Network: Comput. Neural Syst. **11** (2000) 83–102. Printed in the UK PII: S0954-898X(00)11131-5

Odour recognition and segmentation by a model olfactory bulb and cortex

Zhaoping Li† and John Hertz‡

† Gatsby Computational Neuroscience Unit, 17 Queen Square, UCL, London WC1N 3AR, UK
‡ NORDITA, Blegdamsvej 17, 2100 Copenhagen Ø, Denmark

E-mail: zhaoping@gatsby.ucl.ac.uk and hertz@nordita.dk

Received 1 March 1999, in final form 7 October 1999

Abstract. We present a model of an olfactory system that performs odour segmentation. Based on the anatomy and physiology of natural olfactory systems, it consists of a pair of coupled modules, bulb and cortex. The bulb encodes the odour inputs as oscillating patterns. The cortex functions as an associative memory: when the input from the bulb matches a pattern stored in the connections between its units, the cortical units resonate in an oscillatory pattern characteristic of that odour. Further circuitry transforms this oscillatory signal to a slowly varying feedback to the bulb. This feedback implements olfactory segmentation by suppressing the bulbar response to the pre-existing odour, thereby allowing subsequent odours to be singled out for recognition.

1. Introduction

An olfactory system must solve the problems of odour detection, recognition, and segmentation. Segmentation is necessary because the odour environment often contains two or more odour objects. The system must be able to identify these objects separately and signal their presence to higher brain areas. An odour object is defined as an odour entity (which, e.g., the smell of a cat, often contains fixed proportions of multiple types of odour molecules) that enters the environment independently of other odours. Therefore, two odour objects usually do not enter the environment at the same time although they often stay together in the environment afterwards. In cases when different odours do enter the environment together as a mixture, human subjects have great difficulty identifying the components [1]. In this paper we present a model which performs odour segmentation temporally. First one odour object is detected and recognized, then the system adapts to this specific odour so a subsequent one can be detected and recognized.

The odour specificity of this adaptation is the key feature of the operation of the system. This specificity cannot be achieved with simple single-unit fatigue mechanisms [2,3] because of the highly distributed nature of odour pattern representations in the olfactory system: fatiguing neurons that respond to one odour would strongly reduce their response to another one, thereby distorting the pattern evoked by the second odour. In our model a delayed inhibitory feedback signal is directed to the input units in such a way as to cancel out the current input, leaving the system free to respond to new odours as if the first one were not there.

Our model is not intended as a faithful representation of any particular animal olfactory system. Present anatomical and physiological knowledge does not permit such detailed

0954-898X/00/010083+20$30.00

modelling. Rather, our focus is on the computations performed by different groups of neurons, based on general biological findings, which we review briefly here.

In animals, different odour molecules produce different, distributed activity patterns across the neurons of the olfactory nerve, which provide the input to the olfactory bulb [4, 5]. We do not model this part of the processing. We will simply represent different odours as different but overlapping input patterns to the bulb. They are temporally modulated by the animal's sniff cycle (typically 2–4 sniffs per second), i.e., active only during and immediately after inhalation.

The main cell types of the mammalian bulb are the excitatory mitral cells and the inhibitory granule cells. The mitral cells receive the odour input and excite the granule cells, which in turn inhibit them. The outputs of the bulb are carried to the olfactory cortex by the mitral cell axons. In vertebrate animals, odours evoke oscillatory bulbar activity in the 35–90 Hz range, which may be detected by surface EEG electrodes [6, 7]. Different parts of the bulb have the same dominant frequency but different amplitudes and phases [7, 8], and this oscillation pattern is odour-specific [8, 9]. These oscillations are an intrinsic property of the bulb, persisting after central connections to the bulb are cut [10, 11]. (In invertebrates, oscillations exist without odour input but are modulated by odours [12].) Upon repeated presentation of a conditioned odour stimulus, the bulbar oscillations weaken markedly [13]. Since olfactory receptor neurons exhibit only limited adaptation [14, 15], this adaptation must originate either in the bulb or in cortical structures.

The pyriform or primary olfactory cortex receives bulbar outputs via the lateral olfactory tract, which distributes outputs from each mitral cell over many cortical locations [4]. The signals are conveyed to the (excitatory) pyramidal cells of the cortex, both directly and via feedforward inhibitory cells in the cortex. The pyramidal cells send axon collaterals to each other and to feedback interneurons which, in turn, inhibit them. There is thus excitatory–inhibitory circuitry as in the bulb, and oscillatory responses to odours are observed in the cortex, too. However, the cortex differs from the bulb in the much greater spatial range of the excitatory connections and in the presence (or at least the greater extent) of excitatory-to-excitatory connections. This anatomical structure has led a number of workers to model the olfactory cortex as an associative memory for odours [16–21]. Furthermore, the oscillations in the cortex require input from the bulb; they do not occur spontaneously. Cortical output, including the feedback to the bulb, is from pyramidal cells [4]. Some of the feedback is direct, while some of it is via other cortical centres, notably the entorhinal cortex. Most central feedback to the bulb is to the granule cells [5]. Cooling the cortex, presumably reducing or removing the central feedback, enhances the bulbar responses [22].

The basic features outlined here constrain our model: we employ coupled excitatory and inhibitory populations in both bulb and cortex, we wire the network so that odours evoke oscillations in the bulb, which drive similar cortical oscillations through excitatory and inhibitory connections, and we send the central feedback to reduce the bulbar responses.

We will neglect many known features of animal olfactory systems, such as (to name a few) the patterns of connectivity from receptors to mitral cells, the dendrodentritic character of the mitral–granule synapses, and the differing spatial range of connectivity in bulb and cortex. Indeed, the model has no geometry: 'location' and 'distance' have no meaning here. We retain only the basic elements necessary to illustrate the basic operation of the system, in order not to obscure the functions we focus on (detection, recognition, and segmentation).

We will also hypothesize features of the system, in particular the nature of the feedback signal from the cortex to the bulb, for which there is not yet experimental evidence (though they are not incompatible with present knowledge). These assumptions will be necessary in order to make an explicit model that can be tested computationally. Some details of its

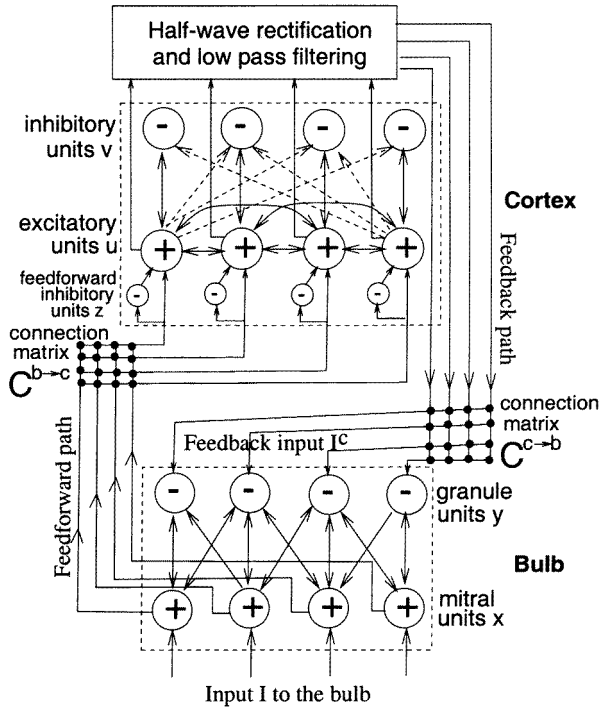

Figure 1. The model. Odour inputs $\boldsymbol{I}$ are fed into the mitral units (x) in the bulb. These interact with the inhibitory granule units (y), both locally (vertical connection lines) and nonlocally, via the connection matrices H and W (diagonal connection lines). The mitral units project their outputs to the cortex via the feedfoward matrix $C^{b\to c}$. The excitatory units in the cortex (u) receive these inputs both directly and indirectly via the feedforward inhibitory units (z). In addition to the local excitatory–inhibitory connections (vertical lines) between the excitatory (u) and the feedback inhibitory units (v), there are nonlocal connections among the excitatory units (J, solid lines) and from excitatory to inhibitory units ($\tilde{W}$, dotted lines). The outputs of the excitatory units are fed back through a matrix $C^{c\to b}$ to the granule units in the bulb after rectification and low-pass filtering. (Details of the rectification/filtering operation are shown in figure A1.)

implementation are neither crucial to the computational function of the model nor intended as explicit neurophysiological predictions. However, the basic framework of the model and the dynamical properties we find for it are subject to experimental test.

In the next section we present the model: its equations of motion and how it detects, recognizes, and segments odour inputs. The following section demonstrates how it works in simulations. In the final section we discuss the implications of our work, including potential experimental tests for this and related models and how they can help us understand the functioning of the olfactory system.

2. The model

Our model consists of two modules, a bulb and a cortex, with feedforward and feedback connections between them. It is depicted schematically in figure 1. The bulb encodes odour inputs as patterns of oscillation. These form the input to the cortex, which acts as an associative memory for odour objects, recognizing them by resonant oscillation in an odour-specific pattern when the input from the bulb matches one of its stored odour memories. The odour-

specific resonant cortical activity pattern is transformed to a feedback signal to the bulb, which approximately cancels the effect of the odour input that elicited it. The system is then able to respond to a newly arrived odour superposed on the previous one. In this way it segments temporally the different odour objects in the environment.

The model is a rate-model network [23], in which we associate each unit with a local population of cells that share common synaptic input (mitral cells for the excitatory units, granule cells for the inhibitory ones). The output (activation) of a unit, representing the average firing rate within the corresponding population, is modelled as a sigmoidal function of the net synaptic input.

In both the bulb and cortex modules, the units occur in pairs, one unit excitatory and the other inhibitory. In the absence of coupling between different such pairs, they form independent damped local oscillators. The coupling between pairs leads to oscillation patterns across the modules, with specific amplitudes for the individual local oscillators and specific phase relations between them. The odour input makes these oscillatory patterns different from odour to odour; thus, these patterns form the internal encoding of the odours. The sizes of the local populations corresponding to our formal units are different for excitatory and inhibitory units; this difference is accounted for in the model by appropriate scaling of the synaptic strengths.

We turn now to the explicit mathematical description of the two modules and the coupling between them.

2.1. The bulb

The bulb model we employ was introduced by Li and Hopfield (1989) [24, 25]. For completeness, we review it here.

The odour input to (mitral) unit i is denoted I_i. (We will also use a vector notation, in which the entire input pattern is denoted $\boldsymbol{I}$.) Adding to this the synaptic input from granule cells within the bulb, we obtain an equation of motion

$$\dot{x}_i = -\alpha x_i - \sum_j H^0_{ij} g_y(y_j) + I_i \tag{1}$$

for the (local population average) membrane potential x_i. Here α^{-1} is the membrane time constant, $g_y(\cdot)$ is the (sigmoidal) activation function of the granule units, y_j is the membrane potential for granule unit j, and H^0_{ij} is the inhibitory synaptic strength from granule unit j to mitral unit i. All the H^0_{ij} are non-negative; the inhibitory nature of the granule cells is represented by the negative sign in the second term on the right-hand side. The signal the bulb sends on to the cortex is carried by the mitral unit outputs $g_x(x_i)$ (with $g_x(.)$ their activation function). We have not included mitral–mitral connections here, because the experimental evidence for them is weak, but including them would not change the properties of the model qualitatively.

For the inhibitory units, representing local populations of granule cells, we have, similarly to (1),

$$\dot{y}_i = -\alpha y_j + \sum_j W^0_{ij} g_x(x_j) + I^{\mathrm{c}}_i, \tag{2}$$

with the mitral-to-granule synaptic matrix W^0_{ij}. Here the external input I^{c}_i represents the centrifugal input (from the cortex), which contains the feedback signal that implements the odour-specific adaptation. In describing the response to an initial odour, it can be neglected or taken as a constant background input.

To see how this network produces oscillatory excitation patterns in response to an odour, start by taking the input $\boldsymbol{I}$ to be static. It determines a fixed point $\bar{x}_i$ and $\bar{y}_i$ of the equations,

i.e., $\dot{x}_i = \dot{y}_i = 0$ at $\bar{x}_i$ and $\bar{y}_i$, which increase with odour input $\boldsymbol{I}$. Taking the deviation from this fixed point as $x_i - \bar{x}_i \to x_i$ and $y_i - \bar{y}_i \to y_i$, linearizing and eliminating the y_i leads to

$$\ddot{x}_i + 2\alpha\dot{x}_i + \alpha^2 x_i + \sum_j A_{ij}x_j = 0, \tag{3}$$

where the matrix $\mathsf{A} = \mathsf{HW}$, with $H_{ij} = H^0_{ij}g'_y(\bar{y}_j)$ and $W_{ij} = W^0_{ij}g'_x(\bar{x}_j)$. This equation describes a coupled oscillator system, with a coupling matrix A. Denoting the eigenvectors and eigenvalues of this matrix by $\boldsymbol{X}_k$ and λ_k, respectively, (3) has solutions $\boldsymbol{x} = \sum_k c_k \boldsymbol{X}_k \exp[-\alpha t \pm \mathrm{i}(\sqrt{\lambda_k}t + \phi_k)]$, with c_k and ϕ_k the amplitude and phase of the kth mode. If A is not symmetric (the general case), λ_k is complex, and the mode has oscillation frequencies $\omega_k \equiv \mathrm{Re}\,(\sqrt{\lambda_k})$. The amplitude for mode k will grow exponentially (in this linearized theory) if $\pm\mathrm{Im}\,(\sqrt{\lambda_k}) > \alpha$. Its growth will be limited by nonlinearities, and it will reach a steady-state saturation value. In this spontaneously oscillating state, the fastest-growing mode, call it the 1st mode, will dominate the output. The whole bulb will oscillate with a single frequency ω_1 (plus its higher harmonics), and the oscillation amplitudes and phases may be approximated by the complex vector $\boldsymbol{X}_1$. Thus, the olfactory bulb encodes the olfactory input via the following steps: (1) the odour input $\boldsymbol{I}$ determines the fixed point $(\bar{\boldsymbol{x}}, \bar{\boldsymbol{y}})$, which in turn (2) determines the matrix A, which then (3) determines whether the bulb will give spontanous oscillatory outputs and, if it does, the oscillation amplitude and phase pattern $\boldsymbol{X}_1$ and its frequency ω_1.

Strictly speaking, this description only applies to very small oscillations. For larger amplitudes, nonlinearities make the problem in general intractable. However, we will suppose that the present analysis gives a decent qualitative guide to the dynamics, checking this assumption later with simulations of the network.

In this model, oscillations arise strictly as a consequence of the asymmetry of the matrix A. The model could be generalized to add intrinsic single-unit oscillatory properties, and these might enhance the network oscillations. However, a model with symmetric A and intrinsic oscillatory properties only at the single-unit level cannot support oscillation patterns in which the phase varies across the units in the network. We will return to this point in section 4.

A word about timescales: the odour input varies on the timescale of a sniff: 300–500 ms. The oscillations are in the 40 Hz range, so the input $\boldsymbol{I}$ hardly changes at all over a few oscillation periods (∼25 ms). We may therefore treat periods of several oscillations as if the input were static within them, and perform the above analysis separately for each such period (adiabatic approximation).

With inhalation, the increasing input $\boldsymbol{I}$ pushes the fixed-point membrane potentials $\bar{x}_i$ from their initial values (where the activation function $g(x)$ has low gain) through a range of increasing gains, thereby increasing the size of some of the elements of the matrix A (recall the definition of A above). This increases the magnitude of both the real and imaginary parts of the eigenvalues λ_k, until the threshold where $|\mathrm{Im}\,(\sqrt{\lambda_k})| = -\alpha$, where oscillations appear. These oscillations increase in amplitude as the input increases further, until the animal stops inhaling and the input $\boldsymbol{I}$ decreases. Then the oscillations shrink and disappear as the system returns toward its resting state. This rise and fall of oscillations within each sniff cycle give the bulb outputs both a slowly varying component (2–4 Hz) and a high-frequency (25–60 Hz) one, as observed experimentally [7].

It is not known how the synaptic connections represented in the model by the matrices H^0 and W^0 develop in the real olfactory bulb, and we do not attempt to model this process here. It is possible that the real bulb acts, to some degree, as an associative memory as a result of this learning. However, our conclusions will not depend on this. Similarly, our analysis does not depend on details of the synaptic matrices, such as their range and degree of connectivity. We

require only that the connections lead to distinct oscillation patterns for different odours, with dissimilar patterns evoked by dissimilar odours.

2.2. The cortex

Our cortical module is structurally similar to that of the bulb. However, there are the following significant differences: (1) the cortex receives an oscillatory input from the bulb, while the bulb receives non-oscillatory (at the timescale of the cortical oscillation) input; (2) the cortex has excitatory-to-excitatory connections, while our bulb module does not.

We focus on the local excitatory (pyramidal) and feedback inhibitory interneuron populations. The units that represent them obey differential equations similar to those for the mitral and granule units of the bulb:

$$\dot{u}_i = -\alpha u_i - \beta^0 g_v(v_i) + \sum_j J^0_{ij} g_u(u_j) - \sum_j \tilde{H}^0_{ij} g_v(v_j) + I^{\rm b}_i, \tag{4}$$

$$\dot{v}_i = -\alpha v_i + \gamma^0 g_u(u_i) + \sum_j \tilde{W}^0_{ij} g_u(u_j). \tag{5}$$

Here u_i represent the the average membrane potentials of the local excitatory populations and v_i those of the inhibitory populations. The synaptic matrix J^0 is excitatory-to-excitatory connections, $\tilde{\mathsf{H}}^0$ is inhibitory-to-excitatory connections, and $\tilde{\mathsf{W}}^0$ is excitatory-to-inhibitory connections. For later convenience, we have written the local terms (the effect of v_i on u_i and vice versa) explicitly, so $\tilde{\mathsf{H}}^0$ and $\tilde{\mathsf{W}}^0$ have no diagonal elements. We also assume $J^0_{ii} = 0$. $I^{\rm b}_i$ are the net inputs from the bulb, both directly and indirectly via the feedforward inhibitory units (see later for the description of this pathway). Like the bulb activity itself, these contain in general both a slow part $I^{\rm b0}_i$, varying with the sniff cycle, and an oscillating (γ-band) part $\delta I^{\rm b}_i$, i.e., $I^b_i \equiv I^{\rm b0}_i + \delta I^{\rm b}_i$.

We can carry out the same analysis as in the bulb, taking the fixed point as $(\bar{\boldsymbol{u}}, \bar{\boldsymbol{v}})$, which are determined by $\boldsymbol{I}^{\rm b0}$, i.e., $\dot{\boldsymbol{u}} = \dot{\boldsymbol{v}} = 0$ at $(\bar{\boldsymbol{u}}, \bar{\boldsymbol{v}})$ when $I^b_i = I^{b0}_i$ with $\delta I^b_i = 0$. Taking $\boldsymbol{u} \to \boldsymbol{u} - \bar{\boldsymbol{u}}$, $\boldsymbol{v} \to \boldsymbol{v} - \bar{\boldsymbol{v}}$, linearizing and eliminating the v_i, we obtain

$$\ddot{u}_i + \sum_j [2\alpha\delta_{ij} - J_{ij}]\dot{u}_j + \sum_j [(\alpha^2 + \beta_i\gamma_i)\delta_{ij} - \alpha J_{ij} + \gamma_i \tilde{H}_{ij} + \beta_i \tilde{W}_{ij} + \sum_k \tilde{H}_{ik}\tilde{W}_{kj}]u_j$$
$$= (\partial_t + \alpha)\delta I^{\rm b}_i. \tag{6}$$

Here $\beta_i = \beta^0 g'_v(\bar{v}_i)$, $\gamma_i = \gamma^0 g'_u(\bar{u}_i)$, $J_{ij} = J^0_{ij} g'_u(\bar{u}_j)$, $\tilde{H}_{ij} = \tilde{H}^0_{ij} g'_v(\bar{v}_j)$, and $\tilde{W}_{ij} = \tilde{W}^0_{ij} g'_u(\bar{u}_j)$. Thus this is a system of driven oscillators coupled by connections J, $\tilde{\mathsf{H}}$, and $\tilde{\mathsf{W}}$ and driven by an external oscillatory signal $\delta\dot{\boldsymbol{I}}^{\rm b} + \alpha\delta\boldsymbol{I}^{\rm b}$, which is proportional to $\delta\boldsymbol{I}^{\rm b}$ for a purely sinusoidal oscillation. A single dissipative oscillator driven by an oscillatory force will resonate to it if the frequency of the driving force matches the intrinsic frequency of the oscillator. A system of coupled oscillators has its intrinsic oscillation patterns—the normal modes determined by the coupling. Analogously, it will also resonate to the input when the driving force, a complex vector proportional to $\delta\boldsymbol{I}^{\rm b}$, matches one of the intrinsic modes, also a complex vector, in frequency and in its pattern of oscillation amplitudes and phases.

It is apparent from equation (6) that the matrices $\tilde{\mathsf{H}}$ and $\tilde{\mathsf{W}}$ play the same roles. Therefore, for simplicity, we will drop the inhibitory-to-excitatory couplings $\tilde{\mathsf{H}}$ from now on, thinking of the fact that the real anatomical long-range connections appear to come predominantly from excitatory cells.

Odour selectivity and sensitivity. In our model, the olfactory cortex functions as an associative memory, as described and modelled by a number of authors [16–21]. It is similar to a Hopfield

model, but instead of stationary patterns it stores oscillating patterns which vary in phase as well as magnitude across the units of the network. The memory pattern for the μth odour is described by a complex vector ξ^μ, whose component ξ_i^μ describes both the relative amplitude and phase of the oscillation in the ith unit. The cortex stores the memories about the odours in the synaptic weights J^0 and $\tilde{\mathsf{W}}^0$, or, effectively, the coupling between oscillators. It then recognizes the input odours, as coded by the oscillating input patterns $\delta \boldsymbol{I}^{\mathrm{b}}$ (which are linearly related to the bulbar oscillatory outputs), by resonating to them, giving high-amplitude oscillatory responses itself. If, however, the input $\delta \boldsymbol{I}^{\mathrm{b}}$ does not match one of the stored odour patterns ξ^μ closely enough, the cortex will fail to respond appreciably.

In the present model the memory pattern ξ_i^μ for odours $\mu = 1, 2, \ldots$ are designed into the synaptic connections J and $\tilde{\mathsf{W}}$. Let ω be the oscillation frequency, $\delta I_i^{\mathrm{b}} \propto \mathrm{e}^{-\mathrm{i}\omega t}$. Once the oscillation reaches a steady amplitude $u_i \propto \mathrm{e}^{-\mathrm{i}\omega t}$, we have $\dot{u}_i = -\mathrm{i}\omega u_i$, $\ddot{u}_i = -\mathrm{i}\omega \dot{u}_i$, so we get

$$\dot{u}_i = \left[-2\alpha - \frac{\mathrm{i}}{\omega}(\beta_i \gamma_i + \alpha^2)\right] u_i + \sum_j \left[J_{ij} - \frac{\mathrm{i}}{\omega}(\beta_i \tilde{W}_{ij} - \alpha J_{ij})\right] u_j + \frac{\mathrm{i}}{\omega}(-\mathrm{i}\omega + \alpha)\delta I_i^{\mathrm{b}}. \tag{7}$$

The second term $[\ldots]$ on the right-hand side gives an effective coupling between the oscillators. From now on in this analysis we will make the approximation that the different local oscillators have the same natural frequencies, i.e. $\beta_i \gamma_i$ is independent of i. Assuming further that the oscillation frequencies for different odours are nearly the same, the odour patterns can then be stored in the matrices in a generalized Hebb–Hopfield fashion as

$$M_{ij} \equiv \left[J_{ij} - \frac{\mathrm{i}}{\omega}(\beta \tilde{W}_{ij} - \alpha J_{ij})\right] = J \sum_\mu \xi_i^\mu \xi_j^{\mu *}, \tag{8}$$

or, with ξ_i^μ expressed in terms of amplitudes and phases as $|\xi_i^\mu| \exp(-\mathrm{i}\phi_i^\mu)$,

$$J_{ij} = J \sum_\mu |\xi_i^\mu||\xi_j^\mu| \cos(\phi_i^\mu - \phi_j^\mu) \tag{9}$$

$$\beta \tilde{W}_{ij} = J \sum_\mu |\xi_i^\mu||\xi_j^\mu|[\omega \sin(\phi_i^\mu - \phi_j^\mu) + \alpha \cos(\phi_i^\mu - \phi_j^\mu)]. \tag{10}$$

Note that here both kinds of connections, J (excitatory-to-excitatory) and $\tilde{\mathsf{W}}$ (excitatory-to-inhibitory), are used to store the amplitude and phase patterns of the oscillation. J is symmetric, while $\tilde{\mathsf{W}}$ is not.

These connections can be obtained by an online algorithm, a simplified version of the full Hebbian learning treated by Liljenström and Wu [20, 21]. Suppose the cortex has effective oscillatory input $\delta \boldsymbol{I}^b = \xi^\mu \mathrm{e}^{-\mathrm{i}\omega t} + \xi^{\mu *} \mathrm{e}^{\mathrm{i}\omega t}$ during learning of the μth pattern. Here we make explicit the real nature of the signals. Suppose also that the J and $\tilde{\mathsf{W}}$ connections inactive, consistent with the picture proposed by Wilson, Bower and Hasselmo [17, 19], who suggested that learning occurs when the long-range intracortical connections are weakened by neuromodulatory effects. Then the linearized (4) and (5) are simply

$$\begin{aligned} \dot{u}_i + \alpha u_i &= -\beta v_i + \xi_i^\mu \mathrm{e}^{-\mathrm{i}\omega t} + \xi_i^{\mu *} \mathrm{e}^{\mathrm{i}\omega t} \\ \dot{v}_i + \alpha v_i &= \gamma u_i, \end{aligned} \tag{11}$$

with solution

$$\begin{aligned} u_i(t) &= \frac{-\mathrm{i}\omega + \alpha}{-\omega^2 + \alpha^2 + \beta\gamma - 2\mathrm{i}\alpha\omega} \xi_i^\mu \mathrm{e}^{-\mathrm{i}\omega t} + \text{c.c.} \\ v_i(t) &= \frac{\gamma}{-\omega^2 + \alpha^2 + \beta\gamma - 2\mathrm{i}\alpha\omega} \xi_i^\mu \mathrm{e}^{-\mathrm{i}\omega t} + \text{c.c.} \end{aligned} \tag{12}$$

where c.c. denotes complex conjugate. In other words, the cortical activities are clamped by the inputs.

For Hebbian learning, $\dot{J}_{ij} \propto u_i(t)u_j(t)$, and, after time averaging, $\delta J_{ij} \propto \int_0^{2\pi/\omega} u_i(t)u_j(t)\,\mathrm{d}t$, leading to

$$\begin{aligned}\delta J_{ij} &\propto \frac{\omega^2+\alpha^2}{|-\omega^2+\alpha^2+\beta\gamma-2\mathrm{i}\alpha\omega|^2}(\xi_i^\mu\xi_j^{\mu*}+\xi_i^{\mu*}\xi_j^\mu)\\ &= 2\frac{\omega^2+\alpha^2}{|-\omega^2+\alpha^2+\beta\gamma-2\mathrm{i}\alpha\omega|^2}|\xi_i^\mu||\xi_j^\mu|\cos(\phi_i^\mu-\phi_j^\mu).\end{aligned} \tag{13}$$

Similarly, $\delta W_{ij} \propto \int_0^{2\pi/\omega} v_i(t)u_j(t)\,\mathrm{d}t$ leading to

$$\begin{aligned}\delta W_{ij} &\propto \frac{\gamma}{|-\omega^2+\alpha^2+\beta\gamma-2\mathrm{i}\alpha\omega|^2}[(\mathrm{i}\omega+\alpha)\xi_i^\mu\xi_j^{\mu*}+(-\mathrm{i}\omega+\alpha)\xi_i^{\mu*}\xi_j^\mu]\\ &= \frac{2\gamma}{|-\omega^2+\alpha^2+\beta\gamma-2\mathrm{i}\alpha\omega|^2}[\omega|\xi_i^\mu||\xi_j^\mu|\sin(\phi_i^\mu-\phi_j^\mu)+\alpha|\xi_i^\mu||\xi_j^\mu|\cos(\phi_i^\mu-\phi_j^\mu)].\end{aligned} \tag{14}$$

Then, if the relative learning rates for J and $\tilde{\mathrm{W}}$ are tuned appropriately, we simply recover the formulae (9) and (10). In actual online learning, we can use high-pass versions of $\boldsymbol{u}$ and $\boldsymbol{v}$ to learn J and $\tilde{\mathrm{W}}$ to remove the baseline value, i.e., the operation point $\bar{\boldsymbol{u}}$ and $\bar{\boldsymbol{v}}$, which does not contain odour information.

To see the selective resonance explicitly, suppose that different patterns ξ^μ are orthogonal to each other. Let us denote the overlap $(1/N)\sum_i \delta I_i^{\mathrm{b}}\xi_i^{\lambda*}$ of the input δI_i^{b} with the stored pattern ξ_i^λ by δI^λ. Then, multiplying (7) by $\xi_i^{\lambda*}$ and summing on i, we find that at steady oscillatory state, the response $u^\lambda \equiv (1/N)\sum_i u_i\xi_i^{\lambda*}$ to pattern ξ^λ obeys

$$\dot{u}^\lambda = -(2\alpha-J)u^\lambda - \frac{\mathrm{i}}{\omega}(\beta\gamma+\alpha^2)u^\lambda + \frac{\mathrm{i}}{\omega}(-\mathrm{i}\omega+\alpha)\delta I^\lambda. \tag{15}$$

This is like an oscillator with oscillation frequency $(\beta\gamma+\alpha^2)/\omega$ and an effective oscillation decay rate $2\alpha - J$. It resonates to external oscillatory input of frequency $\omega \approx \sqrt{\beta\gamma+\alpha^2}$ with a steady state amplitude

$$u^\lambda = \frac{(-\mathrm{i}\omega+\alpha)\delta I^\lambda}{\beta\gamma+\alpha^2-\omega^2-\mathrm{i}\omega(2\alpha-J)} \approx \frac{(1+\mathrm{i}\alpha/\omega)\delta I^\lambda}{2\alpha-J}, \tag{16}$$

However, for an input δI_i^{b} orthogonal to all the stored patterns, $\delta I^\lambda = 0$ for all λ, and the resonance will be washed out when $J < 2\alpha$. For $J > 2\alpha$, the network will support spontaneous oscillations analogous to those in the bulb, but not as observed in the cortex. The effect of the long-range couplings, through the parameter J, is to reduce the damping in the circuit from 2α to $2\alpha - J$ when the input matches a stored pattern, thereby sharpening the resonance as $J \to 2\alpha$ while we keep $J < 2\alpha$. On the other hand, the resonant driving frequency depends only on the single-oscillator parameters α, β and γ.

This oscillatory associative memory enjoys the usual properties that characterize Hopfield networks [26], including rapid convergence (a few oscillation cycles if the presented pattern has reasonable overlap with a stored one), robustness with respect to noise and corrupted input, and a storage capacity of the order of N random patterns, where N is the network size.

2.3. Coupling between bulb and cortex

The model has both feedforward (bulb–cortex) and feedback (cortex–bulb) connections. The former transmit the bulbar encoding of the input odours to the cortex for recognition, while the latter permit segmentation by producing adaptation to recognized odour objects.

Bulb to cortex. As mentioned in the introduction, in the real cortex, the excitatory cells receive input from the bulb both directly from the fibres of the lateral olfactory tract and in a slower pathway via feedforward inhibitory interneurons in the cortex. We model this in the following way. The synapses from local bulb populations j to local cortical populations i are specified by a matrix $C_{ij}^{\mathrm{b}\to\mathrm{c}}$. The values of these connections are not important in the model, and very little is known about them, so we will take them to be random. The resulting signals are then fed to the excitatory cells, both directly and, with the opposite sign, through a parallel low-pass filter, representing the effect of the feedforward inhibitory cells; see figure 1. Details are given in the appendix.

The combination of the direct excitatory and low-pass filtered inhibitory signals makes the feedforward pathway act as a high-pass filter, partially cancelling the slow part $\boldsymbol{I}^{\mathrm{b}0}$ of the bulb output from the cortical input. Consequently, the net input to the cortical excitatory units is dominated by the oscillatory component of the bulb activity, which encodes information about the odour input. (We do not know how well such a cancellation is actually achieved in real olfactory systems, but this could be tested experimentally.)

Cortex to bulb. The odour-specific adaptation that forms the basis for odour segmentation in our model is implemented using a feedback signal from the cortex to the granule units of the bulb. We do not know how such a signal is generated in animals, or even whether it is, although anatomically such a pathway exists. If the signal does exist, it likely also involves areas such as entorhinal cortex, which contributes to the centrifugal input to the bulb. These areas lie outside the scope of the present model, so we will simply construct a suitable signal and explore the consequences.

In exploratory computations, we have found that this form of feedback control only works if the signal is slowly varying in time (on the order of the sniff-cycle time or slower). Merely feeding back the oscillating cortical activities does not appear to permit any kind of robust stimulus-specific adaptation.

Thus, we generate the feedback signal in the following *ad hoc* fashion: first each excitatory cortical output $g_u(u_i)$ is run through a threshold-linear element to remove its non-oscillatory part, which carries no odour information. Then the output of this element is run through a low-pass filter. The time constants of this filter are on the order of the sniff cycle or longer. The net result is a signal pattern which takes a sniff-cycle time or so to grow to full strength. The signal component from excitatory unit i will be proportional to the amplitude of the oscillation of that unit, so this signal will contain information about the odour that evoked the cortical oscillation pattern. The explicit form of the equations used to generate the feedback signal in the simulations is given in the appendix.

Since we rectify and low pass only the excitatory cortical outputs $g_u(u_i)$, the feedback signal includes only the odour information coded in the amplitude but not in the phase pattern of the cortical oscillation. Phase information could be included by (for example) feeding the difference signals $g_u(u_i) - g_u(u_j)$ through the rectification and low-pass processes. However, we have not explored such mechanisms in this work.

The granule units in the bulb respond to the feedback signals by changing their activities proportional to it. These changes are then transmitted to the mitral cells by the synaptic matrix H. As shown by Li [25], a feedback signal

$$\boldsymbol{F} \propto \mathsf{H}^{-1}\boldsymbol{I}, \tag{17}$$

will, when transmitted onward to the mitral units, cancel the odour inputs to the bulb (in linear approximation).

In our model we want to make this cancellation work for all the odour patterns stored in

the cortex. Denoting by G_j^μ the rectified and low-passed cortical output when the system is stimulated by odour pattern I_k^μ, this can be achieved by a Hebbian feedback connection matrix $\mathsf{C}^{\mathrm{c}\to\mathrm{b}}$ that maps $\boldsymbol{G}^\mu$ to feedback signal $\boldsymbol{F}^\mu$ for each odour μ in a single layer network:

$$C_{ij}^{\mathrm{c}\to\mathrm{b}} \propto \sum_\mu F_i^\mu G_j^\mu = \sum_k H_{ik}^{-1} \sum_\mu I_k^\mu G_j^\mu. \tag{18}$$

3. Simulations

We have simulated a network with bulb and cortical modules each consisting of 50 excitatory and 50 inhibitory units. They were coupled as described in section 2.3 and the appendix. The coupled differential equations were integrated using a fourth-order Runge–Kutta routine from *Numerical Recipes* [27].

We used three random odour input patterns I_i^μ. Their elements were drawn independently for each i and μ from a uniform distribution on (0,1]. The elements of the granule-to-mitral synaptic matrix H were taken to have the form $H_{ij} = \text{const.} \cdot \delta_{ij}$. We designed the mitral-to-granule matrix W so as to make the bulb oscillate in response to the three input patterns, taking $W_{ij} \propto \mathrm{Im} \sum_{\mu=1}^{3} \zeta_j^\mu \zeta_j^{\mu *}$. Here the ζ_i^μ are complex, with amplitudes resembling the input odour patterns I_i^μ and with random phases. Since W should have non-negative elements, we simply zeroed out the negative W_{ij} in the construction†. This dilution did not affect the bulb oscillations qualitatively. Other parameters were set as in [24], so the evoked oscillations were in the 40 Hz range.

The cortical design followed section 2.2. The local couplings β^0 and γ^0 were chosen so that the cortical oscillation frequency roughly matched the bulbar one, i.e., $\beta^0\gamma^0 + \alpha^2 \approx \bar{\omega}^2$ (see equation (16), where $\bar{\omega}$ is the average oscillation frequency in the bulbar outputs). The inhibitory units had the sigmoidal activation function used in the model of the bulb [24]. In some of our simulations, the activation function of the excitatory units also had this form. In obtaining the results presented here, however, we used a piecewise linear activation function with gains of 1 and 2, respectively, in the low- and high-input regions above threshold. This choice was made only for convenience in analysing the nonlinear dynamics and is not essential for the function of the network.

The cortical synaptic matrices J and $\tilde{\mathsf{W}}$ were designed to store oscillation patterns for two of the three odour input patterns, in the following way. For each of the two odours, we stimulated the bulb with its input pattern I_i^μ and fed the resulting oscillatory bulb output through the bulb-to-cortex matrix $C_{ij}^{\mathrm{b}\to\mathrm{c}}$ and the subsequent high-pass filtering operation to the cortex, with the intracortical connections J^0 and $\tilde{\mathsf{W}}^0$ set to zero. The resulting oscillation patterns in the cortical units for the two odours were then used as ξ^μ in constructing J and $\tilde{\mathsf{W}}$.

We modified the Hebb rule (equation (8) or equations (9) and (10)) slightly, using, instead, a pseudoinverse formula

$$M_{ij} = J \sum_\mu \xi_i^\mu \eta_j^{\mu *}, \tag{19}$$

where $\sum_i \eta_i^{\mu *} \xi_i^\nu = N\delta_{\mu\nu}$. This was done only to reduce finite-size effects due to mutual overlaps (of order $\sqrt{N}$) between patterns, and would be inessential in sufficiently large networks.

† In the bulb model of Li and Hopfield [24, 25], the idea was that extensive asymmetric random synapses would, for a large network, automatically generate a distributed encoding of odours in the amplitudes and phases of oscillation patterns in the network. Here we will be more concerned with how these patterns are processed by the cortex, so, for convenience, we have engineered particular amplitude patterns in through the bulbar W matrix in this fashion. However, the particular forms used for H and W are not important for the problem that we are studying here, as long as $\mathsf{A} \equiv \mathsf{HW}$ is sufficiently asymmetric.

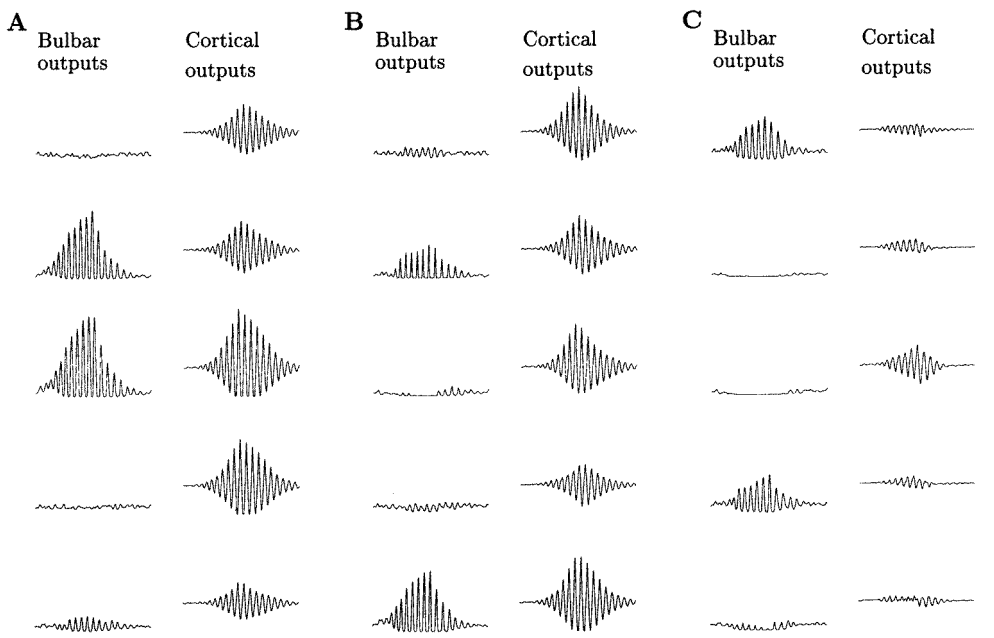

Figure 2. Panels (A)–(C): bulbar and cortical oscillation patterns for odours A, B (stored in the associative memory in the cortex) and C (not stored). In each pattern, we plot temporal traces of outputs from five of the 50 mitral or cortical excitatory units during one sniff cycle lasting 370 ms, roughly the first half of which is inhalation. Note the modulating of oscillation by the sniff cycle, and the different oscillation amplitudes for different units. Oscillation phases also differ between units, though they are not apparent in the figure. The same format is used to display bulbar and cortical responnses in the following figures. Cortex-to-bulb feedback is turned off for the results shown in this figure. Note that the cortex responds little to odour C, since the input does not match any of the stored oscillation patterns.

As explained in section 2.3 and the appendix, the slowly varying feedback signal used for the odour-specific adaptation was generated by a threshold-linear rectification, followed by a pair of simple linear filters. The time constants of these (3 and 0.3 s, respectively) would made it take 10–12 256 ms sniff cycles to generate a full strength feedback signal if the cortical signal were held constant. Similarly, the adaptation takes just as long to wear off after the stimulus is removed.

Like the intracortical M matrix, the cortex-to-bulb matrix $\mathsf{C}^{\mathsf{c}\to\mathsf{b}}$ was modified using the projection-rule algorithm to eliminate finite-size overlap effects between the cortical oscillation patterns of different odours. Thus, in formula (18), we replaced the rectified and low-pass-filtered cortical patterns G_j^μ by $\tilde{G}_j^\mu$, where $\tilde{\boldsymbol{G}}^\mu$ are vectors such that $\tilde{\boldsymbol{G}}^\mu \cdot \boldsymbol{G}^\nu = N\delta^{\mu\nu}$.

Figure 2 shows the bulbar and cortical oscillatory response patterns evoked on five of the 50 mitral or cortical excitatory units by three odours: A, B, and C. Only odours A and B are stored in the cortical memory in the J and $\tilde{\mathsf{W}}$ matrices. Different amplitude response patterns to different odours are apparent. The cortex resonates appreciably to only odours A or B, but not to C, demonstrating the selectivity of the cortical response.

Figure 3 demonstrates odour adaptation to odour A. The response amplitudes decay quicky in successive sniffs, although the oscillation patterns do not change appreciably before the amplitudes decay to zero. The way this comes about is that the feedback signal generated by A, when relayed by the granule cells to the mitral ones, creates an effective extra input signal

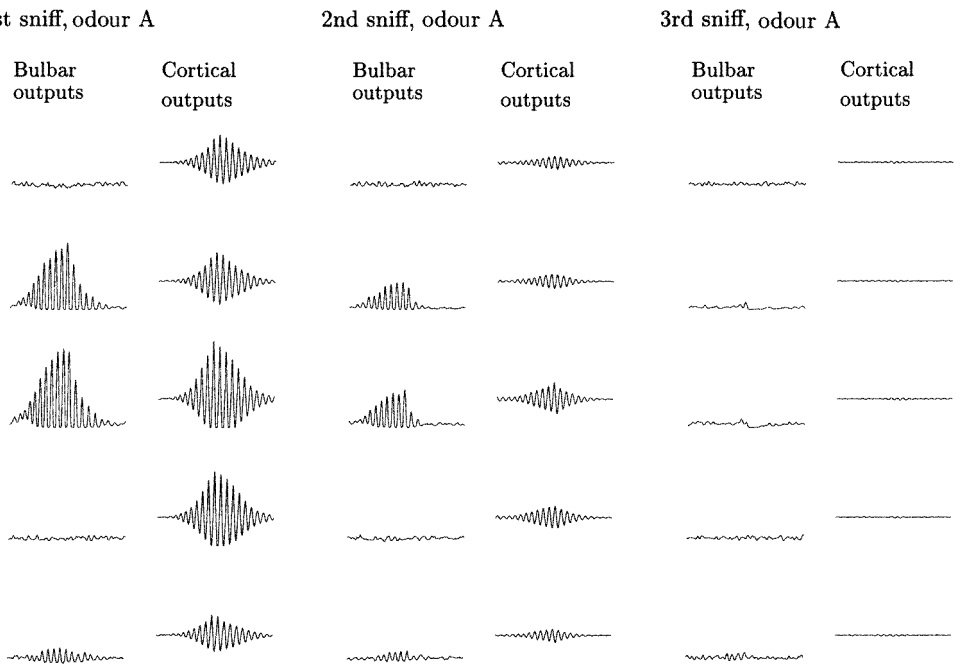

Figure 3. Demonstrating the adaptation to odour A, with the feedback from cortex to bulb active. Plotted are the responses to odour A alone during three successive sniffs. Note that the response magnitudes decay in successive sniffs, but the response pattern, in particular the relative amplitude pattern, stays roughly the same from the first to second sniff before responses disappear at the third sniff.

$\bar{\text{A}}$ (anti-A), and by the third sniff A + $\bar{\text{A}} \approx 0$.

To quantify the similarity between oscillation patterns, we extract an ($N = 50$)-dimensional complex vector $\boldsymbol{O}$ from the temporal Fourier transform of the activity of the cortical excitatory units during the sniff cycle, with the component O_i specifying the amplitude and phase of the oscillations in excitatory unit i. We can measure the similarity between $\boldsymbol{O}$ and $\boldsymbol{O}'$ by the normalized overlap $S_{OO'} = |\langle \boldsymbol{O}|\boldsymbol{O}'\rangle/(|\boldsymbol{O}| \cdot |\boldsymbol{O}'|)$, which is 1 for $\boldsymbol{O} \propto \boldsymbol{O}'$ and near zero ($O(1/\sqrt{N})$) for two unrelated patterns. Calling the pattern vectors $\boldsymbol{A}^0$, $\boldsymbol{A}^1$, $\boldsymbol{A}^2$, and $\boldsymbol{A}^3$ for cortical response to odour A without adaptation and during the first, second, and third sniff cycles of the adaptation respectively, we find $S_{\text{A}^0\text{A}^1} = 0.9997$, $S_{\text{A}^0\text{A}^2} = 0.992$, and $S_{\text{A}^0\text{A}^3} = 0.74$, with response amplitudes $|\boldsymbol{A}^1|/|\boldsymbol{A}^0| = 0.97$, $|\boldsymbol{A}^2|/|\boldsymbol{A}^0| = 0.3$, $|\boldsymbol{A}^3|/|\boldsymbol{A}^0| = 0.08$. Thus, the strength of the response is already significantly weakened after one sniff, but its cortical pattern of variation remains undistorted through several sniffs.

The way this adaptation varies in successive sniffs depends on both the time constants in the feedback circuitry (as discussed above) and the strength of the filtered signal fed back to the bulb. In the simulations shown here, the latter was strong enough that even after one sniff, a large fraction of the input signal is cancelled by the feedback, and after two sniffs the cancellation was nearly complete. A smaller feedback strength and a correspondingly longer time constant of the feedback circuitry would make it take longer for the adaptation to set in. Similarly, the time it takes for the adaptation, once established, to wear off is set by the same time constants (for the values used here, around 3 s or 12 sniff cycles).

Figure 4(*a*) demonstrates the segmentation capability of the system. The response $\boldsymbol{B}^{\text{seg}}$

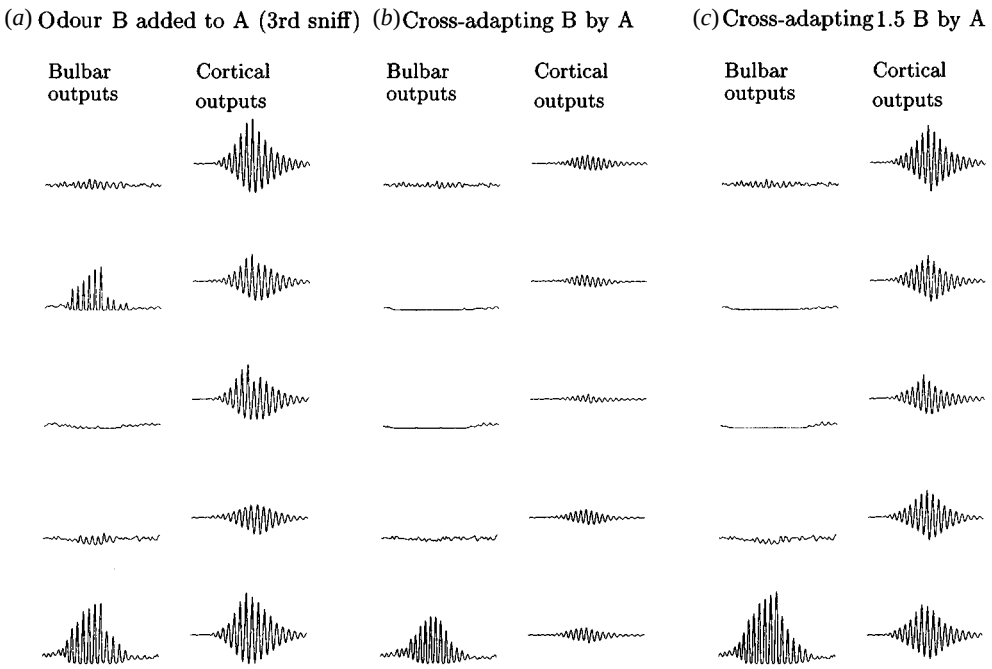

Figure 4. (*a*) Segmenting odours A and B. After two sniffs of A as in figure 3, odour B is added, so the net input is A + B. The response is almost the same as that to B alone (figure 2, middle). (*b*) Cross-adaption: response to odour B after odour A was present in two previous sniffs and then withdrawn. The response is weak and distorted. (*c*) Same as (*b*), except that an odour B 1.5 times as strong is used. This strength is sufficient to evoke a stronger, less distorted response.

to the odour mixture A + B at the third sniff after the first two sniffs of odour A is quite similar to that, $\boldsymbol{B}^0$, to odour B alone: $S_{\mathrm{B}^{\mathrm{seg}}\mathrm{B}^0} = 0.993$, and $|\boldsymbol{B}^{\mathrm{seg}}|/|\boldsymbol{B}^0| = 0.91$. Thus, although A is still present, so is the anti-A, so the net signal to the mitral units is $\mathrm{A} + \bar{\mathrm{A}} + \mathrm{B} \approx \mathrm{B}$. This demonstrates odour-specific adaptation in the model. The system responds with the activity pattern characterizing the new odour, essentially undistorted by the existing odours in the environment, thus effectively achieving odour segmentation. Odour B can be segmented as long as it enters after the adaptation to A is established, in this model at the third or any subsequent sniff.

However, if odour A is suddenly withdrawn at the start of the third sniff, when odour B is introduced, the system response to odour B is weakened and distorted (this is particularly noticable in the bulbar responses). The reason for this is that the effective total input is now $\mathrm{B} + \bar{\mathrm{A}} \approx \mathrm{B} - \mathrm{A}$, which is not at all like B (figures 4(*b*) and (*c*)). This corresponds to the psychophysically observed cross-adaptation—after sniffing one odour, another odour at next sniff smells less strong than it normally would and may even smell different [14]. In the normal olfactory environment, however, such sudden and near complete withdrawal of an odour seldom happens. Let $\boldsymbol{B}^{\mathrm{cross}}$ and $\tilde{\boldsymbol{B}}^{\mathrm{cross}}$ be the cortical response vectors to cross adapted odour B and odour 1.5B. Comparing with the response to odour B alone, we find $S_{\mathrm{B}^0\mathrm{B}^{\mathrm{cross}}} = 0.94$, $|\mathrm{B}^{\mathrm{cross}}|/|\mathrm{B}^0| = 0.23$; $S_{\mathrm{B}^0\tilde{\mathrm{B}}^{\mathrm{cross}}} = 0.97$, $|\tilde{\mathrm{B}}^{\mathrm{cross}}|/|\mathrm{B}^0| = 0.74$. We can understand these results in the following way. The feedback input $\bar{\mathrm{A}} \approx -\mathrm{A}$ acts to move the bulb operating point $\bar{x}_i$ to lower gain values (for units where I_i^{A} is strong), thereby weakening the overall response. For normal-strength B, most of the mitral units in the bulb do not respond much, so the cortical

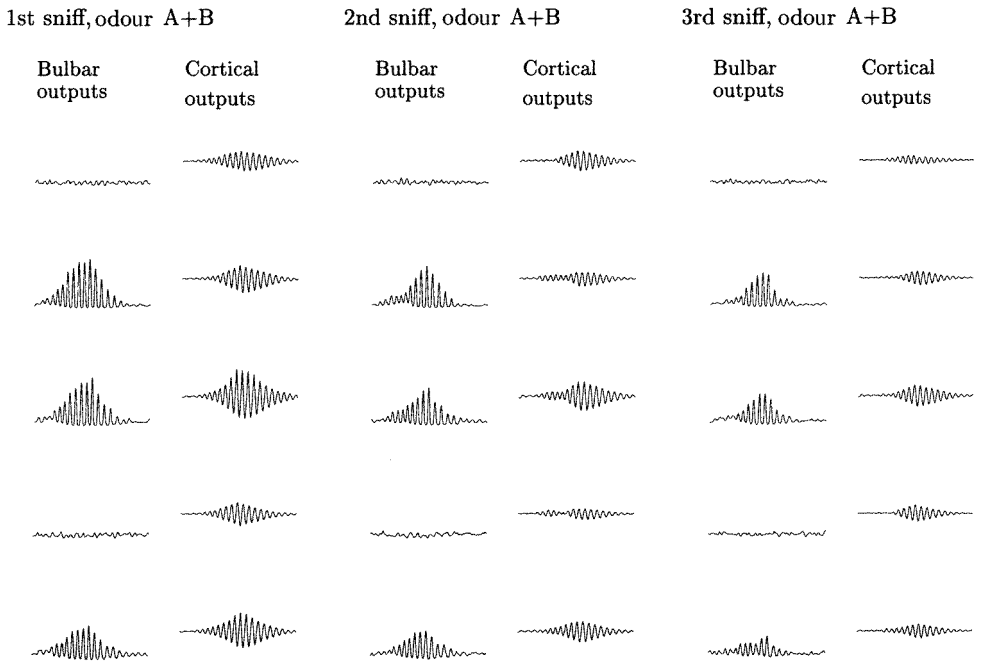

Figure 5. This figure illustrates how adaptation in the model is not effective for the mixture odour (A + B)/2. Responses to this odour are shown for three successive sniff cycles. The cortical response, although initially weaker than that to pure A or B (figure 2) is still appreciable at the third sniff (compare with adaptation to odour A in figure 3).

response is correspondingly weak and distorted relative to that to B in the absence of adaptation. The stronger input 1.5B evokes a stronger bulb response, however, and the cortical response is stronger and better (but still imperfectly) correlated with the unadapted pattern.

Since the olfactory bulb is nonlinear, the odour mixture A + B does not induce a bulbar response equal to the sum of the responses to A and B individually. Consequently, the unadapted cortical response to it (figure 5, left panel) is weaker than that to A or B (the bulb response to the mixture is not embedded in the cortical connections) and not strongly correlated with the responses to the pure odours. The situation is similar to that for any other unstored odour, such as C (figure 2, (C)), to which there is almost no adaptation in the bulb because there is almost no cortical signal to feed back. The unadapted cortical response to A + B is stronger than that to C because the nonlinearity in the bulb here is not strong enough to completely destroy correlations between its reponses to the individual odours A and B and that to their mixture. Nevertheless, the weakness of the cortical response reduces the feedback to the bulb significantly, and the system does not adapt to the mixture as effectively as to individual odours, as shown in the middle and right panels of figure 5. We also note that because the feedback is weak, the attenuation of the signals in both bulb and cortex is also weaker than for pure stored odours (cf figure 3). Thus, the cortical response to the mixed odour, while initially weaker than that to pure stored ones, lasts longer.

4. Discussion

We have presented a computational model for an olfactory system that can detect, recognize and segment odours. Detection is performed in the bulb, which encodes odours in oscillatory activity patterns. Recognition is carried out by the cortex using a resonant associative memory mechanism. Finally, segmentation is implemented by a slowly varying feedback signal which acts to cancel the specific input that evoked the resonant cortical response.

The model is constrained by a few basic anatomical and physiological facts: odours evoke oscillatory activity in populations of excitatory and inhibitory neurons in both bulb and cortex, these two structures are coupled by both feedforward and feedback connections, reducing the cortical feedback enhances the bulbar responses, and the system exhibits odour-specific adaptation. Within these constraints, we have tried to build a minimal model. We have taken the bulb module from earlier work by one of us [24, 25] and augmented it with a model of the pyriform cortex and with feedforward and feedback connections between it and the bulb. We have ignored many further known details of real olfactory systems that do not bear directly on the fundamental property of stimulus-specific adaptation, and when we have had to go beyond current knowledge (as in constructing the feedback signal) we have done so in a purely phenomenological way, avoiding hypothesizing specific details unrelated to the function of the system. From the analysis of the model and the simulations we can see how the basic computations necessary for olfactory segmentation might be carried out by the neural networks of the bulb and cortex.

But do real olfactory systems actually function in this way? This can be tested at the level of both the assumptions we put into the model and the properties we find for it. First of all, we have assumed that the feedback from cortex to bulb is slowly varying (i.e. that firing rates for the feedback fibres vary on the timescale of the sniff cycle, but not of the oscillations found in both the bulb and cortex). Furthermore, we have assumed that this feedback is odour-specific. While the existence of some feedback is well established, neither of these specific hypotheses has been tested experimentally. However they both could be.

Properties we find in the model, beyond the fact that it successfully implements segmentation, can also be tested. These include the following.

First, the fact that the feedback signal requires strong cortical activity to drive it means that unfamiliar (unlearnt) odours will not be adapted to as strongly as familiar ones, so they will not be so easily segmented from subsequently presented ones. As we saw in figure 5, this expectation also applies to unfamiliar mixtures of familiar odours. Furthermore, as we also noted, we expect the weakening of the (initially weaker) responses with adaptation to be slower for such mixtures than for familiar odours.

Second, cross-adaptation, as illustrated in figure 4, is a necessary consequence of the slow feedback: the effective bulb input $\bar{\mathrm{A}} \approx -\mathrm{A}$, from the previous presence of the adapting stimulus, will be present for some time (depending on the time constants of the feedback circuitry) whether odour A remains in the environment or not. Thus the total input to the bulb with A still present will be very different from that with A suddenly removed. If there is odour-specific adaptation of the kind necessary to perform segmentation when A remains in the environment (A cancelled by $\bar{\mathrm{A}}$), then a different response must occur when A is withdrawn. Present evidence on cross-adaptation is rather limited, but psychophysical and electrophysiological investigation of this phenomenon would be helpful in pinning down quantitatively the time constants of the circuitry involved in segmentation.

If odour-specific adaptation is not implemented using our cortical feedback mechanism, how else might it be done? One possibility to consider is single-unit-level adaptation (or fatigue), which can be implemented in a network like ours by making the threshold for each

unit dependent on its own recent activity. In a model with the structure of ours (with bulb and cortical modules) but without feedback, such fatigue would have to be implemented in the bulb; otherwise the activity there would not exhibit adaptation. This presents a problem if the activity patterns of different odours overlap significantly—it is not evident that one can avoid changing the response to a new odour when some of the units active in the normal response to it are to be fatigued. Indeed, in investigations of simple oscillatory associative memories with such adaptation [28], temporal segmentation has been found only for patterns with rather weak mutual overlap. This overlap will be weak for sparse patterns, but it is not clear how sparse real evoked bulb and cortical activity patterns are, when looked at at the level of resolution of the units in our model.

This problem is not present for the mechanism we propose, in which bulb units themselves are not fatigued. Rather, the mechanism cancels the input to bulb units in exactly the degree that they receive input from the adapting odour. It is as if every receptor activated by an odour became adapted by an amount exactly equal to its initial response.

In our model, the feedback connections to the inhibitory bulb units have to have just the right values to produce the necessary cancellation. In real olfactory systems, the strengths of the centrifugal synapses on granule cells are presumably determined by some learning mechanism, and for our model to apply it is necessary that this mechanism find the right values for them. As we know nothing about this mechanism, here in our model we just assumed the necessary form. This form has a degree of plausibility because it is Hebbian, but very little is known yet about learning in these synapses. Investigations could shed important light on the validity of this key element of the model.

Another plausible mechanism, which could implement odour-specific adaptation in the bulb in more or less the right manner, is adaptation of receptor–bulb synapses in such a way that the inputs to bulb capture mainly the transient but not static odour inputs. This would reduce the input signal for the adapting odour directly, at just the right places, and so does not suffer from the problems that single-unit fatigue in the bulb does. However, there is a simple difference between the predictions of such a model and ours, since in ours the cortex, functioning as an associative memory, only sends its feedback to the bulb (or only sends it at full strength) for learnt odours. The receptor–mitral synaptic adaptation model would exhibit the same degree of odour-specific adaptation for all odours, learnt or not. Of course, both mechanisms could be present, and the difference could be large or small according to the relative size of the two contributions.

The fact that we have employed both excitatory-to-excitatory (J) and excitatory-to-inhibitory ($\tilde{W}$) cortical connections enhances the associative memory function by permitting oscillation patterns which differ in phase as well as amplitude. This is of no help for selective adaptation in the model as described here, since phase information is lost in the generation of the feedback signal, but this information could be retained using more elaborate mechanisms, as mentioned in section 2.3.

It is not clear whether real olfactory systems code odours in the phases of their oscillation patterns. However, in any case, a restricted version of our cortex, without $\tilde{W}$, could function with only amplitude-modulated patterns, similarly to the model of Wang *et al* [28]. The addition of intrinsic oscillatory properties for individual units or, implicitly, the individual neurons in the populations they represent, would not change the properties of such a network qualitatively.

The three tasks carried out by the system—detection, recognition, and segmentation—are computationally linked. For example, even if an ambiguous or weak odour is ‘recognized’ by the pyriform cortex in the sense that a characteristic oscillatory response is evoked there, that response may be too weak to suppress further bulbar response. Then the system will continue

to respond to the odour in the same way as if it had not recognized it; that is, the odour-specific adaptation necessary for segmentation can be seen as part of the recognition process.

While our units correspond to functional groups of neurons in real olfactory systems, our model is of higher resolution than that of Ambros-Ingerson *et al* [18]. While we emphasize the coding of odour information in distributed oscillation patterns, their model contains no explicit treatment of dynamics on the 40 hz timscale or of the temporal segmentation problem. They address instead a higher-level problem (hierarchical odour classification) with a higher-level model. In such more complex situations, cortex-to-bulb feedback could be a more general, active phenomenon than in the limited-scope problem we consider, but we do not address such issues here.

Our network performs what might be called 'the simplest cognitive computation'. It is natural to expect that evolution has employed elaborations on this structure in other sensory systems and in central processing. For example, hippocampal processing also employs oscillations, long-range intra-area associative connections, and feedback [29,30]. In another context, work by one of us [31] on visual processing suggests a function for slow feedback to inhibitory neurons from higher areas in modulating the computations carried out in area V1. Our hope is that studying and modelling the olfactory system in the way we have done here will lead to insights into aspects of top-down/bottom-up interactions in other cognitive computations.

Appendix: Bulb–cortex coupling: implementation details

Feedforward

In the feedforward pathway from bulb to cortex, the mitral unit outputs $g_x(x_i)$ are fed both directly to the excitatory cortical units and in parallel, indirectly via feedforward inhibitory units. The process, as indicated schematically in figure 1, is described by the equations

$$L_i = \sum_j C_{ij}^{\mathrm{b}\to\mathrm{c}} g_x(x_j) \tag{20}$$

$$\dot{z}_i = -\alpha_{\mathrm{ff}} z_i + L_i \tag{21}$$

$$I_i^{\mathrm{b}} = L_i - \sigma g_z(z_i). \tag{22}$$

Here L_i is the input signal to the cortical location i, $\mathsf{C}^{\mathrm{b}\to\mathrm{c}}$ is the connection matrix that transforms the mitral outputs to the cortical inputs, z_i are the membrane potentials of the feedforward inhibitory units, $g_z(.)$ is their activation function and $\alpha_{\mathrm{ff}}^{-1}$ is their time constant. I_i^{b} is then the total input signal to the ith cortical excitatory unit in equation (4). In general, this input contains both slowly varying and rapidly oscillating components. The pathway via the inhibitory feedforward units acts like a low-pass filter. Thus, the net effect is that the rapidly varying or high-frequency components, which contain the odour information, are transmitted to the cortex.

In the simulations reported in section 3, we took $g_z(.)$ to have two regions of different gain values, with a smaller gain at smaller input. We designed σ and the parameters of $g_z(.)$ so that the net slow component of I_i^{b} pushed the cortical operation points $\bar{u}_i$ and $\bar{v}_i$ to stable values close to, but below, their high gain region. Thus the cortex had a stable operating point, enabling it to carry out its associative memory function more cleanly that without this engineering refinement.

We make no claims about biological realism for the details of the feedfoward mechanism. However, some kind of effective high-pass filter is essential to the robust functioning of the model. Further experimental investigation of the dynamical properties of the feedfoward pathway would be important for understanding how it actually works.

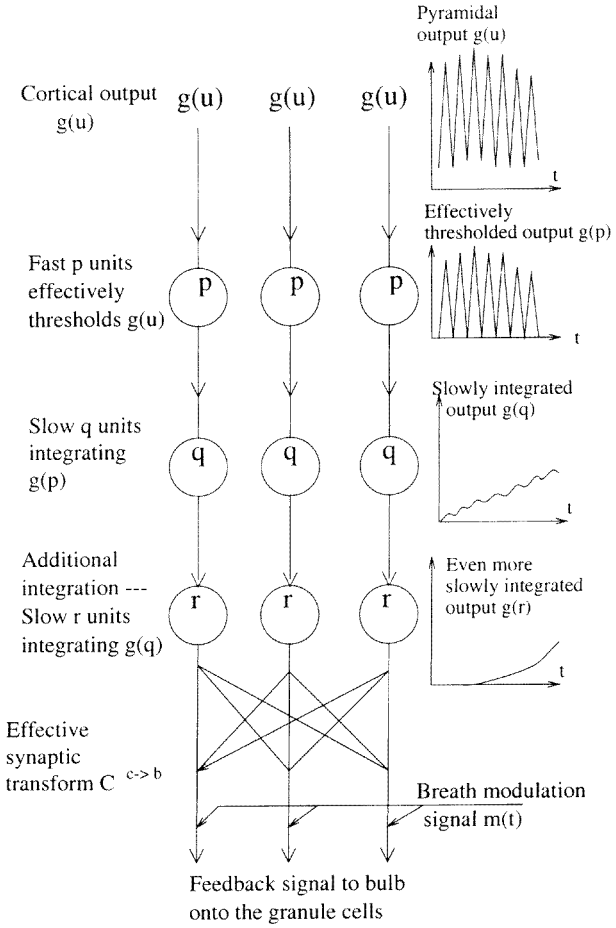

Figure A1. Details of the feedback route. Only the oscillatory components of the cortical outputs $g_u(u)$ contain the odour information. This component is extracted by half-wave rectification by the p units. The oscillatory $g_p(p)$ is converted to slowly varying signals by two successive slow temporal integrating units q and r. The resulting signal $r(t)$ is fed through the matrix $\mathsf{C}^{\mathrm{c}\to\mathrm{b}}$ and modulated with the breathing cycle by a signal $m(t)$ to produce the odour-specific feedback signal to the bulbar granule units. The temporal characteristics of signals from different units are depicted schematically on the right.

Feedback

To generate the half-wave rectified, low-passed feedback signal to the bulb from the cortical excitatory unit outputs, we use three successive groups of units followed by a synaptic matrix, as shown in figure A1:

$$\dot{p}_i = -\alpha_{\mathrm{fast}} p_i + g_u(u_i), \qquad \dot{q}_i = -\alpha_{\mathrm{slow}} q_i + g_p(p_i), \qquad \dot{r}_i = -\alpha'_{\mathrm{slow}} r_i + q_i, \tag{23}$$

$$I_i^{\mathrm{c}} = m(t) \sum_j C_{ij}^{\mathrm{c}\to\mathrm{b}} g_r(r_j), \tag{24}$$

where $m(t)$ is a modulating signal that synchronizes with breathing, increasing during inhalation and decreasing during exhalation.

With a short time constant $1/\alpha_{\mathrm{fast}}$ and a strong nonlinear g_p, the p_i unit has a output $g_p(p_i)$ which is effectively $g_u(u_i)$ thresholded above the average signal level. This 'rectified' output

is then transformed by the two subsequent units q_i and r_i, both with long time constants $1/\alpha_{\text{slow}}$ and $1/\alpha'_{\text{slow}}$, into a slowly varying signal, which is modulated by a function $m(t)$ (representing the breathing rhythm of the animal) and fed back via the connections $C^{c\to b}$ to produce the centrifugal input $\boldsymbol{I}^c$ to the granule units in the bulb.

It is not necessary to use two low-pass filter operations; the model works qualitatively the same with just one. However, adding the second one delays the feedback signal somewhat, giving the oscillations time to establish themselves before the feedback begins to act.

In a more complete model, the large time constants $1/\alpha_{\text{slow}}$ and $1/\alpha'_{\text{slow}}$ could emerge as a dynamic network property of secondary olfactory areas. Similarly, the modulating signal $m(t)$ could arise from additional signals from other parts of the brain.

References

[1] Laing D G 1995 Perception of odour mixtures *Handbook of Olfaction and Gustation* ed R L Doty (New York: Dekker) pp 283–98

[2] Horn D and Usher M 1991 Parallel activation of memories in an oscillatory neural network *Neural Comput.* **3** 31–43

[3] Horn D, Sagi D and Usher M 1991 Segmentation, binding and illusory conjunctions *Neural Comput.* **3** 510–25

[4] Shepherd G M 1990 Computational structure of the olfactory system, in *Olfaction—A Model System for Computational Neuroscience* ed J L Davis and H Eichenbaum (Cambridge, MA: MIT) pp 225–50

[5] Shepherd G M 1990 *The Synaptic Organization of the Brain* 3rd edn (Oxford: Oxford University Press)

[6] Freeman W J 1978 Spatial properties of an EEG event in the olfactory bulb and cortex *Electroencephalogr. Clin. Neurophysiol.* **44** 586–605

[7] Freeman W J and Schneider W 1982 Changes in spatial patterns of rabbit olfactory EEG with conditioning to odours *Psychophysiology* **19** 44–56

[8] Freeman W J and Grajski K A 1987 Relation of olfactory EEG to behaviour: factor analysis *Behav. Neurosci.* **101** 766–77

[9] Adrian E D 1950 Sensory discrimination with some recent evidence from the olfactory organ *Br. Med. Bull.* **6** 330–1

[10] Freeman W J and Skarda C A 1985 Spatial EEG patterns, non-linear dynamics and perception: the Neo-Sherrington view *Brain Res. Rev.* **10** 147–75

[11] Gelperin A and Tank D W 1990 Odour-modulated collective network oscillations of olfactory interneurons in a terrestrial mollusc *Nature* **345** 437–40

[12] Delaney K R, Gelperin A, Fee M S, Flores J A, Gervais R, Tank D W and Kleinfeld D 1994 Waves and stimulus-modulated dynamics in an oscillating olfactory network *Proc. Natl Acad. Sci. USA* **91** 669–73

[13] Bressler S L 1988 Changes in electrical activity of rabbit olfactory bulb and cortex to conditioned odour stimulation *Behav. Neurosci.* **102** 740–7

[14] Moncrieff R W 1967 *The Chemical Senses* 3rd edn (Boca Raton, FL: Chemical Rubber Company)

[15] Ma M, Leinders-Zufall T, Sheperd G M and Zufall F 1997 Two forms of odour adaptation in single olfactory receptor neurons *Soc. Neorosci.* **23** 741 (Abstracts)

[16] Haberly L B 1985 Neuronal circuitry in olfactory cortex: anatomy and functional implications *Chem. Senses* **10** 219–38

[17] Wilson M A and Bower J D 1992 Cortical oscillations and temporal interactions in a computer simulation of piriform cortex *J. Neurophysiol.* **67** 981–95

[18] Ambros-Ingerson J, Granger R and Lynch G 1990 Simulation of paleocortex performs hierarchical clustering *Science* **247** 1344–8

[19] Hasselmo M E 1993 Acetylcholine and learning in a cortical associative memory *Neural Comput.* **5** 32–44

[20] Liljenström H and Wu X-B 1995 Noise-enhanced performance in a cortical associative memory model *Int. J. Neural Syst.* **6** 19–29

[21] Liljenström H 1995 Autonomous learning with complex dynamics *Int. J. Intell. Syst.* **10** 119–53

[22] Gray C M and Skinner J E 1988 Centrifugal regulation of neuronal activity in the olfactory bulb of the waking rabbit as revealed by reversible cryogenic blockade *Exp. Brain Res.* **69** 378–86

[23] Wilson H R and Cowan J D 1972 *Biophys. J.* **12** 1–24

[24] Li Z and Hopfield J 1989 A model of the olfactory bulb and its oscillatory processing *Biol. Cybern.* **61** 379–92

[25] Li Z 1990 A model of odour adaptation and sensitivity enhancement in the olfactory bulb *Biol. Cybern.* **62** 349–61

[26] Hopfield J 1982 *Proc. Natl Acad. Sci. USA* **79** 2554–8
Hopfield J 1984 *Proc. Natl Acad. Sci. USA* **81** 3088–92
[27] Press W H, Flannery B P, Teukolsky S A and Vetterling W T 1988 *Numerical Recipes in C* (Cambridge: Cambridge University Press)
[28] Wang D, Buhmann J and van der Malsburg C 1990 Pattern segmentation in associative memory *Neural Comput.* **2** 94–106
[29] Rolls E T and Treves A 1998 *Neural Networks and Brain Function* ch 6 (Oxford: Oxford University Press)
[30] Hasselmo M E 1995 Neuromodulation and cortical function: modelling the physiological basis of behaviour *Behav. Brain Res.* **67** 1–27
[31] Zhaoping Li 1998 A neural model of contour integration in the primary visual cortex *Neural Comput.* **10** 903–40