OLFACTORY NETWORK DYNAMICS AND THE CODING OF MULTIDIMENSIONAL SIGNALS

Gilles Laurent

The brain faces many complex problems when dealing with odorant signals. Odours are multidimensional objects, which we usually experience as unitary percepts. They are also noisy and variable, but we can classify and identify them well. This means that the olfactory system must solve complicated pattern-learning and pattern-recognition problems. I propose that part of the solution relies on a particular architecture that imposes a dynamic format on odour codes. According to this hypothesis, the olfactory system actively creates a large coding space in which to place odour representations and simultaneously optimizes their distribution within it. This process uses both oscillatory and non-periodic dynamic processes with complementary functions: slow non-periodic processes underlie decorrelation, whereas fast oscillations allow sparsening and feature binding.

Olfactory systems have evolved over millions of years to solve a variety of 'object'-identification problems. Some of these are relatively simple (for example, the identification of CO_2 or oligomolecular mixtures), and in these cases, tight recognition by individual receptors might enable the animal to identify the molecule, measure its local abundance and track it through the use of olfactory LABELLED LINES[1,2]. However, many olfactory problems are more complex: they involve odours that are composed of multimolecular mixtures (sometimes containing hundreds of volatile components[3]). Odour perception tends to bind together rather than to segment the elements of a mixture[4–6]; the olfactory system therefore recognizes odours as patterns. Added complexity arises because the precise composition of an odour often varies during the lifetime of the odour source; fluctuations can be due to noise or to processes such as oxidation or differential volatility of the analytes in a mixture. In addition, the biological chemistry of odour formation (of flower scents, for example) leads to the formation of mixtures of chemically related elements (such as citrus essences). Odour clusters can be defined qualitatively (with many degrees of resolution: aromatic → minty → spearmint) or quantitatively[7,8], noting that concentration changes can also lead to changes in perceived quality. Human psychophysics reveals that such clusters can be identified[4]. I will therefore assume that, through evolution, the olfactory system has found solutions to these pattern-recognition tasks, in which the space of possible signals (perceptually definable odours) is immense and is not smoothly occupied. The magic of olfactory perception is that the brain can achieve cluster separation at (seemingly) many levels of resolution, allowing both gross classification and precise identification. How does it do it?

The olfactory bulb (OB) and its insect analogue, the antennal lobe (AL), are highly interconnected circuits in which inhibition is physically widespread, owing to the projections of either principal or local neurons (FIG. 1). Electrophysiological experiments have revealed many forms of temporal patterning of activity of their output elements, the mitral cells (MCs, in the OB) or projection neurons (PNs, in the AL)[8–18] (FIG. 2). The function of this patterning — seen equally in species in which MCs and PNs are multiglomerular (lower vertebrates and some insects)[8,13,19] and in species in which output neurons are mainly uniglomerular (mammals and

California Institute of Technology, Division of Biology 139-74, 1201 East California Boulevard, Pasadena, California 91125, USA.
e-mail: laurentg@caltech.edu
doi:10.1038/nrn964

other insects)[9–16,20] — is unclear. By exploiting the small sizes of insect (locust, honeybee) and zebrafish brains, we have tried to address this issue from a systems rather than a purely cellular or anatomical perspective; this approach reveals interesting computations that might otherwise remain undetected.

Working hypothesis

I propose that, because of the complexity of the olfactory pattern-recognition problem (the size and landscape of odour space, and the noisiness of odours), the brain exploits circuit dynamics to accomplish at least two objectives. The first is to create, through

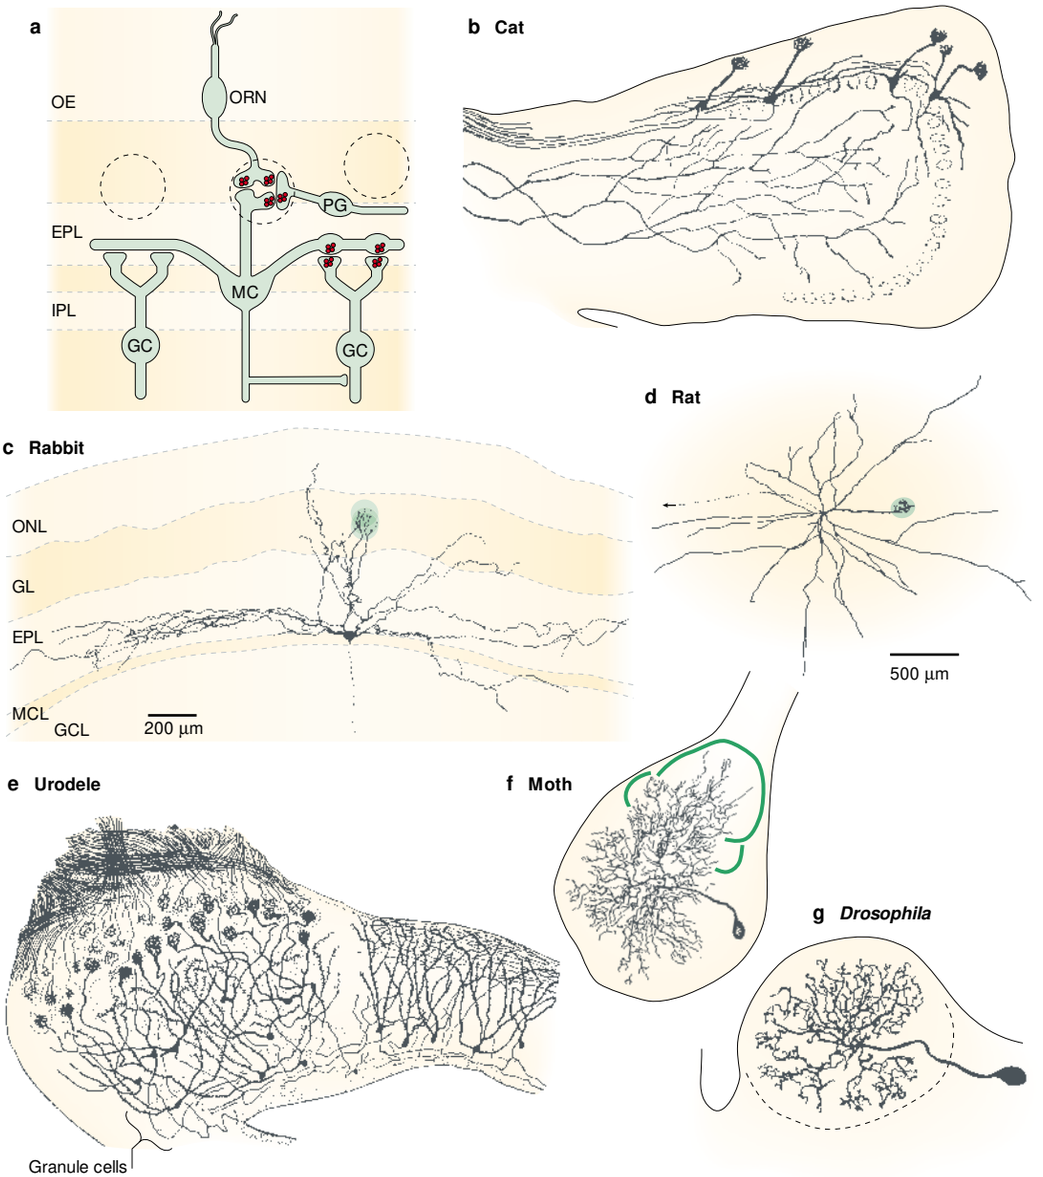

LABELLED LINES
A term that is used to describe a simple connectivity, whereby a set of identically and sharply tuned receptor neurons converge uniquely onto a set of postsynaptic neurons, which in turn project uniquely onto a set of common targets (and so on). Each channel (labelled line) can unambiguously inform the brain about the presence or absence of the signal it conveys.

Figure 1 | **Cellular elements underlying lateral inhibitory connections in the vertebrate OB and insect AL. a** | Schematic diagram of lateral interactions between mitral cells (MCs) and granule cells (GCs) of the olfactory bulb (OB). EPL, external plexiform layer; IPL, internal plexiform layer; OE, olfactory epithelium; ORN, olfactory receptor neuron; PG, periglomerular cell. **b** | Golgi stains of MCs in the OB of a young cat[57]. Note the lateral extent of the secondary dendrites, parallel to the OB surface. **c** | Rabbit MC and its extensive secondary dendrites. Note the lateral projections to more than 1 mm away from the soma. The glomerulus is indicated by a green shadow. GCL, GC layer; GL, glomerular layer; MCL, MC layer; ONL, olfactory nerve layer. **d** | Reconstruction of a rat MC in a plane tangential to the MC layer. Note the extensive dendritic disk and its size compared with the glomerulus (green shadow). **e** | Golgi stain of an urodele OB. Note the multiglomerular MCs and the extensive GC arborizations[57]. **f**,**g** | Stains (intracellular, **f**; Golgi, **g**) of local neurons (the functional analogues of GCs) in the antennal lobe (AL) of two insect species in which projection neurons (the functional equivalents of MCs) are mainly uniglomerular. Part **a** reproduced, with permission, from REF. 58 © 1993 Elsevier Science; part **c** reproduced, with permission, from REF. 59 © 1983 John Wiley & Sons; part **d** reproduced, with permission, from REF. 60 © 1984 John Wiley & Sons; part **f** reproduced, with permission, from REF. 61 © 1994 Springer Verlag; part **g** reproduced, with permission, from REF. 62 © 1994 Springer Verlag.

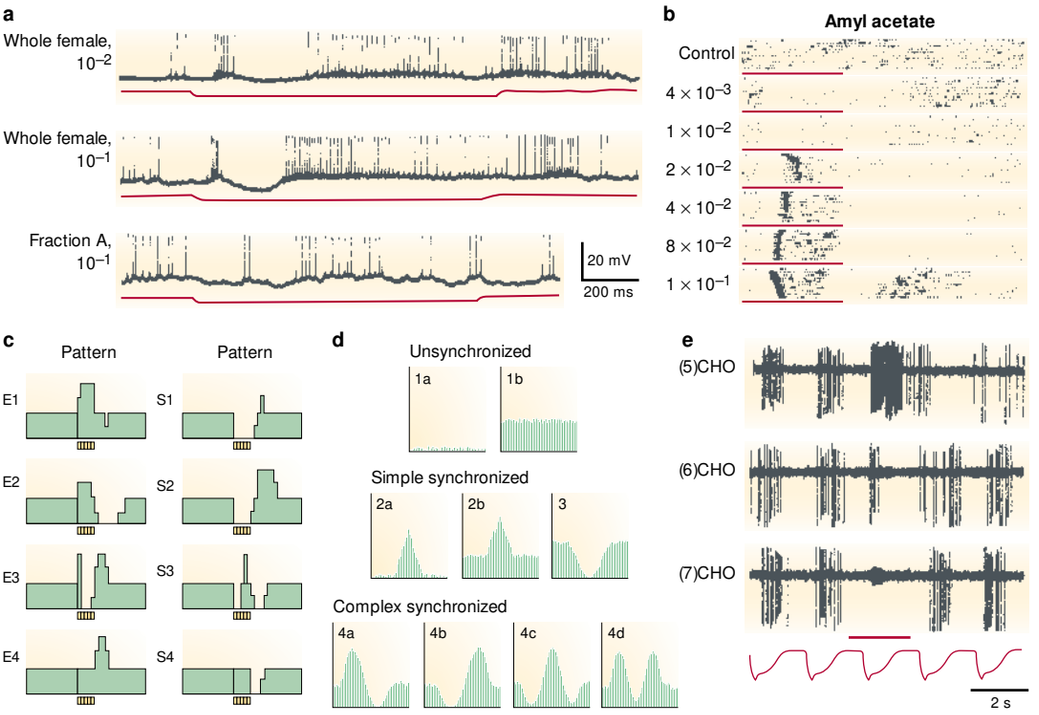

Figure 2 | **Slow temporal patterning of odour responses.** Examples of slow temporal patterning in the responses to odours of uniglomerular mitral/tufted cells (MCs) of the mammalian olfactory bulb (OB) or projection neurons of the insect antennal lobe, indicating that slow response patterning is not limited to systems in which principal neurons are uniglomerular. **a** | Response of a cockroach projection neuron to three pheromonal stimuli. The red line indicates stimulus onset and offset. **b** | Responses of a rat MC to background air (top) and to different concentrations (in M) of amyl acetate. The red line indicates the period of stimulation (500 ms). **c** | Schematic summary of temporal odour-response patterns in hamster OB neurons, including mitral and tufted cells. E and S indicate excitatory and suppressive patterns, respectively. Yellow blocks represent the odour stimulus. **d** | Schematic summary of rat MC response patterns relative to the respiratory cycle (window length corresponds to one cycle). **e** | Responses of a rabbit MC to three different aldehydes (5, 6 and 7 carbons). The bottom trace indicates an artificially imposed respiratory cycle; the bar marks the period of odour presentation. Part **a** modified, with permission, from REF. 10 © 1982 John Wiley & Sons; part **b** modified, with permission, from REF. 9 © 1989 The American Physiological Society; part **c** modified, with permission, from REF. 12 © 1992 The American Physiological Society; part **d** modified, with permission, from REF. 14 © 1972 American Association for the Advancement of Science; part **e** modified, with permission, from REF. 15 © 1995 National Academy of Sciences, USA.

spatiotemporal patterns of neuronal activation, a large CODING SPACE in which to spread representation clusters (FIG. 3a). The large size of this representation space is a consequence of the number of possible spatiotemporal combinations. The goal, however, is not to allow the storage of an infinite number of items; rather, it is to ease the handling of a smaller number of — often unpredictable — items that the animal will, in its lifetime, need to store and recall. The second objective is to use distributed dynamics both to confer stability on each representation in the face of noise and to optimize the filling of the representation space. In this article, I will present evidence that the first olfactory relay can, in two parallel operations, increase the separation between the representations of chemically related odours (decorrelation through slow dynamics) and format those representations so that they can be sparsened in the next station (exploiting oscillatory synchronization) (FIG. 3b). I propose, therefore, that the OB and AL are 'encoding machines' that actively transform a distributed, multidimensional afferent input to allow the formation of compact and easily recalled memories.

CODING SPACE
An abstract space that is defined by the features used to embody the code. If a neural system contains *n* neurons, one coding space can be viewed as an *n*-dimensional space, where each dimension represents the state of each neuron.

LOCAL FIELD POTENTIAL
The extracellular potential between two points in a brain region, resulting from synaptic and other current flow at and around the recording electrodes. It usually reflects input better than output.

The OB and AL as decorrelators

Slow patterns. Recordings from zebrafish MCs[8] and insect PNs[13] indicate — as shown in other species[10–14,21] (FIG. 2) — that the responses of principal neurons are not static. Rather, individual neurons respond with characteristic epochs of increased and decreased firing that are both neuron- and odour-specific. FIGURE 4 shows an example of a locust PN and its response patterns to 16 different airborne odours[22]. Similar patterning was seen across the responses of zebrafish MCs to many amino acids[8]. Because not all responding neurons express the same patterns at the same time, the population representation is dynamic, carried by an assembly of neurons (MCs or PNs) that evolves in a stimulus-specific manner over time (FIG. 3b). In locusts, this evolution can be tracked along a periodic LOCAL FIELD POTENTIAL (LFP; 20–30 Hz), which is caused by the synchronized periodic firing and updating of the participating neurons[23]. LFP oscillations in the same frequency range are also seen in fish, although their development during a response generally lags behind peak MC activity[8].

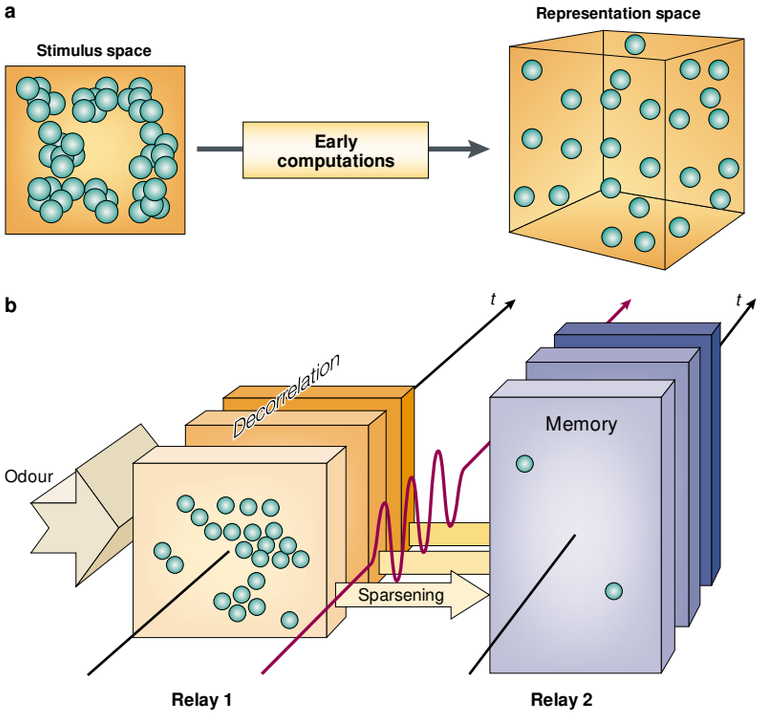

Figure 3 | **Schematic representation of the possible functions of olfactory circuit dynamics and their organization in time and space. a** | The early computations that are carried out in the olfactory bulb (OB)/antennal lobe (AL) and their immediate targets could result in both an expansion of the size of the coding space for odours (using spatiotemporal patterning) and a better use of that coding space for the distribution of odour representations. Each sphere in the stimulus space represents a combination of chemicals; each sphere in the representation space embodies one (or a family of) spatiotemporal pattern(s). **b** | As an odour is processed by the first relay (OB or AL), its representation by afferent neurons (pattern of glomerular activation) is given a spatiotemporal format because of dynamics that result from internal connectivity within that circuit. This patterning results in a decorrelation of representations (overlap reduction) over time. At the same time (at least in the locust), the spatial patterns of projection neuron activation at each oscillation cycle are compressed into patterns of few active neurons in a large population (relay 2; here, the mushroom body). This transformation results in an increase in the specificity of individual neurons' responses, and in a sparsening of representations[8,22]. The diagrammatic slabs along the time axis represent short time epochs, approximately equivalent to one half of a local field potential oscillation cycle. Each such epoch represents approximately the integration time of neurons in relay 2. So, neurons in relay 2 take short 'snapshots' of the state of relay 1, at times determined by the periodic output of relay 1.

Decorrelation. Successive time epochs in a sustained response involve different but stimulus-specific assemblies of active projection cells. What are the functional consequences of this dynamic and distributed activity? In zebrafish, spatiotemporal patterning results in a rapid decorrelation of odour representations[8]; however, for this to be uncovered, representations must be considered across MC assemblies. Decorrelation means that the overlap between the representations of related odours (for example, several aromatic amino acids) decreases with time, corresponding to divergent redistributions of activity over time across the MC array[8]. In this study, the representation size remained constant, on average. The trajectory followed by each stimulus-evoked evolution was reliable from trial to trial, and importantly, short response segments late in a response were more reliable than early ones for stimulus identification[8]. Early epochs offered reliable clues for odour classification, whereas later ones allowed precise stimulus identification by an observer.

Mechanisms and possible formal principles. The mechanisms that underlie this slow population patterning must involve interactions within the OB (or the AL), because afferent output shows no odour-specific or olfactory receptor neuron (ORN)-specific patterning, and no decorrelation over time[8,24]. In the locust, these mechanisms seem to be independent of fast inhibition in the AL[25], and do not involve feedback from downstream areas. Slow patterning in the locust AL therefore results from both slow inhibition (the mediation of which is still not fully understood) and, possibly, lateral excitatory interactions within the AL. A computational model of the AL and its connections was used to explore the minimum cellular, synaptic and network requirements for generating realistic population dynamics[26,27]. This revealed that distributed dynamics similar to those observed experimentally arise naturally in networks with realistic slow synapses and distributed lateral connections. A more abstract approach with smaller networks was used to explore fundamental aspects of these dynamic phenomena[28]. This work proposes, within the framework of nonlinear dynamical systems theory, a 'weak chaotic' regime called 'winnerless competition'[28] (WLC), in which the activity of the responding population follows an ORBIT that links unstable states. Orbits are highly sensitive to input, explaining the amplification over time of small input differences. Owing to the dissipative properties of motion, these orbits are stable, such that the population trajectory can be resistant to noise in the participating neurons, possibly explaining the trial-to-trial reliability of population patterns despite probabilistic responses in each neuron and epoch[28]. Qualitatively, each odour is represented by a constantly changing assembly in which each active neuron both participates in the dynamics of the others and benefits from the global stability of the assembly, preventing large individual response deviations. The possible link between WLC dynamics and experimental observations needs to be strengthened. However, this approach provides a simplified framework for exploring olfactory responses, their causes (such as asymmetrical inhibitory coupling) and their computational consequences (including representation optimization and stability). It is hoped that this combination of experiment and theory, focused on small model systems, will help to reveal principles of broad relevance.

Fast oscillations and sparse representations

Oscillatory synchronization in the olfactory system was first described using electroencephalographic (EEG) and LFP recordings in mammals[17]. It has since been found in most other systems (visual, auditory, somatosensory and motor)[29,30], including other olfactory systems[18,29,31,32]. We are attempting to provide a high-resolution description of the cellular, synaptic and circuit events that underlie these oscillatory LFPs[13,23,25–27]

ORBIT
The trajectory that is defined by a dynamical system, or its motion within state space. When applied to a system of neurons, an orbit is an abstract description of the states of all the neurons and the evolution of those states as a function of time.

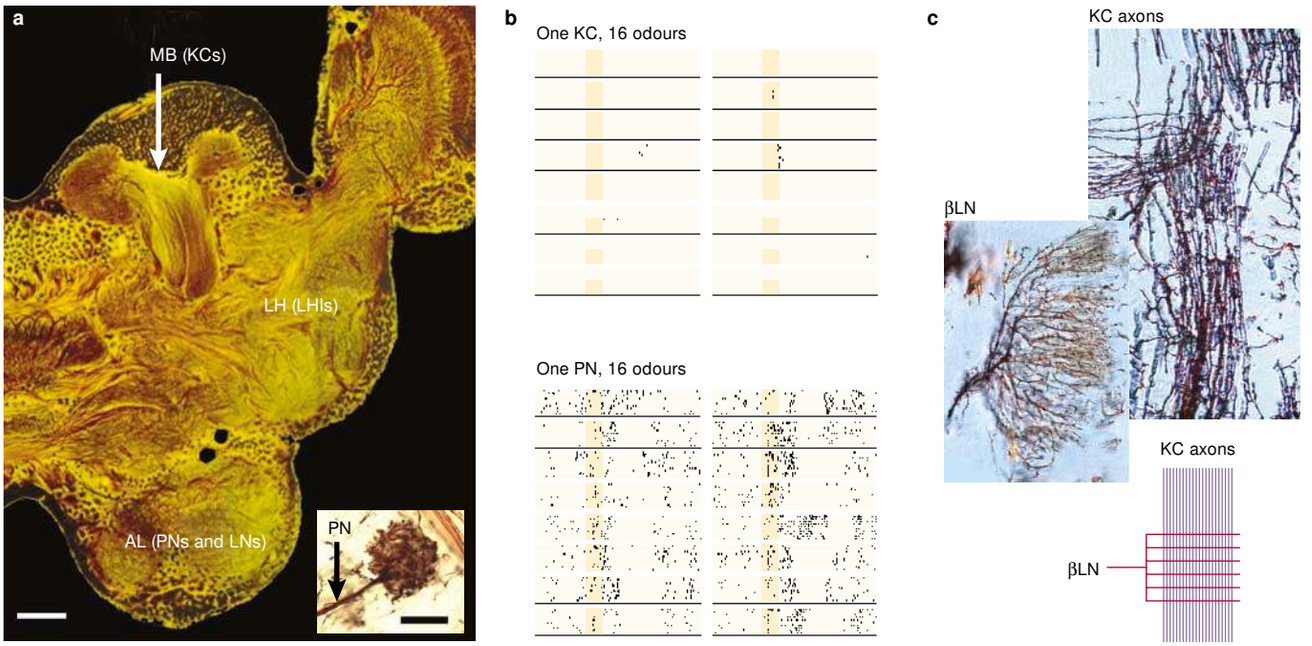

Figure 4 | **The locust olfactory circuits and the transformation of response properties between the first and second relay. a** | BODIAN STAIN of a locust brain (transverse section) showing the antennal lobe (AL), mushroom body (MB) and lateral horn (LH). The inset shows the terminal dendrite of one projection neuron (PN) in one glomerulus; in locusts, each PN sends dendrites to ~15 of ~1,000 glomeruli. KC, Kenyon cell; LHI, lateral horn interneuron; LN, local inhibitory neuron. Scale bar, 80 μm (inset, 20 μm). **b** | The bottom panel shows a tetrode recording from one PN in the AL and its responses to 16 different odours (ten trials each; odour pulses of 1 s are indicated by yellow shaded areas). Note the high baseline rates, high probability of response and odour-specific temporal patterning. The top panel shows a tetrode recording from one KC in the MB and its responses to the same 16 odours. Note the very low baseline rate, high specificity and brevity of response[22]. **c** | Golgi stains of KC axon tracks in the MB beta lobe and of beta lobe neuron (βLN) dendrites, sampling across those axons (see diagram in inset). βLNs are few (possibly several tens to hundreds) relative to the number of their KC inputs (50,000). The anatomy of the KC–βLN circuit is reminiscent of the microcircuit between beams of many parallel fibres onto the dendrites of Purkinje cells in the cerebellum[53].

in insects, and to test their functional relevance[21,33]. This system offers the prospect of understanding encoding, decoding and functional/behavioural aspects of periodic and synchronized activity in a brain area.

Causes and behavioural relevance. Oscillatory synchronization in the locust AL arises through the action of local inhibitory GABA (γ-aminobutyric acid) neurons (LNs) with widespread output to other LNs and to the PNs (FIG. 1f,g). Synchronization can be blocked by the local infusion of Cl^- channel blockers into the AL[25]. However, blocking fast inhibition leaves untouched the slow inhibition that is important for generating slow response patterning[25]. Consequently, the global, patterned PN population output can be maintained while disrupting periodic synchronization[25]. This dichotomy of inhibitory actions allowed us to establish the relevance of oscillatory synchronization, using both behavioural (in honeybees)[21] and physiological (in locusts) assays[25,33].

Hidden activity. If oscillations are functionally relevant, how and why are they useful? Let us first examine what an oscillatory LFP indicates. An LFP is a weighted average of local potential fluctuations that are caused by events (in our case, mainly synaptic currents) in the vicinity of the sampling site. Oscillations in the LFP indicate that groups of neurons with outputs close to the electrode tend to fire in common periodic epochs. The LFP does not, however, reveal any details of the activity that causes it. For example, in the simplest case, it could be caused by a subgroup of neurons that all fire together, periodically and precisely in phase throughout the epoch of oscillation. Alternatively, it could be caused by a neuronal group, the members of which change as the response progresses, but which, when they fire, do so in the proper phase. It could also be caused by a dynamic group of cells in which some, but not all, phase-lock to each other. Because the LFP is a mean, the influence of the phase-locked neurons (even if they are few) on the LFP waveform can be greater than that of the neurons that fire independently. Finally, LFP oscillations could also result from a dynamic assembly in which the active neurons can produce both locked and non-locked spikes, at different times in the response. Although seemingly baroque, this is what occurs in the locust AL[13,23,34]. Because PN output is temporally patterned (see above), not all PNs fire in the same epochs; when individual PN firing events, collected over many trials, are compared with the LFP, the spikes produced by individual PNs in some epochs of a response tend to be locked, whereas spikes produced earlier or later by the same PNs are not[13]. These epochs of locking are

BODIAN STAIN
A reduced-silver impregnation technique that is used for neuroanatomical studies of fixed brain tissue.

different for different PNs and for different stimuli. Possible reasons for this conditional locking are indicated by modelling experiments[26,27]: the strength of locking of any PN spike is correlated with the number of presynaptic LNs that are active in the short period before that spike. Consequently, the epochs during which the spikes of a PN are locked to the LFP can be determined continuously by the instantaneous state of the network and of the LNs presynaptic to that PN.

These observations matter for several reasons. First, they indicate that detecting pairwise correlations between any two neurons can be difficult: it is easy to miss those few oscillation cycles in which the two examined PNs fire synchronously[13,23], and it is similarly easy to erase the existence of a transient correlation by improper data analysis (for example, by measuring inter-cell correlations over time windows that exceed the average duration of pairwise correlation). Realizing the transient nature of pairwise synchronization is crucial to understanding the decoding of these signals (see below). Second, they indicate that LFP oscillations arise from a large number of spikes, of which only a fraction is locked at any one cycle. This is also important because, as we will see, the decoding circuits will not react to all spikes equally. The overall PN output during an odour response is therefore a complex distributed pattern in which PN spikes can be found at any time, but with a bias towards some periodic epochs, imposed by collective LN activity. The cells that are active together at different cycles change throughout a response and the spatiotemporal patterns differ for different odours.

Decoding. How is all this decoded? Recent intracellular and TETRODE recordings from Kenyon cells (KCs) — the intrinsic neurons of the mushroom body (MB) — indicate a marked transformation of representations between the AL and the MB[22] (FIGS 3 and 4). PN responses are long lasting, patterned, transiently locked, highly probable when tested over a set of ~20 randomly selected odours, and superimposed on a baseline firing rate of ~4 spikes s^{-1}. By contrast, most KC responses are extremely brief (~2 spikes), consequently unpatterned, locked to the LFP, highly improbable over the same odour sets, and superimposed on a baseline firing rate of 0.005–0.025 spikes s^{-1} (REF. 22). FIGURE 4 shows the responses of a typical PN and KC to the same set of 16 odours. The information content of a KC spike is clearly much higher than for a typical PN. Because KC responses are rare and because the MB contains many more KCs than there are PNs (50,000 compared with 830, respectively), odour representations in the MB are sparse[22]. The MB therefore seems to be sparsening (in space and in time) odour representations.

How is this accomplished? Some of the basic mechanisms are summarized in FIG. 5a,b. KCs and the circuits that surround them act as coincidence detectors on the dynamic PN input[22]. This results from several cooperating sets of features. First, the olfactory input to KCs is a complex pattern of PN firings distributed in space and in time. Second, anatomy indicates that individual PNs diverge, on average, to ~600 KCs. Third, given this fan-out ratio, the number of PNs (830) and the number of KCs connected to PNs (25,000–50,000), each KC must receive convergent input from 10–20 PNs, on average. Fourth, KC responses to PN spikes can be amplified by voltage-dependent nonlinearities that also shorten excitatory postsynaptic potentials (EPSPs) when the input causing them is strong enough[22,32]. So, KCs will summate EPSPs preferentially if their timing is synchronized. Fifth, a short feedforward circuit through inhibitory lateral horn interneurons (LHIs) produces inhibitory postsynaptic potentials onto KCs that are out of phase (FIG. 5b) with the EPSPs caused by synchronized PNs[22]. Because individual LHIs respond to most odours, because LHIs are few (~60) and because they diverge extensively in the MB, they can collectively inhibit the KCs during half of every oscillation cycle caused by any odour. This ensures that KCs can summate PN input only briefly and periodically during the other half of each cycle; that is, before the LHIs fire (FIG. 5b).

How, then, do KCs respond at all? This is explained by the limited convergence of PNs onto any KC, by the transient nature of the PN output during odour stimulation, and by a presumably high KC firing threshold. Only when a sufficiently high proportion of the PNs presynaptic to a KC fire synchronously does that KC fire an action potential. The brevity of the KC response could be explained by two observations: first, a large and long-lasting AFTERHYPERPOLARIZATION follows each KC spike[32], making EPSP summation less effective and further firing unlikely; and second, the evolving nature of the PN output ensures that, within a few cycles, the set of co-active PNs has changed. So, sparsening results from an asymmetrical influence of periodic excitation and inhibition on each KC: excitation is highly specific, whereas inhibition is not (FIG. 5a,b). These results show that neurons can act as coincidence detectors[35,36], and they reveal how oscillatory synchronization underlies an important computation.

Significance. Oscillatory synchronization and a set of appropriately tuned ancillary mechanisms can, in one step, convert a dense, distributed and redundant stimulus representation into a sparse one. But are oscillations necessary? I would argue that shaping synthetic and specific responses might not be easy, especially if it must be achieved in only one step. Synthetic tuning implies the convergence of many inputs onto one neuron. If that neuron must respond only to the co-activation of all (or most) of its converging inputs (a logical 'AND'), it must be able to ignore a pattern in which only a subset of these inputs is vigorously active, but respond when all inputs are equally active. In other words, it must be able to select against temporal summation and for spatial summation of (coincident) input; input synchronization and active shortening of the integration window, as found here, is a solution to this problem. It will be interesting to determine whether solutions that do not use synchrony are equally efficient.

TETRODE
An extracellular electrode that comprises four juxtaposed recording channels, which can be used to disambiguate the signals emitted by individual point sources. Because each neuron occupies a unique position in space, its spikes are 'seen' slightly differently by each electrode, providing a unique signature. This technique allows the identification of many more neurons than there are sampling electrodes.

AFTERHYPERPOLARIZATION
The membrane hyperpolarization that follows the occurrence of an action potential.

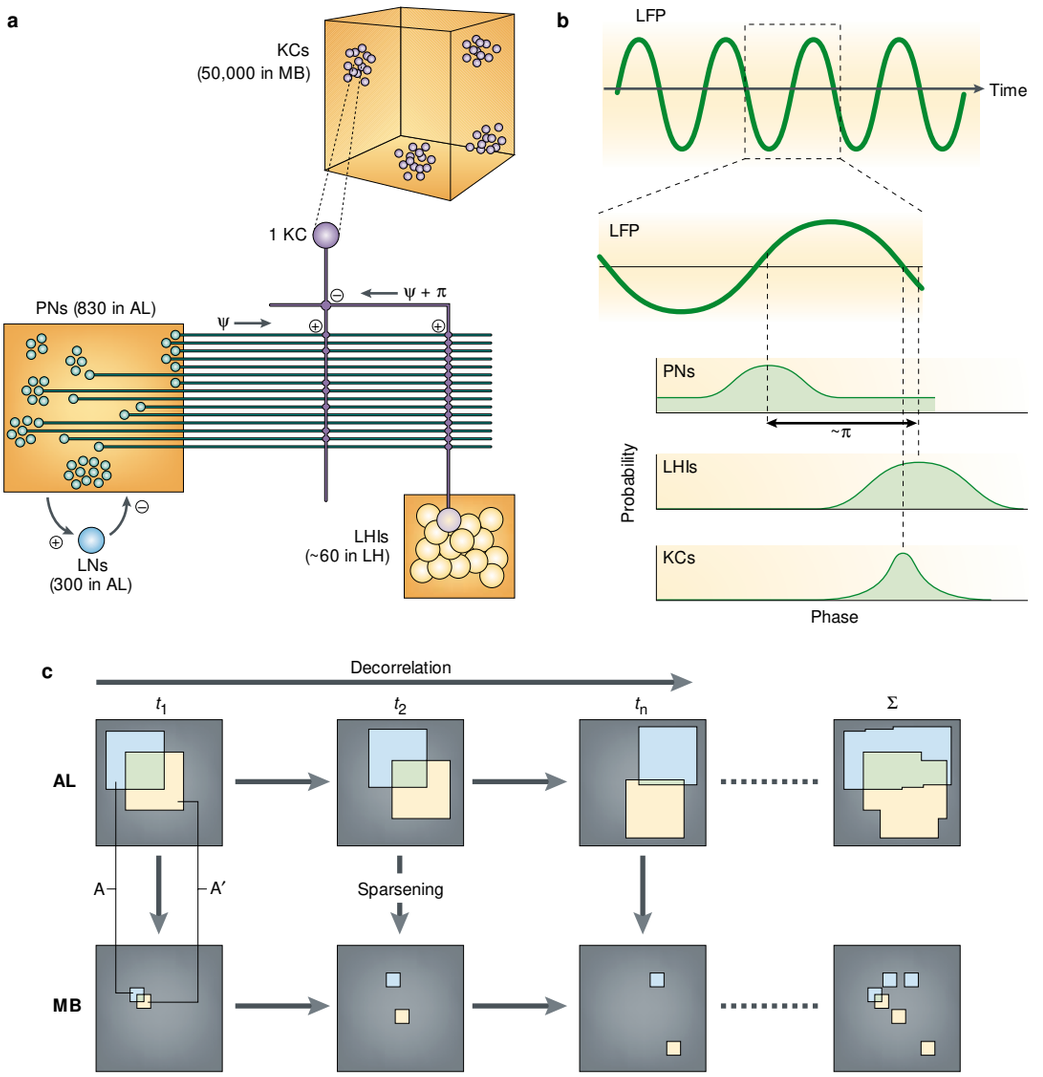

Figure 5 | **Mechanisms and possible consequences of sparsening of sensory representations by oscillatory patterning and coincidence detection. a** | The main cellular elements in the early olfactory system of locusts. The antennal lobe (AL) contains about 830 excitatory projection neurons (PNs) and 300 local inhibitory neurons (LNs). Interactions between PNs and LNs within the AL generate the patterned PN output. The 830 PN axons run through the mushroom body (MB) and terminate in the lateral horn (LH). In the MB, each PN contacts ~600 Kenyon cells (KCs), which are distributed throughout the MB. Each KC receives inputs from an estimated average of just 10–20 PNs. A single KC therefore samples the states of a small subset of PNs. The same PN axons terminate in the LH, which contains about 60 GABA (γ-aminobutyric acid) interneurons (LHIs)[22]. LHIs send branched and divergent projections to the MB, contacting KC dendrites. Because LHI responses to odours show little stimulus specificity, each KC receives reliable inhibition at each oscillation cycle, out of phase with the excitation from a subset of PNs (**b**). During each oscillation cycle, an ensemble of PNs is activated, among which some are tightly locked. Most KCs will receive weak excitation caused by the few active PNs that each KC is connected to; these KCs will remain silent. For a subset of KCs, however, a significant proportion of the PNs presynaptic to them will be activated. The action potentials of these KCs encode the co-activation of these PN sets. At least two mechanisms ensure that these KCs produce an action potential only if their presynaptic PNs fire at the right time. KC intrinsic properties amplify and shorten excitatory postsynaptic potentials (EPSPs) when the input is sufficiently strong; this automatically reduces the KC effective temporal integration window at each oscillation cycle. Second, the feedforward circuit through the LH inhibits all the KCs during each half of each oscillation cycle (**b**). This ensures that most KCs are actively kept silent during an odour presentation and that PN-evoked EPSPs can summate only during the short period of each cycle that corresponds to synchronized PNs firing. Spurious and ill-timed PN spikes must compete with an inhibitory postsynaptic potential and are therefore less effective. **b** | Schematic of the firing time probability of PN, LHI and KC populations relative to the local field potential. **c** | Schematic indicating how the combination of decorrelation and sparsening in early olfactory circuits could allow a simple decoding of odour identity by slow temporal integration (Σ) by neurons such as the MB beta lobe neurons (FIG. 4c). Odours A (blue) and A′ (yellow) evoke overlapping PN response patterns in the AL. The overlap between them decreases with time, but because of the dense, distributed representation mode in the AL, the sums of the A and A′ patterns over time overlap. If, by contrast, the AL patterns are sparsened in the MB, the sum of the A and A′ patterns in the MB can have very little overlap, implying that simple integrators (one for A and one for A′) downstream of the MB could easily differentiate them. A temporal pattern could therefore be decoded without recourse to sequence decoding.

There are practical considerations. The detection of sparse representations, when they exist, can be difficult, simply because spikes might be extremely rare. If there is no independent reason to suspect the existence of such responses (such as intracellular recordings that indicate stimulus-related subthreshold activity), it is easy to miss these few, highly informative action potentials. However, KC action potentials are highly significant only because they ride on a very low baseline firing rate. So, mechanisms must be invested to secure the contrast between response and no-response. This mode of representation therefore has an associated cost that it would be interesting to estimate.

Sparsening has many advantages, especially if it occurs in a structure that is implicated in learning (such as the MB). As well as reducing overlaps, sparse representations could facilitate storage (fewer synapses need to be modified), pattern matching (fewer elements need to be compared) and pattern association: the different attributes of a percept (for example, shape, colour, texture, odour, identity and category) should, in principle, be more easily associated if they require the linking of fewer neuronal elements. Conversely, by combining many converging inputs, specific neurons (KCs in our case) could contribute to the formation of the complex associations that underlie perceptual binding[29,30,37]. Sparse, synthetic representations are useful, but they eliminate the detail and segmentability of a representation (Gestalt). This is consistent with behavioural and psychophysical observations[5,6,38] in olfactory perception. Another possible advantage of the phenomena that we describe is that they are adaptive. The feedforward inhibitory loop that sharpens KC tuning could also be viewed as a compensatory mechanism for the sloppiness of the oscillatory clock: the LFP oscillation frequency usually varies between 15 and 30 Hz from cycle to cycle[18]. Because feedforward inhibition is locked, cycle by cycle, to each ongoing wave of excitation, delays or advances in the PN output are always compensated for adaptively by LHIs. The relative lack of formatting precision in the PN output can therefore be corrected automatically.

Finally, our results imply that not all spikes are alike: whether a PN spike succeeds in activating its targets will be determined by the timing of that spike relative to the timings of other spikes that are produced by other neurons at around the same time. The relevant information content of a PN spike is therefore determined by its temporal correlation with the spikes of other PNs that share the same targets; it cannot be measured meaningfully without the knowledge of these spatiotemporal relationships. The existence of oscillatory synchronization can therefore indicate a selective filtering of throughput.

Are slow dynamic patterns features of a code?

So far, it seems that decoding of the AL output by KCs makes no explicit use of the dynamic features of PN responses. KCs do not seem to accomplish any kind of sequence decoding across the incoming PN input; rather, each KC selectively assesses the state of a small part of the PN assembly, one fraction of an oscillation cycle at a time, with no apparent memory of activity in previous cycles. For this reason, slow dynamics might seem to be irrelevant. But we believe them to be essential. First, sequence decoding could be accomplished downstream of KCs — for example, by extrinsic neurons in the alpha and beta lobes of the MB[33,39,40] (FIG. 4c) — using unknown spatiotemporal integration mechanisms. Alternatively, the relevance of slow dynamics in the AL might be implicit in the KC responses. We have proposed that the AL/OB output is a self-organized process, the outcome of which becomes less ambiguous with time[8,34]. I would argue that the evolution of AL/OB patterns might need no decoding *per se*. According to this perspective, dynamics are crucial for the optimization of the code, but need not be the code itself (that is, a feature to be decoded); the complicated patterning we observe in AL/OB neurons might simply be part of the process through which the format of the message is actively optimized for further processing (learning, association, recall) by downstream areas. The more time there is available for optimization (that is, the longer the stimulus), the easier the discrimination. The decoding could therefore occur piecewise over time, becoming increasingly refined as it is updated with each oscillation cycle, or alternately, it could be achieved by simple temporal integration, as explained below.

Temporal patterning without sequence decoding

The existence of a sparsening stage in the representation considerably simplifies the read-out of spatiotemporal patterns. Consider odour A, which is represented in the periphery by a physical array of activated glomeruli. This representation overlaps with that of A′, a related odour. By imposing a temporal structure on this representation, the OB/AL unfolds the spatial patterns of A and A′, and reduces overlaps (FIG. 5c). In the locust, using the periodic output of the AL, the MB sparsens these representations at each oscillation cycle across a large assembly of KCs. Both transformations decrease the probability of overlap between representations. This implies that slow postsynaptic integration of KC outputs (sensitive to the identity but not to the order of activation of the responding neurons) might suffice to separate A from A′, even if these assemblies contain a few common elements, as a consequence of their relatedness or common root. (This would not be possible with PNs in the AL, especially if many patterns needed to be stored, because PN representations are dense.) Neurons in the alpha and beta lobes, the dendrites of which sample the axons of hundreds to thousands of KCs[33,39,40] (FIG. 4c), could carry out this simple integration. In conclusion, one could imagine that slow temporal patterns, although crucial for the separation of representations, are never actually decoded as such. More generally, the creation of spatiotemporal representations by circuit dynamics might be a transient phase in signal processing, used simply to spread out those representations in a larger coding space and to facilitate decoding (for example, sparsening followed by conventional spatiotemporal integration). Note that I do not exclude the possibility of sequence decoding

Box 1 | Sparse memories, the cerebellum and the mushroom body

Marr[41] and Kanerva[42] proposed related theoretical approaches to the problem of memory storage in neural circuits. Their proposed data structures (for the cerebellum) are reminiscent of those described here. Marr proposed a scheme in which three features of connectivity — divergence of mossy fibres onto granule cells (GCs), convergence of mossy fibres onto GCs and very high GC number — are elements of a sparsening design. Each mossy fibre input pattern to a GC is called a codon, and each GC fires only when all of its afferents are active. Simple calculations show that this naturally leads to overlap reduction between input patterns. This rule is reminiscent of the projection patterns and input transformations between the antennal lobe and the mushroom body (MB)[22].

Kanerva[42] imagines n binary neurons that define a space, $\{0,1\}^n$, where $n \gg 10^2$. If each memory (representation) is an n-bit word, 2^n memories can be encoded. A HAMMING DISTANCE d (0–n bits) is used to measure the separation between any two points (dissimilarity between two memories). Given an arbitrary point x, most of the space lies $d = n/2$ bits away from x (most points are uncorrelated with x). If n = 1,000 neurons, only 10^{-10} of the space will be within 400 bits of x. If we define memory item X as anything within 400 bits of x — if we aim to store a number of items that is small relative to $2^{1,000}$ — then recognition can be almost guaranteed if the state of each neuron can be determined with a success probability of 0.6 (600 correct bits out of 1,000). So, if memories are spread out in a large coding space, recognition could occur even with a very partial match between object and target patterns.

This idea is appealing if applied to projection neurons (PNs). At each odour-induced oscillation cycle, we can picture the PN assembly as defining a 830-bit word in which up to $n/4$ of the bits are '1's. At later cycles, PN patterns become decorrelated — the mean distance between words created by similar odours increases. Provided that the occupation of the coding space is optimized (decorrelation) and the decoding is carried out correctly, recognition could occur with a very partial match.

The decoding stage is crucial. The 2^{830} patterns that PNs could produce per oscillation cycle cannot each be assigned to an individual neuron. Kanerva's solution is a sparse and distributed memory, where 'sparse' means that only a fraction of possible words will be represented physically in a 'storage location': he considers that to be the GCs. For our purpose, imagine the storage location to be the Kenyon cells (KCs). KC numbers range from 2,500 (*Drosophila*) to 4×10^5 (cockroaches) per MB. The distribution of the n-bit words embodied by these storage locations should ideally be homogeneous in $\{0,1\}^n$ space. The memory should also be distributed, because the distribution of points in $\{0,1\}^n$ means that any two points have many common neighbours. Confusion can be avoided by representing an input across a combination of storage locations. In Kanerva's implementation, the memory item is copied across several GCs, and identification converges rapidly to the stored item, if it exists. Convergence also occurs when recalling input sequences, which is the format of the PN input to KCs. In Kanerva's implementation, the 'bit locations' (where memories are stored as synaptic weights) are the synapses between GCs and Purkinje cells (PCs). A single GC diverges onto many PCs, and a single PC receives many GC inputs (parallel fibres). The PC acts as the output line and sums GC input during recall. If we substitute KCs for GCs, KC axons for parallel fibres, and extrinsic MB neurons for PCs, we obtain a prototypical MB: KC axons make many *en passant* synapses in the pedunculus and lobes. The recipients of these contacts are 'extrinsic' neurons with stratified dendrites that comb through the KC axon arrays (FIG. 4c), just as PCs do across parallel fibres[33,39,40,53]. Marr and Kanerva propose a rationale for this architecture. However, theory and experiments do not fit perfectly. For example, Kanerva's scheme implies that learning occurs at the GC–PC synapse. Although odour memories probably reside in KCs[54], synaptic transmission by KCs might not be required during odour learning[55,56]. In addition, Kanerva's scheme requires that most GCs (KCs) receive $n/2$ inputs (n being the number of input lines: mossy fibres for Kanerva, PNs for us). This is true neither of the cerebellum nor of the MB: convergence is ~ 5:1 onto cerebellar GCs and ~ 10–20:1 onto KCs[22]. In olfactory circuits, we believe that this small convergence ratio is a constraint to sparsen odour-activated patterns. Perhaps this is also true of GC activation patterns.

mechanisms. I simply point to a realistic solution, possible here only because AL and MB processing results in advantageous representation formats.

Circuit dynamics as a mechanism for recall

Forming and storing representations (for odours as for any other feature) is clearly not an end in itself for the brain. Memories are formed because they might be needed later to help in decision-making and action. In other words, the format of what we call a 'response' probably depends as much on the process that forms a representation as it does on the process that will lead to recall of that representation at a later stage: when an animal explores and samples the world, it might not always choose between acquisition and recognition modes. These two operations must therefore be able to occur through the same machinery and the same process. Presumably, evolution exerted some selective pressure on brain mechanisms that serve both pattern formation/memorization and recall equally well. So, we should consider the possibility that dynamics might be useful not only for representation, but also for recognition. For example, reinforcing, through learning, the connections between the neurons that form sequences in a spatiotemporal pattern might facilitate reactivation of the sequence by a corrupted input, and thus recognition. This idea, sometimes called 'pointer chain', is implicit in Marr's paper[41] and is developed in Kanerva's book[42] (BOX 1). Circuit dynamics might also be crucial when the animal is actively looking for a particular feature: top-down influences might bias and facilitate the recognition of the searched item by more peripheral circuits[43]. In brief, circuit dynamics might have functions that we do not yet understand or even suspect, because

HAMMING DISTANCE
The number of bits by which two n-bit vectors differ. For example, the Hamming distance between 001101 and 001110 is 2. It is also the square of the Euclidian distance.

our functional framework, defined by experimental constraints, is often much narrower than that in which brains normally operate.

The problem with noise

Much of the processing described above exploits mechanisms that should, in principle, be very sensitive to noise. Input decorrelation, for example, requires an operation akin to the amplification of small input differences, but not of ones that arise from natural, noisy fluctuations of the stimulus. Also, pattern encoding by KCs relies on rare but highly informative action potentials that must not be polluted by spurious spikes. How are these problems solved? We do not know, but there are some hints to the possible solutions. The convergence of many (in some cases thousands of) ORNs onto single glomeruli (and therefore few output neurons)[44–46], and the distributed sprinkling of these ORNs on the receptive sheet[44,47], limiting the probability of correlated noise, could allow the averaging necessary to increase signal-to-noise ratios. In addition, slow and diffuse communication within individual glomeruli could contribute to averaging or adaptive gain control. This issue clearly needs careful attention. A second potential mechanism for noise reduction is a form of fast learning that is seen in the locust[48]. AL circuits seem to undergo stimulus-specific modifications (the underlying mechanisms of which remain unknown) to the extent that successive responses of PNs to the same stimulus rapidly decrease in intensity, but become more precise and coordinated with those of other PNs[48]. Because olfaction is generally intermittent, and because the detection of an odour at one time predicts the presence of the same odour in the very near future (odours rarely disappear suddenly), the AL circuits might operate at low detection threshold at 'rest' (explaining the high responses in a naive state), but immediately 'focus' on a signal once it has been detected (explaining the refined representation after just a few trials). If this form of learning exploits short-term changes in the synapses formed by the activated neurons, only those synapses that are repeatedly activated over successive samplings could be reinforced. In other words, unreliable contamination occurring on some but not all trials (noise) would be averaged out. Third, network mechanisms (especially the connectivity matrix of OB/AL circuits) could have a crucial role in ensuring stability in the collective output. Recall that the dynamic evolution of the OB/AL output is forced as long as the stimulus lasts. So, the input signal is not an initial condition (as it is in HOPFIELD NETWORKS), but rather, an ongoing signal that generates an ongoing dynamical pattern across the AL. During that time, the stimulus could define a state space in which the ATTRACTOR is unique. There might be particular rules of connectivity that confer stability on their dynamical evolution in response to noisy stimulation. This is an area in which theory, modelling, neuroanatomy and physiology will all be necessary. Fourth, the powerful slow inhibition that is mediated by LNs (in addition to fast, synchronizing inhibition)[25] creates a means by which to average out noise and also actively suppress the baseline noise that results from non-activated PNs. Modulatory neurons and their potential role for plasticity in ALs could be similarly important. Finally, the mechanisms that ensure that KCs fire only when they detect the right input combination probably require delicate fine-tuning of the balance between excitation and inhibition. Inhibitory feedback from the output of the MB to the dendrites of KCs is known to exist[40]. These pathways might contribute to the adaptive control of KC excitability and therefore to a more-or-less constant representation sparseness. In conclusion, the experimental results and proposed principles that are summarized above require mechanisms that can discriminate noisy from meaningful differences. There are a few promising candidate mechanisms that could ensure or promote resistance to noise, but much work is needed in this crucial area.

Decorrelation and perceptual clusters

Although the decorrelation of input representations is useful in principle, it could also have undesirable effects. For example, the perceptual relatedness of odours (for example, all citrus-like smells) might be lost. Similarly, individual odours at different concentrations usually retain, at least over some range, the same perceptual identity. But if the patterns evoked by different concentrations of the same odour differ even slightly from one another, decorrelation would enhance those differences and possibly preclude their perceptual grouping. How is this potential conflict resolved? Again, we do not know yet, but we can suggest several possible solutions, each of which needs to be explored. The first is that perceptual grouping could be a high-level, learned property. In this scheme, input patterns that end up in the same perceptual group (for example, several concentrations of jasmine) do not actually evoke related patterns after decorrelation. But because the animal experiences all these concentrations within a given epoch (perhaps while visiting a given cluster of flowers), and because all patterns experienced during this period are equally meaningful (they are all associated with a reward), all patterns evoked by the different concentrations are lumped, downstream of the circuits responsible for decorrelation, as 'meaning' the same thing (in this example, jasmine). This is a high-level grouping, by contingency.

A second possibility draws on the fact that decorrelation is a temporal process, so early phases of a representation (say, the first 100 ms) are similar across related stimuli[8]. Provided that the brain can hold this (fleeting) information, it could use it for perceptual grouping or, conversely, ignore it for precise identification using decorrelated patterns. This supposes the existence of several read-out streams — for example, one for early patterns and one for the entire pattern — and a top-down system to decide which stream to listen to. A third hypothesis is that network dynamics and sparsening never completely orthogonalize representations. In this scheme, the dynamics would be designed such that the spacing between representation clusters increases faster than the spacing between representations within a

HOPFIELD NETWORK
A type of trainable, asynchronous artificial neural network with symmetrical connections that defines sets of attractor states. Given a certain input set, a Hopfield network can therefore be made to settle into a given attractor, in a process akin to pattern completion.

ATTRACTOR
Given a dynamical system and the state space in which it lives (that is, all the possible states that this system can occupy), an attractor is a preferred region of state space; that is, a state or set of states to which the system moves inexorably as time approaches infinity.

Box 2 | **An active and systems perspective on sensory processing**

The contrast between odour representations in the antennal lobe and in the mushroom body is enormous. An observer looking at activity in these two areas might never suspect that they are separated by only one synapse. Similarly, observing individual Kenyon cell responses alone does not hint at integrative properties that depend on presynaptic spike timing and on the coordination of input arrival. In other words, understanding the computations that take place in a circuit can be difficult if we fail to consider individual neurons as parts of a system in action. I would argue that our lexicon introduces subtle but real biases in our thinking about sensory processing. Our predisposition as sensory physiologists is to call 'responses' the spike patterns that follow a stimulus; we then use these responses to define 'receptive fields'. In doing so, we forget that these terms are meant only to be operational. Our thinking about sensory integration seems to be much too linear and passive: stimulus *a* leads to a response in area *x*, which produces a response in area *y*, and so on. In reality, neural circuits are often massively interconnected and reciprocally connected. Similarly, our thinking generally ignores the fact that, with the exception of motor neurons, a given neuron is never an end-point or its 'response' an end-product. So, how a neuron behaves might be relevant not as a response *per se* (something to be analysed by us to estimate the information it contains about a stimulus, although this is, of course, useful knowledge), but as part of a transformation (possibly extremely complex and distributed) to help further processing (for example, optimization, storage, recognition and retrieval) in the area in which the neuron lies (for example, decorrelation in circuits of the olfactory bulb) or in 'target' circuits. Thinking about sensory integration in these active terms (considering 'responses' not only as products, but also as ongoing transformations towards some other goal) might be helpful as we try to understand some brain operations.

cluster. So, patterns that are similar (different concentrations or related chemicals) remain more similar to one another than to any random pattern. This type of biased decorrelation might require a particular circuit architecture.

Conclusion

Much integrative work is needed to understand the computational organization of olfactory systems. I propose a systems perspective that is based on experimentation with small olfactory brains; by exploiting their relative simplicity, we have shown that circuit dynamics over several timescales and correlation rules have an integral role in optimizing stimulus representations. Dynamic formatting might be a transient phase in the processing of these signals: once representations have been optimized, their apparent initial complexity could be reduced to simple responses that are carried by a few specific neurons. Circuit dynamics might have other advantages for processes such as memory recall, but experimental support for this is so far lacking. Advances in our knowledge of how natural stimuli are distributed within odour space and what multiple tasks olfactory systems must be capable of (for example, learning, recognition and classification) will help us to understand better why olfactory computations are the way they are; I would argue that a traditional, passive 'stimulus–response' view of sensory processing hinders our understanding of seemingly complicated modes of operation (BOX 2). Note that the computational framework that I propose need not be valid only when examining complex (multimolecular) odour processing. Indeed, what is crucial is simply that the input signal — be it mono- or multimolecular — be transduced and processed in parallel by interacting 'channels' (which might represent the several receptor types and postsynaptic neurons that are activated by a single molecule).

Finally, one should resist the temptation to propose an artificial and misleading dichotomy between spatial and temporal aspects of olfactory codes. As emphasized previously[34], I believe that spatial (that is, identity related) and temporal aspects are two sides of the same coin. What matters is that we understand the functions of each one of these facets, and the features of cellular and population activity that matter to the neurons doing the 'decoding'. It will be interesting to learn whether the principles and mechanisms that I propose here apply to other olfactory systems, including those of mammals, and possibly also to other brain systems that are involved in the processing of multidimensional signals, such as vision[49–51] or action[52].

1. Kauer, J. S. & White, J. Imaging and coding in the olfactory system. *Annu. Rev. Neurosci.* **24**, 963–979 (2001).
2. Hansson, B. S. (ed.) *Insect Olfaction* (Springer, Berlin, 1999).
3. Knudsen, J. T., Tollsten, L. & Bergstrom, L. G. Floral scents — a checklist of volatile compounds isolated by headspace techniques. *Phytochemistry* **33**, 253–280 (1993).
4. Laing, D. G. in *The Human Sense of Smell* (eds Laing, D. G., Doty, R. L. & Breipohl, W.) 241–259 (Springer, Berlin, 1991).
5. Chandra, S. & Smith, B. H. Analysis of synthetic processing of odour mixtures in the bee (*Apis mellifera*). *J. Exp. Biol.* **201**, 3113–3121 (1998).
6. Livermore, A. & Laing, D. G. Influence of training and experience on the perception of multicomponent odor mixtures. *J. Exp. Psychol.* **22**, 267–277 (1996).
7. Duchamp-Viret, P. & Duchamp, A. Odor processing in the frog olfactory system. *Prog. Neurobiol.* **53**, 561–602 (1997).
8. Friedrich, R. & Laurent, G. Dynamical optimization of odor representations in the olfactory bulb by slow temporal patterning of mitral cell activity. *Science* **291**, 889–894 (2001).
   A paper showing the decorrelation of odour representations by MC assemblies resulting from slow MC temporal patterning. The data indicate that activity across MCs is redistributed across the population so that overlap between the representations of similar odours decreases as a function of time within the first ~800 ms of a response.
9. Wellis, D. P., Scott, J. W. & Harrison, T. A. Discrimination among odorants by single neurons of the rat olfactory bulb. *J. Neurophysiol.* **61**, 1161–1177 (1989).
10. Burrows, M., Boeckh, J. & Esslen, J. Physiological and morphological properties of interneurons in the deutocerebrum of male cockroaches with responses to female pheromones. *J. Comp. Physiol. A* **145**, 447–457 (1982).
11. Meredith, M. Patterned response to odor in mammalian olfactory bulb: the influence of intensity. *J. Neurophysiol.* **56**, 572–597 (1986).
12. Buonviso, N., Chaput, M. A. & Berthommier, F. Temporal pattern analyses in pairs of neighboring mitral cells. *J. Neurophysiol.* **68**, 417–424 (1992).
13. Laurent, G., Wehr, M. & Davidowitz, H. Temporal representations of odors in an olfactory network. *J. Neurosci.* **16**, 3837–3847 (1996).
14. Macrides, F. & Chorover, S. L. Olfactory bulb units, activity correlated with inhalation cycles and odor quality. *Science* **185**, 84–87 (1972).
    References 14 and 19 are among the first papers to indicate that mammalian and non-mammalian vertebrate MCs show slow temporal patterning in response to odours.
15. Yokoi, M., Mori, K. & Nakanishi, S. Refinement of odor molecule tuning by dendrodendritic synaptic inhibition in the olfactory bulb. *Proc. Natl Acad. Sci. USA* **92**, 3371–3375 (1995).
16. Motokizawa, F. Odor representation and discrimination in mitral tufted cells of the rat olfactory bulb. *Exp. Brain Res.* **112**, 24–34 (1996).
17. Adrian, E. Olfactory reactions in the brain of the hedgehog. *J. Physiol. (Lond.)* **100**, 459–473 (1942).
    One of the first papers to indicate the existence of oscillatory dynamics in the mammalian olfactory system.
18. Laurent, G. & Davidowitz, H. Encoding of olfactory information with oscillating neural assemblies. *Science* **265**, 1872–1875 (1994).
19. Kauer, J. S. & Moulton, D. Responses of olfactory bulb neurones to odour stimulation of small nasal areas in the salamander. *J. Physiol. (Lond.)* **243**, 717–737 (1974).
20. Spors, H. & Grinvald, A. Spatio-temporal dynamics of odor representations in the mammalian olfactory bulb. *Neuron* **34**, 301–315 (2002).
21. Stopfer, M., Bhagavan, S., Smith, B. H. & Laurent, G. Impaired odour discrimination on desynchronization of odour-encoding neural assemblies. *Nature* **390**, 70–74 (1997).
    This paper indicates that oscillatory synchronization is required for fine odour discrimination. The authors used selective pharmacological blockade of oscillatory synchronization in the honeybee AL, and behavioural assessments to determine its consequences for odour discrimination.

22. Perez-Orive, J. *et al.* Oscillations and sparsening of odor representations in the mushroom body. *Science* **297**, 359–365 (2002).
The first description of responses of MB KCs to odours. This paper provides direct physiological evidence that oscillatory synchronization is a dynamic mechanism used by a brain circuit to bind separate elements in a sensory representation and so sparsen that representation.
23. Wehr, M. & Laurent, G. Odor encoding by temporal sequences of firing in oscillating neural assemblies. *Nature* **384**, 162–166 (1996).
24. Friedrich, R. W. & Korsching, S. I. Combinatorial and chemotopic odorant coding in the zebrafish olfactory bulb visualized by optical imaging. *Neuron* **18**, 737–752 (1997).
25. MacLeod, K. & Laurent, G. Distinct mechanisms for synchronization and temporal patterning of odor-encoding neural assemblies. *Science* **274**, 976–979 (1996).
26. Bazhenov, M. *et al.* Model of transient oscillatory synchronization in the locust antennal lobe *Neuron* **30**, 553–567 (2001).
27. Bazhenov, M. *et al.* Model of cellular and network mechanisms for temporal patterning in the locust antennal lobe. *Neuron* **30**, 569–581 (2001).
28. Rabinovich, M. *et al.* Dynamical encoding by networks of competing neuron groups: winnerless competition. *Phys. Rev. Lett.* **87**, 068102 (2001).
29. Gray, C. Synchronous oscillations in neuronal systems, mechanisms and function. *J. Comput. Neurosci.* **1**, 11–38 (1994).
30. Engel, A. K., Fries, P. & Singer, W. Dynamic predictions, oscillations and synchrony in top-down processing. *Nature Rev. Neurosci.* **2**, 704–716 (2001).
References 29 and 30 are good reviews of synchronous periodic phenomena in various brain areas and circuits.
31. Gelperin, A. & Tank, D. W. Odour-modulated collective network oscillations of olfactory interneurons in a terrestrial mollusc. *Nature* **345**, 437–440 (1990).
32. Laurent, G. & Naraghi, M. Odorant-induced oscillations in the mushroom bodies of the locust. *J. Neurosci.* **14**, 2993–3004 (1994).
33. MacLeod, K., Bäcker, A. & Laurent, G. Who reads temporal information contained across synchronized and oscillatory spike trains? *Nature* **395**, 693–698 (1998).
The authors examine the consequences of a pharmacological blockade of synchronization on the specificity of neuronal responses downstream of the desynchronized assemblies. This paper complements references 21 and 22 in testing directly the possible function of oscillatory synchronization in a brain circuit.
34. Laurent, G. *et al.* Odor encoding as an active, dynamical process: experiments, computation and theory. *Annu. Rev. Neurosci.* **24**, 263–297 (2001).
35. Abeles, M. Role of the cortical neuron, integrator or coincidence detector? *Isr. J. Med. Sci.* **18**, 83–92 (1982).
This paper and reference 36 make the case for coincidence detection as a key to understanding transformations by cortical neurons. According to this view, neurons exploit cellular and biophysical properties, as well as the spatiotemporal format of their inputs, to transform representations.
36. König, P., Engel, A. K. & Singer, W. Integrator or coincidence detector? The role of the cortical neuron revisited. *Trends Neurosci.* **19**, 130–137 (1996).
37. Von der Malsburg, C. & Schneider, W. A neural cocktail party processor. *Biol. Cybern.* **54**, 29–40 (1986).
38. Linster, C. & Smith, B. H. Generalization between binary odor mixtures and their components in the rat. *Physiol. Behav.* **66**, 701–707 (1999).
39. Schidberger, K. Local interneurons associated with the mushroom bodies and the central body in the brain of *Acheta domesticus*. *Cell Tissue Res.* **230**, 573–586 (1983).
40. Grünewald, B. Morphology of feedback neurons in the mushroom body of the honeybee, *Apis mellifera*. *J. Comp. Neurol.* **404**, 114–126 (1999).
41. Marr, D. A theory of cerebellar cortex. *J. Physiol. (Lond.)* **202**, 437–470 (1969).
Marr examines the consequences of cerebellar architecture on the possible format of memories and on the management of overlaps between memories.
42. Kanerva, P. *Sparse Distributed Memory* (MIT Press, Cambridge, Massachusetts, 1988).
Kanerva examines (more generally than Marr) the case of sparse and distributed representations and the application of this thinking to cerebellar architecture. This book and reference 41 are interesting reading in general, and especially in view of experimental findings summarized in the current review.
43. Freeman, W. J. *Neurodynamics: an Exploration in Mesoscopic Brain Dynamics* (Springer, London, 2000).
44. Buck, L. B. Information coding in the vertebrate olfactory system. *Annu. Rev. Neurosci.* **19**, 517–544 (1996).
45. Axel, R. The molecular logic of smell. *Sci. Am.* **273**, 130–137 (1995).
46. Mombaerts, P. *et al.* Visualizing an olfactory sensory map. *Cell* **87**, 675–686 (1996).
47. Vassar, R. *et al.* Topographic organization of sensory projections to the olfactory bulb. *Cell* **79**, 981–991 (1994).
48. Stopfer, M. & Laurent, G. Short-term memory in olfactory network dynamics. *Nature* **402**, 664–668 (1999).
49. Dong, D. W. & Atick, J. J. Temporal decorrelation — a theory of lagged and nonlagged responses in the lateral geniculate nucleus. *Netw. Comput. Neural Syst.* **6**, 159–178 (1995).
50. Dan, Y., Atick, J. J. & Reid, R. C. Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. *J. Neurosci.* **16**, 3351–3362 (1996).
51. Vinje, W. E. & Gallant, J. L. Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* **287**, 1273–1276 (2000).
52. Bergman, H. & Bar-Gad, I. Stepping out of the box, information processing in the neural networks of the basal ganglia. *Curr. Opin. Neurobiol.* **11**, 689–695 (2001).
53. Schürmann, F. W. On the functional anatomy of the corpora pedunculata in insects. *Exp. Brain Res.* **19**, 406–432 (1974).
54. Zars, T. *et al.* Localization of a short-term memory in *Drosophila*. *Science* **288**, 672–675 (2000).
55. Dubnau, J., Grady, L., Kitamoto, T. & Tully, T. Disruption of neurotransmission in *Drosophila* mushroom body blocks retrieval but not acquisition of memory. *Nature* **411**, 476–480 (2001).
56. McGuire, S. E., Le, P. T. & Davis, R. L. The role of *Drosophila* mushroom body signaling in olfactory memory. *Science* **293**, 1330–1333 (2001).
References 54–56 report recent experiments geared towards identifying the locus or loci of olfactory/associative memories in *Drosophila*. All point to KCs as being crucial cellular elements, although it is still unclear which synaptic sites are modified and what their roles are for memory or recall.
57. Cajal, S. R. *Histology of the Nervous System* (Oxford Univ. Press, New York, 1995).
58. DeVries, S. H. & Baylor, D. A. Synaptic circuitry of the retina and olfactory bulb. *Cell* **72**, 139–149 (1993).
A good review summarizing the proposal that local circuits in the OB serve to sharpen the tuning curves of MCs in a process akin to that which occurs in retinal local circuits. This paper and reference 15 are focused on single neuron data and on the traditional tuning-curve assessment of neuronal responses; they are to be contrasted with the systems and dynamic perspective that is proposed in the current review.
59. Mori, K., Kishi, K. & Ojima, H. Distribution of dendrites of mitral, displaced mitral, tufted and granule cells in the rabbit olfactory bulb. *J. Comp. Neurol.* **219**, 339–355 (1983).
60. Orona, E., Rainer, E. C. & Scott, J. W. Dendritic and axonal organization of mitral and tufted cells in the rat olfactory bulb. *J. Comp. Neurol.* **226**, 346–356 (1984).
61. Hansson, B. S., Anton, S. & Christensen, T. A. Structure and function of antennal lobe neurons in the male turnip moth, *Agrotis segetum* (Lepidoptera, Noctuidae). *J. Comp. Physiol. A* **5**, 547–562 (1994).
62. Stocker, R. F. The organization of the chemosensory system in *Drosophila melanogaster* — a review. *Cell Tissue Res.* **275**, 3–26 (1994).
63. Chen, W., Midtgaard, J. & Shephard, G. Forward and backward propagation of dendritic impulses and their synaptic control in mitral cells. *Science* **278**, 463–467 (1997).

Acknowledgements

The work from my laboratory reviewed here was funded by the National Science Foundation, the National Institute on Deafness and other Communication Disorders, and the McKnight, Keck, Sloan and Sloan-Swartz Foundations. I thank M. Stopfer, R. Friedrich, K. McLeod, M. Wehr, J. Perez-Orive, O. Mazor, S. Cassenaer, R. Wilson, G. Turner, C. Pouzat, V. Jayaraman, S. Farivar, H. Davidowitz, R. Jortner, A. Holub, M. Rabinovich, H. Abarbanel, R. Huerta, T. Nowotny, V. Zighulin, A. Bäcker, M. Bazhenov, P. Perona and E. Schuman for the privilege of working on these problems with them. I thank P. Cariani for pointing me to Kanerva's book on sparse distributed memories, K. Heyman for secretarial assistance and S. Farivar for Golgis in figure 4.

Online links

FURTHER INFORMATION
Encyclopedia of Life Sciences: http://www.els.net/olfaction | olfactory receptor neurons
Sense of Smell Institute: http://www.senseofsmell.org/
The Laurent Lab: http://marvin.caltech.edu/
Access to this interactive links box is free online.