as 5 nV/cm. To achieve this amazing sensitivity, the electrosensory system must extract weak extrinsic signals from a relatively large reafferent background signal associated with the animal's own movements. Ventilatory movements, in particular, strongly modulate the firing rates of primary electrosensory afferent nerve fibers, but this modulation is greatly suppressed in the medullary electrosensory processing nucleus, the dorsal octavolateral nucleus. Experimental evidence suggests that the neural basis of reafference suppression involves a common-mode rejection mechanism supplemented by an adaptive filter that fine tunes the cancellation. We present a neural model and computer simulation results that support the hypothesis that the adaptive component may involve an anti-Hebbian form of synaptic plasticity at molecular layer synapses onto ascending efferent neurons, the principal output neurons of the nucleus. Parallel fibers in the molecular layer carry a wealth of proprioceptive, efference copy, and sensory signals related to the animal's own movements. The proposed adaptive mechanism acts by canceling out components of the electrosensory input signal that are consistently correlated with these internal reference signals.

Key words Adaptive filter · Anti-Hebbian synaptic plasticity · Computer simulation · Dorsal octavolateral nucleus · Neural model

M. E. Nelson (✉)
Department of Physiology and Biophysics and Beckman Institute for Advanced Science and Technology, University of Illinois, Urbana-Champaign, USA

M. G. Paulin
Department of Zoology and Centre for Neuroscience, University of Otago, Dunedin, New Zealand

Abbreviations *AEN* ascending efferent neuron · *AFF* primary afferent nerve fiber · *DGR* dorsal granular ridge · *DON* dorsal octavolateral nucleus · *ELL* electrosensory lateral line lobe · *GABA* γ-aminobutyric acid · *IN* inhibitory interneuron · *ISI* interspike interval · *ST* stellate cell

Introduction

Reafference refers to the component of sensory input that an animal receives as a consequence of its own movements or actions (von Holst and Mittelstaedt 1950). For example, movements that change the position or orientation of the eye in space give rise to visual reafference as a result of image motion across the retina; swimming movements in fish give rise to mechanosensory reafference due to lateral line excitation. Some sensory systems are specialized for extracting information from reafferent signals, as in the case of active sensory systems like bat echolocation or fish electrolocation, where a motor action (a vocal emission or an electric organ discharge) gives rise to a reafferent sensory signal that conveys useful information about the external environment. In other sensory systems, such as the passive electric sense discussed here, the reafferent component of the signal is largely a contaminating background which, unless it is somehow suppressed, would interfere with the detection of weak extrinsic signals of interest.

Elasmobranchs (sharks, skates, and rays) possess an electrosensory system that is remarkably sensitive to weak electric fields (Bodznick and Boord 1986). Elasmobranchs use their electric sense to detect the bioelectric fields generated by prey and to detect small voltage gradients of geomagnetic and geochemical origin for navigation (Murray 1960; Kalmijn 1974; Paulin 1995). The amazing sensitivity of this system has been

demonstrated in behavioral experiments in which sharks were induced to strike at artificial electric dipoles that mimic prey (Kalmijn 1982). Based on the distance at which orienting responses were observed, it was determined that voltage gradients as small as 5 nV/cm could elicit a behavioral response. Electric fields are detected by specialized electroreceptor organs called ampullae of Lorenzini (Murray 1974). An individual ampullary organ consists of a subdermal cluster of receptor cells connected to a skin pore by a jelly-filled canal. Ampullary organs are grouped in clusters beneath the skin with the canals radiating out from each cluster in many directions.

The extreme sensitivity of elasmobranch ampullary organs presents a significant signal processing challenge for the animal. Because the receptor organs are so sensitive, their output is strongly modulated by the bioelectric fields associated with the animal's own movements. Ventilatory movements, in particular, have been shown to strongly modulate the firing rates of primary electrosensory afferent nerve fibers (Montgomery 1984b; New and Bodznick 1990). The degree of ventilatory modulation at the level of primary afferents is striking. In some cases, electrosensory afferents are driven through a significant fraction of their full dynamic range (typically, 0–80 spikes/s) on each ventilatory cycle. For comparison, a threshold behavioral stimulus is expected to cause a change of much less than 1 spike/s in the firing rate of an individual afferent fiber. How does the nervous system manage to extract useful information about weak external sources in the presence of an overpowering reafferent signal due to ventilation?

The first hints of an answer were provided by neurophysiological experiments which demonstrated that the ventilatory modulation which dominates primary afferent activity is greatly suppressed in certain secondary sensory neurons in the dorsal octavolateral nucleus (DON), the region of the medulla that is responsible for the initial processing of electrosensory information in the central nervous system (Montgomery 1984b; New and Bodznick 1990). Furthermore, while these secondary neurons were less sensitive to ventilatory modulation, they were shown to be more sensitive to external electrical stimuli than were the primary afferent fibers. Thus the neural processing associated with the DON circuitry appears to be effective in suppressing the reafferent signal and enhancing the sensitivity to external signals. A simplified schematic and brief description of the relevant DON circuitry is provided in Fig. 1.

The first round of experimental data on ventilatory suppression in the DON obtained by Montgomery (1984b) and New and Bodznick (1990) appeared to be consistent with a relatively simple common mode rejection mechanism acting at the level of the DON. This mechanism, originally suggested by Kalmijn (1974), proposes that components of the signal that are common across afferents can be effectively suppressed by combining signals differentially in the central nervous system, such that the signal from one afferent (or group of afferents) is subtracted from another, thus suppressing common mode signals in a manner analogous the operation of a differential amplifier. It has been demonstrated that ventilatory modulation is, to good approximation, a common mode signal across the entire afferent population. The amplitude and phase of the modulation is nearly identical in all afferents, independent of position and orientation of the peripheral receptor organs (Montgomery 1984b; New and Bodznick 1990). This uniformity of influence is presumed to come about because ventilatory modulation is thought to primarily influence the internal body potential of the animal which is common to all the ampullary organs. Thus the reaffernt ventilatory signal is a good candidate to be removed by a common mode rejection mechanism. Subsequent to their initial experiments, Bodznick and Montgomery have obtained further experimental evidence that support a form of common mode rejection acting at the level of the DON (Bodznick et al. 1992; Bodznick and Montgomery 1992; Montgomery and Bodznick 1993).

The most recent experimental results from this system suggest that ventilatory suppression in the DON actually involves a more powerful signal processing principle than simple common mode rejection. Experiments reported recently by Bodznick (1993) and Montgomery and Bodznick (1994) indicate that ventilatory suppression also has an adaptive filtering component. To demonstrate this effect, experiments were conducted in which an external electrical stimulus was repeatedly paired with the fish's own ventilatory movements over a number of consecutive ventilatory cycles. Secondary neurons in the DON, called ascending efferent neurons (AENs), initially gave a strong response to the external stimulus, but continued pairing of the stimulus at a fixed phase in the ventilatory cycle caused this response to gradually diminish over the course of several minutes. If the external stimulus was subsequently removed after several minutes of pairing, the AEN exhibited a "negative image" of the original response pattern, suggesting that the adaptive filtering may involve an additive mechanism. Bodznick (1993) has shown that this adaptive suppression had temporal and spatial specificity; pairing of an external stimulus only affects AEN responsiveness in a limited temporal window and only when the stimulus is presented within the AEN's spatial receptive field. Thus the observed reduction in AEN responsiveness to the external stimulus is not due to some global adaptation or suppression taking place in the DON.

The general features of adaptive ventilatory suppression in elasmobranchs are strikingly similar to adaptive reafference suppression phenomena that have been observed in the electrosensory system of weakly electric teleosts. Bell (1981, 1982) was the first to report and characterize such a phenomenon in his work on

a modifiable efference copy mechanism in the ampullary region of the electrosensory lateral line lobe (ELL) of weakly electric mormyrids. Bell and colleagues have recently obtained experimental evidence suggesting that the adaptive changes they observe may be mediated by an anti-Hebbian form of synaptic plasticity in the ELL (Bell et al. 1993). Bastian (1994) has recently reported adaptive suppression of reafferent responses in the ELL of weakly electric gymnotids associated with changes in trunk and tail position. Montgomery and Bodznick (1994) also report a similar adaptive phenomenon in the lateral line mechanosensory system of the scorpion fish. The associated medullary sensory nuclei in these systems have striking parallels in terms of their structural and functional organization, suggesting that adaptive reafference suppression may involve similar neural mechanisms in all of these cases (Montgomery et al. 1995)

The goal of the modeling work presented here is to gain a better understanding of the signal processing principles and associated neural mechanisms that mediate reafference suppression in the elasmobranch DON. We present a simplified model of the DON circuitry that is sufficient to capture the essential features of both the common mode rejection and adaptive filtering aspects of the experimental data. In terms of mechanism, these two aspects of ventilatory suppression appear to have separate neural substrates. Common mode rejection seems likely to be mediated by a class of inhibitory interneurons in the central zone of the DON that are monosynaptically activated by electrosensory afferents, while adaptive filtering seems to involve a descending parallel fiber pathway that carries proprioceptive and efference copy signals related to the animal's ventilatory movements. From a signal processing point of view, however, we suggest that both the common mode rejection and adaptive filtering aspects of ventilatory suppression can be understood within a unified framework that sees the DON as a filter that removes predictable regularities from electrosensory data.

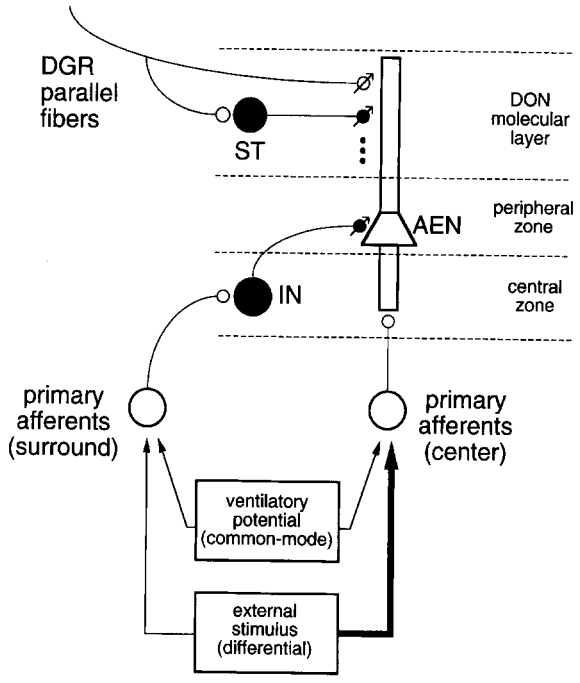

Fig. 1 A simplified schematic representation of the neural circuitry of the elasmobranch dorsal octavolateral nucleus (*DON*). Primary afferent nerve fibers (AFF) from peripheral ampullary organs project to the DON, where they contact ascending efferent neurons (*AEN*), the principal output neurons of the nucleus. Afferent fibers making direct synaptic contacts onto the AEN give rise to an excitatory receptive field "center". Afferents that influence the AEN via an interposed inhibitory interneuron (*IN*) give rise to an inhibitory "surround". AEN receptive fields typically have a single, focal excitatory center comprising a small number of ipsilateral ampullary organs. The spatial organization of the inhibitory "surround" is more complex, possibly including one or more focal inhibitory subfields as well as broad, diffuse subfields. Both focal and diffuse inhibitory subfields can include contralateral components. AENs also receive parallel fiber inputs from an overlaying mass of granule cells known as the dorsal granular ridge (*DGR*). DGR fibers synapsing directly on the AEN have an excitatory influence. DGR inputs can also have an inhibitory influence on AENs via stellate (*ST*) interneurons. *Arrows* drawn through synapses indicate possible sites of synaptic plasticity in our model

Methods

The results presented here are based on computer simulation studies of the schematic circuitry of the elasmobranch dorsal octavolateral nucleus (DON) shown in Fig. 1. The simulation software was written in C++ and simulations were carried out on Sun SPARCstation 10/30 and DECstation 5000/200 Unix workstations. The structure of the model is described below along with references to experimental data used to constrain the model parameters.

General model structure and convergence ratios

We model a single ascending efferent neuron (AEN) and its associated inputs as illustrated in Fig. 1. The elements of the model and their connections are summarized as follows: *Ascending efferent neuron* (*AEN*) : 1 total, receiving 64 direct primary afferent (AFF) inputs, 16 inhibitory interneuron (IN) inputs, 32 stellate (ST) inputs and 256 dorsal granular ridge (DGR) parallel fiber inputs; *Primary electrosensory afferents* (*AFF*) : 192 total, 64 converging directly onto the AEN plus 16 groups of 8 converging onto 16 inhibitory interneurons; *Inhibitory interneurons* (*IN*) : 16 total, each receiving input from 8 different primary afferents, all 16 INs converge onto the single AEN; *DGR parallel fibers*: 256 total, converging onto the AEN; of the thousands of parallel fibers that converge onto an AEN, the 256 fibers in the model are intended to represent only the subset that carries signals related to ventilatory movements, not the entire population. *Stellate cells* (*ST*) : 32 total, converging onto the AEN; we lack anatomical and physiological data to constrain the convergence of parallel fibers onto stellate cells and the stellate cell input-output relationship, so we therefore model only the stellate cell outputs. In general, the numbers of elements and convergence ratios in the model are consistent with expectations for the actual DON circuitry (D. Bodznick and J. Montgomery, personal communication).

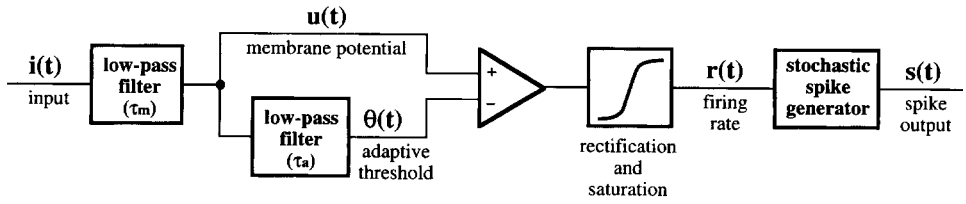

Fig. 2 Structure of the adaptive threshold neuron model used in these simulations. The model includes the effects of low-pass filtering by the membrane time constant, firing rate adaptation in response to sustained stimuli, firing rate rectification, firing rate saturation and stochastic spike generation. The mathematical description of the model is given in equations 1a–c and 2a–b

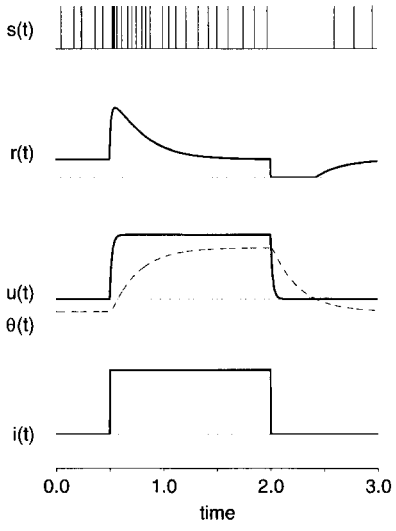

Fig. 3 Typical response of the adaptive threshold neuron model to a step change in input. The membrane potential $u(t)$ is a low-pass filtered version of the input $i(t)$ with a filter time constant τ_m. The threshold $\theta(t)$ is low-pass filtered version of the membrane potential $u(t)$, with an adaptation time constant τ_a which is typically longer than the membrane time constant τ_m. The difference between the membrane potential and threshold is passed through a sigmoidal nonlinearity to give rise to a continuous-valued representation of the instantaneous firing rate $r(t)$. The firing rate is converted into a train of action potentials by a stochastic spike generating element. Parameter values: $G = 1.0$, $\tau_m = 0.020$ s, $A = 1.0$, $\tau_a = 0.250$ s, $\theta_0 = -0.2$, $\lambda = 1.0$, $R_{max} = 50.0$ Hz, $n_\gamma = 8$

General modeling approach

Before describing the modeling methods in detail, we provide a brief conceptual overview of the general modeling approach that we have taken. The focus of this study is on understanding information processing principles that arise from network level interactions in the neural circuitry of the DON. Thus we are primarily concerned with the way in which neural and synaptic components interact with one another, rather than with their intrinsic properties. For example, our model will not specify what particular biophysical mechanisms might underlie adaptation in DON neurons, or what mechanisms might contribute to the regularity of firing in primary afferent fibers. This is not to imply that the biophysical details of neuronal function are inconsequential, but rather that they are not the primary focus of the questions that we are addressing here. In selecting a level of abstraction to describe the neural and synaptic components of our model, we choose concise phenomenological descriptions that capture the essential dynamic characteristics without specifying details of the underlying mechanisms that do not directly impact network level interactions. This approach allows us to more easily identify and isolate the key features of the model that directly relate to the ability of the DON circuitry to suppress reafferent components of the input signal.

General neuron model

The structure of the neuron model used in these simulations is shown in Fig. 2 and a typical response of the model to a step stimulus is illustrated in Fig. 3. The input to the model is a time-dependent variable $i(t)$, that can be thought of as a current. The output is a binary spike train $s(t)$, where $s(t) = 1$ is taken to represent the occurrence of an action potential at time t and $s(t) = 0$ otherwise. Internally, the model has three state variables: a membrane potential $u(t)$, an adaptive threshold $\theta(t)$ and an instantaneous firing rate $r(t)$. The model has a total of eight parameters: a gain parameter G that is analogous to neuronal input resistance, a membrane time constant τ_m, an adaptation parameter A that controls the degree of adaptation ($A = 0$ is non-adapting; $A = 1$ is fully adapting), an adaptation time constant τ_a, a resting threshold level θ_0 (i.e., θ_0 is the steady-state value of the threshold in the absence of input), a maximum firing rate R_{max}, a slope parameter λ that controls the steepness of a sigmoidal nonlinearity that incorporates the effects of firing rate rectification and saturation, and a parameter n_γ that controls the regularity of a stochastic spike generator. The mathematical formulation for the continuous-valued portion of the model, up to the spike generator, is given by the following three equations:

$$\frac{du(t)}{dt} = -\frac{u(t)}{\tau_m} + G\frac{i(t)}{\tau_m} \tag{1a}$$

$$\frac{d\theta(t)}{dt} = -\frac{\theta(t) - \theta_0}{\tau_a} + A\frac{u(t)}{\tau_a} \tag{1b}$$

$$r(t) = \begin{cases} R_{max}\tanh\left(\dfrac{u(t) - \theta(t)}{\lambda}\right) & u(t) \geq \theta(t) \\ 0.0 & u(t) < \theta(t) \end{cases} \tag{1c}$$

The parameters in equations 1a–b can be directly related to experimentally observable aspects of the time course of a neuron's response to a step stimulus. Equation 1c is a sigmoid-shaped function describing the instantaneous firing rate $r(t)$ in terms of the instantaneous values of the membrane potential $u(t)$ and the threshold $\theta(t)$. For $u < \theta$, the firing rate is zero; for $u > \theta$, the firing rate saturates at R_{max}. The parameter λ determines the steepness of the sigmoid function; as λ increases, the transition from zero firing rate to maximum firing rate occurs over a larger range of membrane potential u.

Stochastic spike generator

A stochastic spike generating element converts the instantaneous firing rate $r(t)$ into a binary spike train $s(t)$. The regularity of spike generation in this model is controlled with a single parameter n_γ. For a constant input firing rate $r(t)$, the spike generator produces a gamma distribution of interspike intervals of order n_γ. The gamma distribution of order n is defined as the waiting time to the n^{th} event in a Poisson process. For a model neuron firing at a constant rate, $n_\gamma = 1$ produces a Poisson distribution of interspike intervals; larger values of n_γ produce progressively narrower and more bell-shaped distributions as illustrated in Fig. 4. The spike generating element is implemented by explicitly modeling an underlying Poisson subprocess with an instantaneous rate $r_{\text{sub}}(t)$ given by:

$$r_{\text{sub}}(t) = n_\gamma r(t) \tag{2a}$$

where $r(t)$ is the instantaneous spike rate given in (Eq. 1c) above. The probability of a subprocess event occurring in a single simulation time step of duration Δt is:

$$P_s(t, \Delta t) = 1 - e^{-r_{\text{sub}}(t)\Delta t} \tag{2b}$$

The spike generator produces one output spike for every n_γ subprocess events.

Primary afferent (AFF) model

Primary electrosensory afferent nerve fibers are responsive to sinusoidal modulation frequencies from near DC (0.1 Hz) up to about 10–15 Hz, with a peak sensitivity around 4–8 Hz (Montgomery 1984a; New 1990) and they adapt slowly to sustained stimuli (Bodznick et al. 1993). In recordings from paralyzed animals, afferents have a regular spontaneous discharge rate of about 15–18 spikes/s (at 17°C). The firing rate is modulated by applied electric fields in a nearly linear fashion, with a slope of about 4 (spikes/s)/(μV/cm) (Montgomery 1984a). We incorporate these properties into a neural model of the form described in equations 1a–c where we model the input $i(t)$ as a voltage gradient in units μV/cm, rather than as an input current. Parameter values for the afferent model are $G = -1.0$, $\tau_m = 0.015$ s, $A = 1.0$, $\tau_a = 4.0$ s, $\theta_0 = -3.6$ μV/cm, $\lambda = 19.0$ μV/cm, $R_{\max} = 80.0$ Hz. The gain G is negative to reflect the fact that primary afferents respond with an increase in firing rate to a cathodal (negative) stimulus. The regularity of afferent nerve activity, which is presumed to arise in part from the convergence of multiple electroreceptor cells onto a single afferent nerve fiber, is modeled by setting $n_\gamma = 8$ in the stochastic spike generating element.

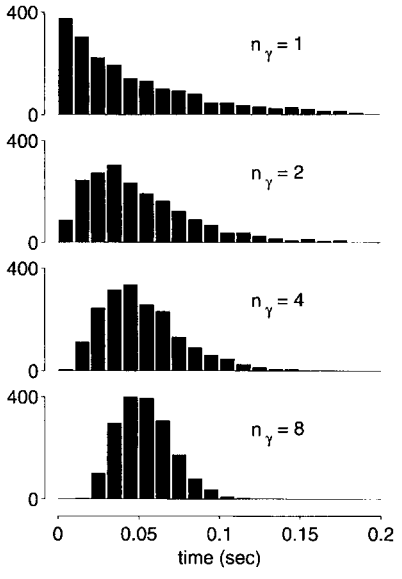

Fig. 4 Histograms of interspike intervals (ISI) produced by the stochastic spike generating element for different values of the parameter n_γ which controls spike regularity. For $n_\gamma = 1$, the spike generator produces a Poisson ISI distribution; for larger values of n_γ the distribution becomes progressively narrower and more bell-shaped. These histograms were generated using a constant input rate $r(t) = 20$ Hz, resulting in a mean interspike interval of 0.05 s in all cases

Central zone inhibitory interneuron (IN) model

Montgomery and Bodznick (1993) described a class of central zone interneurons in *Raja* on which our model interneurons are based. These neurons are monosynaptically activated by afferent electrosensory input and respond well to both common mode signals, as well as extrinsic E-fields. In general, interneuron response properties appear similar to primary afferents, but interneurons appear to be more phasic and have a somewhat lower peak firing rate. Parameter values for the interneuron model are: $G = 1.0$, $\tau_m = 0.005$ s, $A = 1.0$, $\tau_a = 4.0$ s, $\theta_0 = -0.08$, $\lambda = 0.60$, $R_{\max} = 75.0$ Hz, $n_\gamma = 8$. Since we do not have neurophysiological data to constrain the magnitude of the synaptic input currents or the magnitude of intracellular potential changes, we treat $i(t)$ and $u(t)$ as dimensionless variables in the model, hence G, A, θ_0, and λ do not have explicit units.

Ascending efferent neuron (AEN) model

In comparison to primary afferents, ascending efferent neurons are more sensitive to extrinsic stimuli, much less sensitive to common-mode stimuli, and saturate at a lower firing frequency; in addition, their frequency tuning curves are not appreciably different from primary afferents (Montgomery 1984a; New 1990). These observed characteristics arise from a combination of both the intrinsic neuron properties and network properties, so parameters for the AEN neuron model had to be appropriately tuned within the context of a working network model. Parameter values for the AEN model are: $G = 1.0$, $\tau_m = 0.005$ s, $A = 1.0$, $\tau_a = 4.0$ s, $\theta_0 = -0.01$, $\lambda = 0.23$, $R_{\max} = 50.0$ Hz, $n_\gamma = 2$.

DGR parallel fiber model

Parallel fibers arising from granule cells of the dorsal granular ridge (DGR) carry proprioceptive, efference copy, electrosensory and other sensory signals to the DON. Extracellular recordings have revealed that many DGR units are modulated by the animal's ventilatory activity with a period that matches the ventilatory cycle, but with variable offset and duration (New and Bodznick 1990; Hjelmstad et al. 1993; Conley and Bodznick 1994). While these recordings are likely to come from DGR afferents and Golgi cells, rather than from the small granule cells that give rise to the parallel fibers, we expect that modulation patterns of parallel fiber activity will be similar. We model the subset of DGR parallel fibers whose activity is correlated with ventilation using phase-locked half-cycle sinusoids of variable offset and duration with respect to the ventilatory cycle. The firing rate for an individual DGR fiber is given by:

$$r_i(t) = r_{\text{spont}} + \begin{cases} r_{\max} \sin\left(\dfrac{2\pi(t_{\text{vent}} - \delta_i)}{T_i}\right) & 0 \le (t_{\text{vent}} - \delta_i) \le \dfrac{T_i}{2} \\ 0.0 & \text{otherwise} \end{cases} \tag{3}$$

where $r_i(t)$ is the instantaneous firing rate for the i^{th} fiber, t_{vent} is the time from the beginning of the ventilatory cycle, δ_i is an offset with respect to the beginning of the ventilatory cycle, and T_i is the full-cycle period. These rate values are then passed to the stochastic spike generator to produce trains of action potentials. Parameter values used for the DGR model in this simulation: $r_{spont} = 2.0$ Hz, $r_{max} = 40.0$ Hz, $T_i = 0.2$–1.0 s, $\delta_i = 0.0$–2.0 s and $n_\gamma = 8$. Values of T_i and δ_i were randomly generated with a uniform distribution over the specified range of values. A representative sample of 64 of the 256 DGR spike trains produced by this model is shown in Fig. 5.

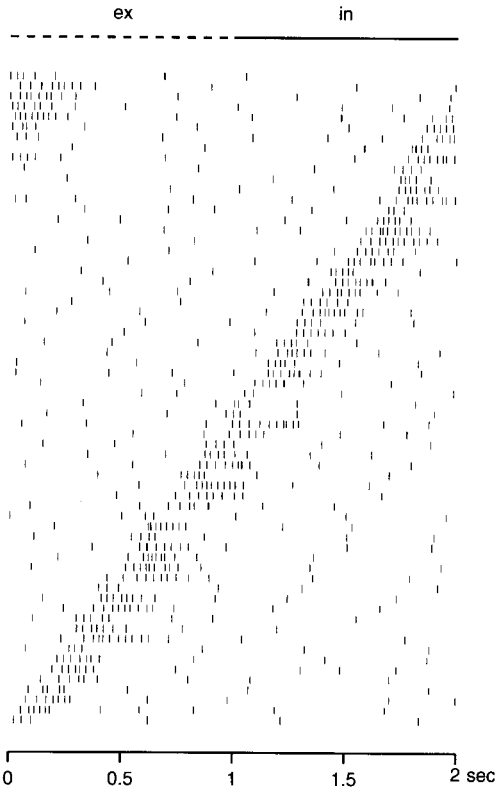

Fig. 5 Representative DGR parallel fiber spike trains for a single ventilatory cycle. DGR activity is phase-locked to the ventilatory cycle, but the modulation patterns have variable offset and duration. The phase of the ventilatory cycle is indicated by the bar at the top of the figure (*ex*: exhalation *in*: inhalation). In the figure, the spike trains have been sorted by offset for ease of visualization, although no such sorting takes place in the simulation. The mathematical description of the DGR fiber model is given in equation 3

Stellate cell (ST) model

Due to a lack of anatomical and physiological constraints, we do not model the convergence of parallel fibers onto the stellate cells and the stellate cell input-output relationship, but rather model only the stellate cell outputs. As is the case for parallel fibers, we expect that the activity patterns of the subset of stellate cells that are correlated with ventilation will show variable offset and duration with respect to the ventilatory cycle and can be modeled as phase-locked half-cycle sinusoids as given by equation 3. Parameter values used for the ST model: $r_{spont} = 2.0$ Hz, $r_{max} = 40.0$ Hz, $T_i = 1.0$ s, $\delta_i = 0.0$–2.0 s and $n_\gamma = 8$. Uniformly spaced values of δ_i were chosen to span the full range of the ventilatory cycle.

Synapse model

Synapses are modeled as being either excitatory or inhibitory and as being either fixed or plastic. The strength of each synapse is characterized by a weight value w which is positive for excitatory synapses and negative for inhibitory synapses. For plastic synapses, the weight is a time-dependent variable $w(t)$. Each synapse also has a state variable $g(t)$, analogous to a synaptic conductance, which is driven by presynaptic spike activity. A single presynaptic spike gives rise to a postsynaptic conductance change described by

$$g_\alpha(t) = \alpha^2 t e^{-\alpha t} \tag{4a}$$

This form, often referred to as the alpha function, provides a simple functional form which is capable of approximating the time course of experimentally measured synaptic currents using a single free parameter (Jack and Redman 1971). The normalization of the alpha function is such that a single postsynaptic conductance event has unit area (i.e., $\int_0^\infty g_\alpha(t)\,dt = 1$). The parameter α can be interpreted in terms of a synaptic time constant $\alpha = 1/\tau_{syn}$, where τ_{syn} is the time for a unitary conductance change to reach its peak value. In our model we use $\tau_{syn} = 0.010$ s for excitatory synapses and $\tau_{syn} = 0.050$ s for inhibitory synapses.

A series of presynaptic spikes gives rise to a summed conductance of the form:

$$g(t) = \sum_{i=1}^{n} g_\alpha(t - t_i) = \alpha^2 \sum_{i=1}^{n} (t - t_i) e^{-\alpha(t - t_i)} \tag{4b}$$

where t_i is the time of the i^{th} presynaptic spike. In the simulation, numerical values for $g(t)$ were computed analytically using the running sums method of Srinivasan and Chiel (1993). Each synapse contributes a postsynaptic current $i_{syn}(t)$ given by:

$$i_{syn}(t) = w(t)\, g(t) \tag{4c}$$

Note that we do not include a separate term for the driving force in equation 4c (i.e., $u(t) - E_{syn}$, where E_{syn} is the reversal potential) as is typical, but choose to absorb the time-averaged value of the driving force into the definition of w. This choice was made to minimize the number of model parameters for which we have no explicit biological constraints (i.e., synaptic reversal potentials and the overall normalization of the intracellular potential). We have verified that this simplification in the form of the synaptic current has a negligible effect on the system dynamics in this particular model.

Anti-Hebbian learning rule

For plastic synapses, the weight $w(t)$ in equation 4c evolves in time according to the following differential equation:

$$\frac{dw}{dt} = \begin{cases} 0 & w \geq w_{max} \\ \eta(r_{pre} - \rho_{pre})(r_{post} - \rho_{post}) & w_{min} < w < w_{max} \\ 0 & w \leq w_{min} \end{cases} \tag{5a}$$

where η is the learning rate, $r_{pre}(t)$ and $r_{post}(t)$ are the pre- and postsynaptic firing rates as defined in equation 1c, and $\rho_{pre}(t)$ and $\rho_{post}(t)$ are baseline levels of activity which in this model are taken to be low-pass filtered (e.g. time-averaged) versions of the pre- and postsynaptic firing rates:

$$\frac{d\rho(t)}{dt} = -\frac{\rho(t)}{\tau_{avg}} + \frac{r(t)}{\tau_{avg}} \tag{5b}$$

where τ_{avg} is the low-pass filter time-constant. In our simulations, we set $\tau_{avg} = 2.0$ s and use a learning rate $\eta = -0.0004$. The weight limits w_{min} and w_{max} in equation 5a are used to constrain excitatory synapses to positive values and inhibitory synapses to negative values.

When the learning rate η is negative, equation 5a implements an anti-Hebbian learning rule which can be qualitatively described as the following set of rules: (a) if the presynaptic element is active (relative to its baseline level) at the same time the postsynaptic element is active (relative to its baseline level), then make the synaptic weight more negative (i.e., decrease the strength of an excitatory synapse or increase the strength of an inhibitory synapse); (b) if the presynaptic neuron is less active and the postsynaptic neuron more active than baseline levels, or vice versa, then make the weight more positive (i.e., strengthen an excitatory synapse or weaken an inhibitory synapse); and finally, (c) if both pre- and postsynaptic elements are less active than their baseline levels, make the synaptic weight more negative. In the current model, rule (c) does not play much of a role because presynaptic modulation patterns are modeled as increases above a low baseline level. Including rule (c), however, makes the learning rule more general, such that it would properly handle the biologically plausible situation in which a presynaptic element had a high spontaneous rate and exhibited only decreases in activity with respect to its baseline level.

Ventilatory potentials and external stimuli

Primary afferents in the model receive input from two sources (Fig. 1). All afferents receive a common-mode ventilatory potential input $\phi_{vent}(t)$ which is modeled as a continuous sinusoid:

$$\phi_{vent}(t) = A_{vent}\sin\left(\frac{2\pi t}{T_{vent}}\right) \tag{6a}$$

with $A_{vent} = -6$ μV/cm and $T_{vent} = 2.0$ s. In addition, afferents can receive input from an external stimulus which can have a differential effect on afferents in the center and surround of the AEN receptive field. The external stimulus is typically modeled as a single-cycle sinusoid with amplitude A_{ext}, period T_{ext} and phase offset δ_{ext}. When the external stimulus is phase-locked to ventilation, the time t_{vent} is measured with respect to the onset of the most recent ventilatory cycle:

$$\phi_{ext}(t) = \begin{cases} A_{ext}\sin\left(\dfrac{2\pi t_{vent}}{T_{ext}} + \delta_{ext}\right) & \text{center} \\ f_{cent-surr}A_{ext}\sin\left(\dfrac{2\pi t_{vent}}{T_{ext}} + \delta_{ext}\right) & \text{surround} \end{cases} \tag{6b}$$

The parameter $f_{cent-surr}$ is a scaling factor for the amplitude of the stimulus in the surround relative to the center of the receptive field.

Numerical integration

Differential equations for neuron and synapse state variables were numerically integrated using an exponential update technique. This technique casts the differential equations into the following first-order form:

$$\frac{dy}{dt} = A - By \tag{7a}$$

Assuming that the coefficients A and B can be considered constant over the update interval Δt, equation 7a can be integrated directly to give the following update rule (MacGregor 1987):

$$y_{i+1} = y_i e^{-B\Delta t} + \frac{A}{B}(1 - e^{-B\Delta t}) \tag{7b}$$

For most of our simulations, we used an integration time step $\Delta t = 0.001$ s. We verified that this value of Δt gives accurate and stable solutions by comparing with simulation runs where Δt was an order of magnitude smaller.

Results

Sensitivity to external stimuli

Figure 6 shows the response sensitivity of primary afferents (AFF), inhibitory interneurons (IN) and ascending efferent neurons (AEN) in the model to a 1 Hz sinusoidal stimulus with intensities ranging between 0–20 μV/cm. The response amplitude is measured as the increase in firing rate above the spontaneous rate. Spontaneous rates in the model are AFF: 15 spikes/s, IN: 10 spikes/s and AEN: 2 spikes/s. The measurements shown in Fig. 6 were made under conditions in which there was no ventilatory input to the DON model, thus simulating the "paralyzed" condition in the experimental preparation in which ventilatory movements are suppressed using curare (Montgomery 1984a; New 1990). As will be discussed in more detail later, synaptic weights in the molecular layer from dorsal granular ridge (DGR) parallel fibers and stellate (ST) cells onto AENs are adaptive. Thus individual synapses will have different steady-state values under "paralyzed" and "ventilating" conditions. In the simulation, we predetermine and store appropriate weight

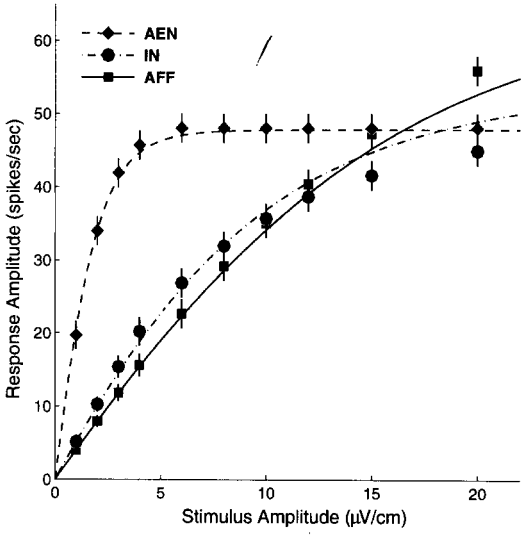

Fig. 6 Stimulus-response curves for primary afferent (*AFF*), inhibitory interneuron (*IN*), and ascending efferent neurons (*AEN*) in the model. The stimulus is a 1 Hz sinusoid; stimulus amplitude is measured as 1/2 the peak-to-peak value. Data points were obtained by fitting a sinusoid to poststimulus time histograms of spike activity accumulated over 10 stimulus cycles. *Smooth curves* indicate expected values computed from model parameters using the continuous firing rate approximation. (For a comparison with experimental data, refer to Fig. 3 of Montgomery 1984a)

settings for each condition and then select the appropriate weight set when initializing a simulation run. In the case of the simulation runs associated with Fig. 6, molecular layer weights were initialized to the paralyzed weight set.

Primary afferents in the model have an initial response slope of about 4 (spikes/s)/(μV/cm) and show a fairly linear increase in response amplitude with increasing stimulus intensity with only modest saturation effects becoming apparent over the range of intensities tested. In comparison, ascending efferent neurons (AENs) in the model have a much higher initial slope of about 20 (spikes/s)/(μV/cm), but rapidly reach saturation for stimulus intensities above a few μV/cm. These results are in good qualitative agreement with experimental results obtained by Montgomery (1984a) and New (1990). Much less is known about inhibitory interneuron (IN) response properties, other than the qualitative observations that they are more afferent-like than AEN-like. In our model, INs have an initial response slope of about 5 (spikes/s) /(μV/cm), which is slightly higher than for primary afferents and they show much less of a saturation effect than AENs. The difference in slope of the stimulus-response curves for INs and AENs relative to primary afferents is due to the combined effects of three factors: the degree of input convergence, the synaptic strengths of the inputs and the intrinsic sensitivity of the neuron, the latter being inversely related to the slope parameter λ in (1c).

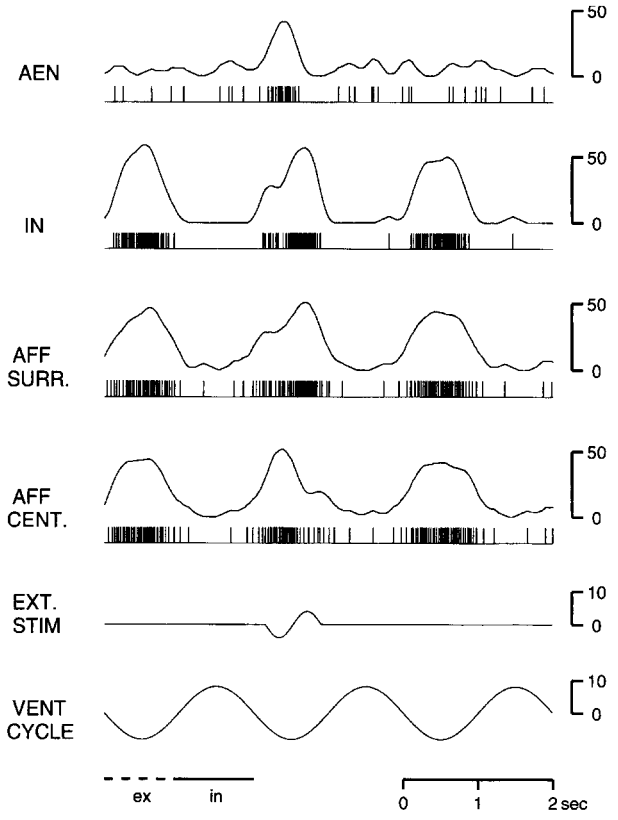

Fig. 7 Typical responses of DON model neurons to ventilatory modulation plus a superimposed external stimulus shown for three periods of the ventilatory cycle. The external stimulus is applied only during the second cycle. Ventilatory modulation is readily apparent in primary afferent (*AFF*) and inhibitory interneuron (*IN*) activity, but is absent in the ascending efferent neuron (*AEN*) activity. The AEN only responds to the external stimulus. The stimulus is a single cycle sinusoid with an amplitude of 5 μV/cm and a period of 0.75 s. *Continuous curves* above the spike train records are estimates of the instantaneous firing rate constructed using a triangular rate estimating filter (Paulin 1992). The firing rate estimates are included as an aid for visualizing the neural activity patterns, they are not part of the neural model. Units for vertical scale bars are μV/cm for the ventilatory and external stimuli and spikes/s for the neuron activity records

Ventilatory suppression

Figure 7 illustrates the ventilatory suppression performance of the DON model. Typical responses are shown for one representative primary afferent in the AEN receptive field center, one primary afferent in the surround, one inhibitory interneuron and the sole AEN in the model. All afferents receive a ventilatory modulation signal plus a superimposed external stimulus which is applied during the second of the three ventilatory cycles shown in the figure. Note that ventilatory modulation is readily apparent in the AFF and IN spike trains, but is absent in the AEN spike activity. The AEN only responds to the external stimulus. Above each spike train record in Fig. 7 is an estimate of the instantaneous firing rate constructed using a triangular rate estimating filter with a full width at half maximum (FWHM) of 200 ms. The triangle filter method gives more accurate rate estimates than standard rate histogram techniques, particularly when extracting rate estimates from a single trial (Paulin 1992). The firing rate estimates are included in Fig. 7 as an aid in visualization of the activity patterns, they are not part of the neural model itself. From these rate estimates, or by looking carefully at the corresponding spike trains, we can see that the primary afferent in the AEN receptive field center (AFF CENT.) is excited by the negative phase of the external stimulus and inhibited by the positive phase. The pattern of excitation and inhibition from the external stimulus is superimposed on the ongoing quasi-sinusoidal pattern of ventilatory modulation. The external stimulus causes the opposite effect in the surround. The primary afferent in this region (AFF SURR.) and the representative inhibitory interneuron (IN) are inhibited by the negative phase and excited by the positive phase of the external stimulus. This relationship between center and surround modulation patterns would arise, for example, when the center and surround receptive fields were oriented 180° apart and a uniform field stimulus was presented along that axis. The AEN responds to the external stimulus with an increase in firing during the negative

phase of the external stimulus and an inhibition during the positive phase. Conceptually, the AEN response arises from the difference between the direct afferent input conveyed from the center of the receptive field (AFF CENT.) and the surround input conveyed via inhibitory interneurons (IN). This subtraction of center and surround activity is the basis of the common-mode rejection mechanism that suppresses the AEN response to ventilation and can also enhance the AEN response to external stimuli.

Demonstration of adaptive suppression

To achieve the degree of AEN ventilatory suppression illustrated in Fig. 7, it was necessary to first invoke the anti-Hebbian learning rule to fine-tune the molecular layer synaptic weights in the DON model. Figure 8 demonstrates how this synaptic learning rule is capable of adaptively canceling out a component of the input signal that is consistently correlated with the animal's ventilatory cycle. The stimulus presentation used to generate the data for Fig. 8 is similar to the protocol used by Montgomery and Bodznick (1994) to experimentally demonstrate adaptive cancellation in this system. Initially, in the prestimulus period (0 min), the AEN is unmodulated by ventilation alone because the common-mode rejection and adaptive filtering mechanisms have already canceled out the ventilatory component of the input signal. Beginning at $t = 0$, an external stimulus is presented phase-locked to the ventilatory cycle. The stimulus is a single cycle sinusoid with an offset relative to the ventilatory cycle of 0.15 s and a duration of 0.75 s (identical to the stimulus used in Fig. 7). The AEN initially responds well to the stimulus with a period of excitation followed by inhibition, as seen in the first record following stimulus onset (0.5 min). This response gradually diminishes over the course of several minutes (1, 2, 4 min), illustrating the effects of the adaptive cancellation mechanism. When the stimulus is turned off after 4 mins of coupling, a "negative image" of the original response pattern is apparent in the subsequent data record (4.5 min). Where the AEN was originally excited it is now inhibited, and vice versa. This negative image, which was not present under identical external conditions (i.e., stimulus off) in the prestimulus record (0 min), has been created by changes in molecular layer synaptic weights from DGR parallel fiber and stellate cells. The anti-Hebbian learning rule has caused weight changes at molecular layer synapses whose presynaptic activity was consistently correlated with AEN activity. The induced weight changes were of an appropriate magnitude and direction to effectively cancel out the AEN modulation caused by the external stimulus. If the stimulus remains off for several minutes, the negative image decays away and the AEN adapts back to its initial flat response profile. The time course of molecular layer weight changes that occur during this adaptive tuning process are illustrated in Fig. 9 for a representative set of synapses.

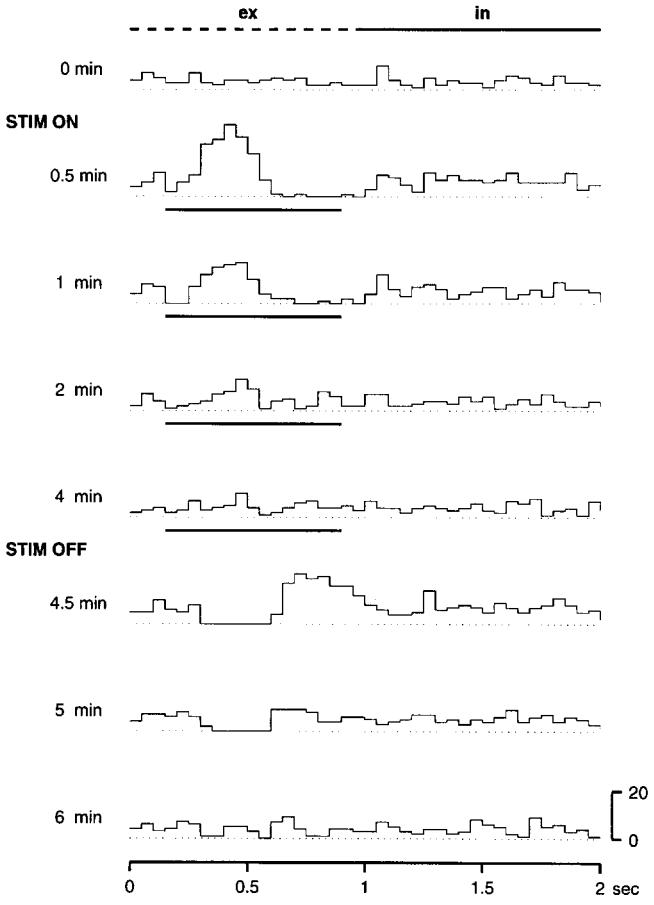

Fig. 8 Adaptive cancellation of an external electrosensory stimulus that is phase-locked to the fish's own ventilatory cycle. Responses of the model ascending efferent neuron (AEN) are shown as cycle histograms with respect to the ventilatory cycle (15 trials each). Times along the left-hand side of the figure indicate the time at the end of the block of 15 trials that contribute to the corresponding histogram. The phase of the ventilatory cycle is indicated by the *bar* at the top of the figure (*ex*: exhalation *in*: inhalation). Initially, the AEN is unmodulated by ventilation alone (0 min). Beginning at t = 0, an external stimulus is presented phase-locked to the ventilatory cycle. The stimulus is a single cycle sinusoid with an offset relative to the ventilatory cycle of 0.15 s and a duration of 0.75 s (identical to the stimulus in Fig. 7). The timing of the stimulus is indicated by the *horizontal bars* beneath the individual records. The AEN initially responds well to the stimulus with an initial period of excitation followed by a period of inhibition (0.5 min). This response gradually diminishes over the course of several minutes (1, 2, 4 min). When the stimulus is turned off after 4 min of coupling, a "negative image" of the original response pattern is apparent (4.5 min). After several minutes with the stimulus off, the negative image decays and the AEN adapts back to its initial flat response profile (5, 6 min). (cf. Fig. 1, Montgomery and Bodznick 1994)

Contribution of molecular layer inputs to ventilatory suppression

Having demonstrated the functional characteristics of the molecular layer adaptive filter, we can now return

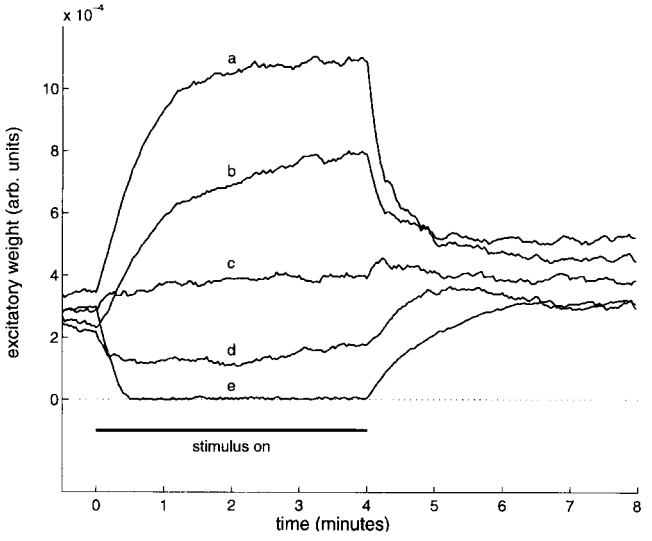

Fig. 9 Time course of DGR parallel fiber synaptic weight changes during the adaptive cancellation experiment shown in Fig. 8. This figure shows 5 representative cases out of the 256 parallel fiber synapses in the simulation. Some weights increase in magnitude, some decrease, and others remain relatively unchanged. One of the synapses shown here undergoes an approximate 3-fold increase in synaptic strength and one gets turned off completely; these two cases represent the extremes for weight changes observed in this particular simulation run. The temporal offsets of the individual parallel fiber responses with respect to the beginning of the ventilatory period are: (*a*) 0.52, (*b*) 0.45, (*c*) 1.35, (*d*) 0.40, and (*e*) 0.24 s

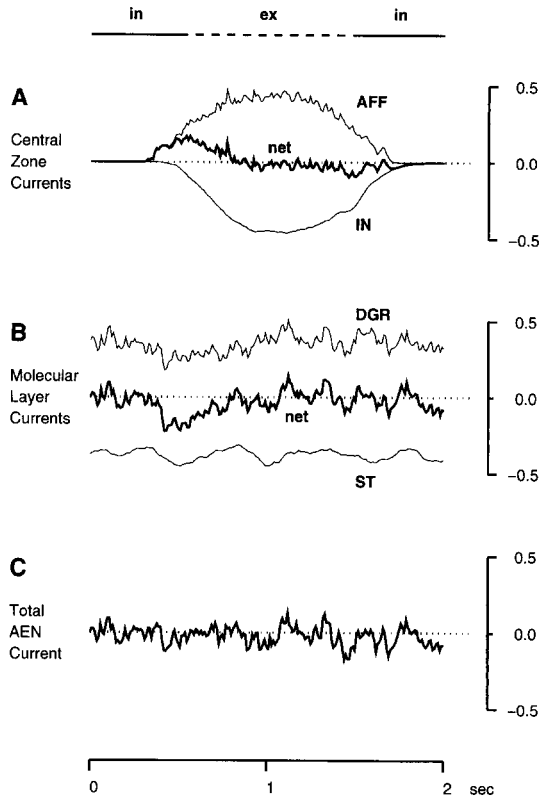

Fig. 10 Contributions of individual synaptic inputs to the total AEN current during a ventilatory cycle; currents are displayed separately for (**A**) central zone, (**B**) molecular layer, and (**C**) total current. The net current in each region is indicated by a *heavy solid line*. The total AEN current in **C** is the sum of the net currents in **A** and **B**. Excitatory inputs (*AFF* and *DGR*) give rise to positive synaptic currents; inhibitory inputs (*IN* and *ST*) give rise to negative currents; the zero current level is indicated by the *horizontal dotted lines*. The phase of the ventilatory cycle is indicated by the *bar* at the top of the figure (*ex*: exhalation, *in*: inhalation)

to a more complete analysis of the mechanisms underlying the AEN ventilatory suppression that was illustrated in Fig. 7. Figure 10 shows the contributions made by the various synaptic inputs to the AEN during a single ventilatory cycle, with no external stimulus present. In the central zone, the direct afferent input from the center of the receptive field (AFF) and the inhibitory interneuron (IN) input from the surround combine to produce a net current that is close to zero. However, as discussed in more detail below, the inhibitory interneuron pathway produces a delayed inhibition relative to the direct excitation, so there is a residual positive net current from the central zone inputs at the beginning of the ventilatory cycle (Fig. 10A). By virtue of the adaptive mechanism demonstrated above, the molecular layer weights become adjusted so as to produce a net negative image current (Fig. 10B) that counterbalances the residual positive current from the central zone. The overall result is that the total AEN current (Fig. 10C) is approximately zero over the entire ventilatory cycle, and hence the AEN shows no ventilatory modulation.

Another important feature to note in Fig. 10 is that different synaptic inputs exhibit markedly different noise levels. The source of the variance in each input signal arises from a combination of effects, including convergence ratios, synaptic time constants, presynaptic spike regularity and presynaptic firing rates. The impact of this noise on overall system performance will be discussed more fully below.

Discussion

Common-mode rejection and adaptive filtering

The suppression of ventilatory reafference in this model of the DON circuitry can be considered in two stages. First, as shown in Fig. 10A, the central zone circuitry takes care of the bulk of the suppression by wiring up the AEN to make differential measurements of the input signal. Rather than responding to afferent input directly, the AEN responds to differences between one group of afferents (the excitatory center) which converge directly onto the AEN and a second group of afferents (the inhibitory surround) which act on the AEN through an intervening inhibitory interneuron (IN). This wiring pattern achieves a form of

common-mode rejection in which signals that are common to both sets of inputs are suppressed via subtraction, in the same manner as a differential amplifier suppresses common-mode signals appearing on its positive and negative input terminals. However this common-mode rejection mechanism is not completely effective in removing all of the reafferent signal. In the model, a second stage of signal filtering is invoked to take care of cleaning up the residual reafferent signal that makes it through the common-mode rejection circuitry. This is the adaptive cancellation circuitry associated with the DGR parallel fiber system. The parallel fibers can be thought of as providing a set of internal reference signals, some of which are related to the animal's ventilatory movements. The anti-Hebbian synaptic modification rule acts to cancel out any components of the residual signal after common-mode rejection that are consistently correlated with any of the internal reference signals carried by the DGR parallel fibers. Signals of extrinsic origin, such as the bioelectric fields generated by prey, will not be correlated with these internal reference signals and will not be canceled out by this secondary stage of DON filtering.

Limitations of common-mode rejection circuitry

One of the first insights from our simulation studies was that if one imposes biologically-realistic constraints on the model parameters, it is difficult to achieve good cancellation of the reafferent signal using just the central zone common-mode rejection circuitry alone. To understand why this is so, it is useful to reconsider the differential amplifier analogy. For a high-gain differential amplifier to achieve good common-mode rejection of a background signal (for example, to effectively eliminate 50/60 Hz interference from AC power lines), the signal pathways to the positive and negative terminals must be well-matched in terms of their electrical characteristics. Any mismatch in gain or phase characteristics between the two input pathways will create differences in the signals arriving at the positive and negative terminals which will then be amplified. If one thinks of the AEN as a high-gain differential amplifier, then it is clear from Fig. 1 that the positive and negative input pathways are unlikely to be well-matched. The afferent signal from the inhibitory surround passes through an intervening inhibitory interneuron, which not only introduces extra transmission delay, but also affects the temporal filtering characteristics of the pathway since it involves additional time constants associated with activation of the inhibitory neuron and subsequent activation of inhibitory synapses onto the AEN. In order to effectively counterbalance the excitatory postsynaptic currents generated by direct afferent input, we suspect that this inhibition may involve K^+ conductance changes, rather than shunting-type Cl^- conductances. This suggests that the inhibition may be mediated by the G-protein-coupled $GABA_B$ receptor, in which case the activation time constants for the inhibitory synapses would be an order of magnitude longer than for excitatory synapses (Otis et al. 1993). Thus, in response to an increase in afferent input, the synaptic current from the excitatory inputs would typically start earlier and rise faster than the synaptic current from the inhibitory inputs. In the model, this mismatch in the central zone input pathways gives rise to a robust transient positive AEN current as shown in Fig. 10A, that is present in the simulation over a broad range of biologically-realistic parameter values.

Advantages of a second-stage adaptive filter

In principle, the mismatch between the signal transmission properties of the central zone excitatory and inhibitory pathways onto the AEN could be minimized by judicious selection of parameter values (generally, increasing excitatory and decreasing inhibitory time constants), although this would push the parameter values relatively far from their normally expected values. If the common-mode circuitry in the central zone of the DON were indeed acting alone, without a secondary stage of signal filtering, and if the quality of the common-mode rejection were of great significance to the overall system performance, then one might expect that selective pressures could have favored such parameter value shifts over the course of evolution of the DON circuitry. However, the presence of a secondary stage of adaptive filtering in the system alleviates the need for precise balancing of the excitatory and inhibitory pathways, since any residual reafferent signal making it through the common-mode stage can be "trimmed" away by the adaptive filter. While this adaptive trimming of the common-mode rejection circuitry is certainly a useful feature, the signal processing capabilities of this secondary stage of adaptive filtering are much more extensive than this simple task indicates. Rather than just eliminating incoming signals that are common across afferents, the adaptive filtering stage can eliminate any component of the incoming signal that is consistently correlated with one or more of the internal reference signals carried by the DGR parallel fibers. For example, changes in body and fin position, or opening and closing of the mouth, can give rise to localized reafferent electrosensory signals that affect certain afferents more than others. These localized reafferent signals can be removed by the adaptive filter, even though they are not common-mode, so long as there is a set of appropriately correlated reference signals provided by DGR parallel fibers. Thus the addition of a second-stage adaptive filter greatly enhances the ability of the DON circuitry to suppress a broad range of reafferent sensory signals as well as certain other non-reafferent, but predictable components of the sensory input.

Possible sites of synaptic plasticity

In the DON circuit model shown in Fig. 1, there are three classes of synapses that seem to be good candidates for involvement in an anti-Hebbian based mechanism for reafferent suppression: (1) excitatory synapses from DGR parallel fibers onto AENs, (2) inhibitory synapses from molecular layer stellate (ST) neurons onto AENs and (3) inhibitory synapses from central zone interneurons (INs) onto AENs. For the first two classes, an important issue is whether molecular layer plasticity is confined to just the excitatory or inhibitory synapses, or whether it is present in both sets of synapses. We have modeled all three possibilities (excitatory plasticity alone, inhibitory plasticity alone and both excitatory and inhibitory plasticity) and demonstrated that all three configurations can produce results consistent with existing experimental data. As a general rule, simulations with combined excitatory and inhibitory plasticity were the most robust and performed well under all conditions. For simulations where only one type of synapse was allowed to change, the initial weights had to be biased at a relatively high level such that both positive and negative weight changes of sufficient magnitude could occur without reaching zero, otherwise the adaptive mechanism would exhibit asymmetric cancellation of excitatory and inhibitory components of the input. Increasing the bias levels on the initial synaptic weights, however, generally has a detrimental effect on the signal-to-noise ratio of the system, since increasing the synaptic weights amplifies the synaptic noise (see Fig. 10) associated with those inputs.

The fact that the central zone interneuron synapses are included in this list of possible sites of anti-Hebbian plasticity is actually an outgrowth of this modeling study. Originally, we formulated the model assuming that the central zone synapses were fixed and that plasticity was limited to the molecular layer. However, in the course of our simulation studies, it became clear that it could be advantageous for the system to employ an adaptive mechanism to adjust the strength of inhibitory inputs in the central zone to balance the direct excitatory input. This new hypothesis could be tested using an experimental paradigm in which unbalanced stimuli from two separate dipole sources are repeatedly presented to the excitatory and inhibitory receptive fields of an AEN in a paralyzed animal, thus creating a persistent imbalance in the excitatory and inhibitory central zone inputs to the AEN. If the central zone synapses are plastic, then the AEN would be able to cancel this imbalance. If such an effect were observed it would be highly suggestive of central zone plasticity, but additional studies would be needed to rule out contributions from electroreceptive feedback signals carried by the DGR parallel fiber system. For example, this might be accomplished by using lidocaine injections into the DGR to knock out parallel fiber inputs to the DON or by employing brief electrical stimuli such that descending electrosensory feedback signals would arrive too late to aid in cancellation of the imbalance.

Having added central zone interneuron synapses to the list of possible sites of anti-Hebbian plasticity in the model, one might ask why not make the afferent to AEN or afferent to interneuron synapses plastic too. The reason that these synapses are not included on the list is due to the fact that when all of the synaptic inputs onto a target neuron are allowed to change via an anti-Hebbian learning rule, it is likely that the system will move toward a steady-state solution in which all of the synaptic weights are zero. In the case where all weights are plastic, correlations between pre- and post-synaptic activity can be completely eliminated by simply learning to ignore all of the inputs. Thus we are led to suspect that there should always be at least one set of synaptic inputs to a target neuron which do not obey an anti-Hebbian update rule. In the case of the AEN, we assume that it is the afferent synapses that are fixed.

Anti-Hebbian learning rules

In general, anti-Hebbian learning rules implement a form of synaptic plasticity which can be summarized as follows: if pre-synaptic activity is correlated with post-synaptic activity, then make the corresponding synaptic weight more negative (i.e., reduce the strength of an excitatory synapse or increase the strength of an inhibitory synapse) and, conversely, if pre- and post-synaptic activities are anti-correlated, make the weight more positive. When expressing this basic idea as an explicit set of equations or update rules, many variations can be formulated depending on how one chooses to define activity (firing rate, local membrane potential, calcium concentration, nitric oxide concentration, etc.), whether or not one computes activity levels relative to zero or relative to a non-zero baseline, how one chooses to handle situations where one or both of the neurons is inactive or below baseline activation, whether or not one includes activity-independent terms such as synaptic weight decay, and whether or not synaptic weights are bounded. Some of these issues have been reviewed in the context of Hebbian learning rules by Dayan and Sejnowski (1993) and Montague and Sejnowski (1994). The goal of our study is not to try to distinguish between alternative forms of anti-Hebbian learning, but rather to demonstrate that this general class of learning rules is sufficient to account for the adaptive reafference suppression phenomenon observed in the elasmobranch DON. For our simulation studies we have selected a particular form of the learning rule that we have found to be robust in terms of its stability and convergence properties. In our particular implementation, activity is measured in terms of the firing frequency of the pre- and post-synaptic elements relative to an adaptive baseline level that tracks the

average activity levels with a time-constant on the order of the ventilatory period. Such a learning rule is quite plausible from a biological perspective, but we do not intend to imply that this particular form is the only, or even the most likely, possibility. In fact, we have explored several alternative forms of anti-Hebbian learning rules in our simulation studies and found that most of them are suitable for reproducing the desired phenomenon. Eludication of the details of the synaptic modification rule in the DON will require further experimental studies targeted at that specific issue.

Importance of convergence ratios, spike regularity, and synaptic filtering

At the outset of this modeling study, we were reasonably confident that the common-mode rejection and adaptive filter mechanisms would work well if the signals in the system were smoothly varying, continuous functions of time. We were not certain, however, how these mechanisms would perform when the signals were conveyed by trains of discrete, stochastically generated action potentials, since we expected this would introduce a significant source of variance into the signal components. Again thinking in terms of the high-gain differential amplifier analogy for AEN function, any uncorrelated noise appearing at the inputs to the positive and negative terminals will be amplified by the system. As expected, reasonably large signal fluctuations do occur in our model, as can be seen in several of the AEN input current records in Fig. 10. The degree of variance for a particular AEN input signal depends on a number of factors including convergence ratios, synaptic time constants, presynaptic spike regularity and presynaptic firing rates. In general, decreasing any of these factors increases the corresponding signal variance and potentially degrades the overall system performance. In our model, we have been able to achieve good performance using biologically reasonable values for these factors. However, our model operates close to the edge of acceptable noise levels. If the factors listed above are changed too much in an unfavorable direction (e.g., decreasing convergence ratios, decreasing synaptic time constants, decreasing afferent spike regularity), then the total input current to the AEN can become very noisy, thus decreasing the ability of the system to detect weak extrinsic signals.

Parameter insensitivity

Other than the factors listed above, the DON model performance is extremely robust to parameter variations. Because the system is adaptive, it automatically adjusts to parameter changes that affect the response properties of individual neurons, such as gain, membrane time constants and adaptation time constants. While we have tried to accurately tune our neuron models to facilitate comparison with experimental data, the adaptive reafference suppression performance of the model is largely insensitive to our parameter choices.

Unified view of common mode rejection and adaptive filtering

Throughout this paper, we have discussed common mode rejection and adaptive filtering as two independent, but cooperative, mechanisms that contribute to reafference suppression in the DON. Certainly these two mechanisms appear to have seperate neural substrates, with common mode rejection being mediated by inhibitory interneurons in the central zone and adaptive filtering involving parallel fiber and stellate cell inputs in the molecular layer. However, if the central zone inhibitory synapses also exhibit anti-Hebbian plasticity as we have suggested above, then common mode rejection and adaptive filtering can be unified into a single processing principle acting at the level of the AEN. In this unified view, each AEN receives a principal input signal from a set of fixed weight synapses in the center of its receptive field, then based on an anti-Hebbian learning rule, the weights of all other synaptic inputs to the AEN are adjusted to remove any correlations, whether the correlated signals arise from peripheral inputs in the receptive field surround or from centrally generated internal reference signals carried in the molecular layer. In this manner, each AEN filters out predictable components of the incoming electrosensory signal before conveying information on to subsequent stages of electrosensory processing in the brain.

Acknowledgements We wish to thank Drs. David Bodznick and John Montgomery for their many helpful insights and suggestions. This work supported by the US-NZ Cooperative Science Program (NZ-MORST RST-08-2-1-c: SP, NSF INT-9120572) and by NIMH 1-R29-MH49242 (MN).

References

Bastian J (1994) Pyramidal-cell plasticity in weakly electric fish: A mechanism for attenuating responses to reafferent electrosensory inputs. J Comp Physiol A 176: 63–78

Bell CC (1981) An efference copy modified by reafferent input. Science 214: 450–453

Bell CC (1982) Properties of a modifiable efference copy in an electric fish. J Neurophysiol 47: 1043–1056

Bell CC, Caputi A, Grant K, Serrier J (1993) Storage of a sensory pattern by anti-Hebbian synaptic plasticity in an electric fish. Proc Natl Acad Sci USA 90: 4650–4654

Bodznick D (1993) The specificity of an ada ptive filter that suppresses unwanted reafference in electrosensory neurons of the skate medulla. Biol Bull 185: 312–314

Bodznick D, Boord RL (1986) Electroreception in Chondrichthyes. In: Bullock TH, Heiligenberg W (eds) Electroreception. Wiley, New York, pp 225–256

Bodznick D, Montgomery JC (1992) Suppression of ventilatory reafference in the elasmobranch electrosensory system - medullary neuron receptive fields support a common mode rejection mechanism. J Exp Biol 171: 127–137

Bodznick D, Montgomery JC, Bradley DJ (1992) Suppression of common mode signals within the electrosensory system of the little skate, *Raja erinacea*. J Exp Biol 171: 107–125

Bodznick D, Hjelmstad G, Bennett MVL (1993) Accommodation to maintained stimuli in the ampullae of Lorenzini: How an electric fish achieves sensitivity in a noisy world. Jp J Physiol 43: 231–237

Conley RA, Bodznick D (1994) The cerebellar dorsal granular ridge in an elasmobranch has proprioceptive and electroreceptive representations and projects homotopically to the medullary electrosensory nucleus. J Comp Physiol A 174: 707–721

Dayan P, Sejnowski TJ (1993) The variance of covariance rules for associative memories and reinforcement learning. Neural Computation 5: 205–209

Hjelmstad GO, Parks G, Bodznick D (1993) A corollary discharge of the ventilatory motor command in the dorsal granular ridge of skates; implications for electrosensory processing. Soc Neurosci Abstr 19: 374

Holst E von, Mittelstaedt H (1950) Das Reafferenzprinzip. Naturwissenschaften 37: 464–476

Jack JJB, Redman SJ (1971) The propagation of transient potentials in some linear cable structures. J Physiol (Lond) 215: 283–320

Kalmijn AJ (1974) The detection of electric fields from inanimate and animate sources other than electric organs. In: Fessard A (ed) Electroreceptors and other specialized receptors in lower vertebrates (Handbook of sensory physiology, vol III/3) Springer, Berlin Heidelberg New York, pp 147–200

Kalmijn AJ (1982) Electric and magnetic field orientation in elasmobranch fishes. Science 218: 916–918

MacGregor RJ (1987) Neural and Brain Modeling, Academic Press, New York

Montague PR, Senjnowski TJ (1994) The predictive brain: temporal coincidence and temporal order in synaptic learning mechanisms. Learning & Memory 1: 1–33

Montgomery JC (1984a) Frequency response characteristics of primary and secondary neurons in the electrosensory system of the thornback ray. Comp Biochem Physiol 79A: 189–195

Montgomery JC (1984b) Noise cancellation in the electrosensory system of the thornback ray; common mode rejection of input produced by the animal's own ventilatory movements. J Comp Physiol A 155: 103–111

Montogomery JC, Bodznick D (1993) Hindbrain circuitry mediating common mode suppression of ventilatory reafference in the electrosensory system of the little skate, *Raja erinacea*. J Exp Biol 183: 203–215

Montgomery JC, Bodznick D (1994) An adaptive filter that cancels self-induced noise in the electrosensory and lateral line mechanosensory systems of fish. Neurosci Lett 174: 145–148

Montgomery JC, Coombs S, Conley RA, Bodznick D (1995) Hindbrain sensory processing in lateral line, electrosensory and auditory systems: A comparative overview of anatomical and functional similarities. Auditory Neurosci (in press)

Murray RW (1960) Electrical sensitivity of the ampullae of Lorenzini. Nature (Lond) 187: 957

Murray RW (1974) The ampullae of Lorenzini. In: Fessard A (ed) Electroreceptors and other specialized receptors in lower vertebrates (Handbook of sensory physiology, vol III/3) Springer, Berlin Heidelberg New York, pp 125–146

New JG (1990) Medullary electrosensory processing in the little skate. I. Response characteristics of neurons in the dorsal octavolateralis nucleus. J Comp Physiol A 167: 285–294

New JG, Bodznick D (1990) Medullary electrosensory processing in the little skate. II. Suppression of self-generated electrosensory interference during respiration. J Comp Physiol A 167: 295–307

Otis TS, DeKoninck Y, Mody I (1993) Characterization of synaptically elicited $GABA_B$ responses using patch-clamp recordings in rat hippocampal slices. J Physiol (Lond) 463: 391–407

Paulin MG (1992) Digital filters for firing rate estimation. Biol Cybern 66: 525–531

Paulin MG (1995) Electroreception and the compass sense of sharks. J Theor Biol (in press)

Srinivasan R, Chiel HJ (1993) Fast calculation of synaptic conductances. Neural Computation 5: 200–204