RESEARCH ARTICLE

Optimality of sparse olfactory representations is not affected by network plasticity

Collins Assisi[1]*, Mark Stopfer[2], Maxim Bazhenov[3]

1 Division of Biology, Indian Institute of Science Education and Research, Pune, India, **2** NICHD, National Institutes of Health, Bethesda, Maryland, United States of America, **3** Department of Medicine, Division of Pulmonary, Critical Care and Sleep Medicine, University of California San Diego, La Jolla, California, United States of America

* collins@iiserpune.ac.in

Abstract

The neural representation of a stimulus is repeatedly transformed as it moves from the sensory periphery to deeper layers of the nervous system. Sparsening transformations are thought to increase the separation between similar representations, encode stimuli with great specificity, maximize storage capacity of associative memories, and provide an energy efficient instantiation of information in neural circuits. In the insect olfactory system, odors are initially represented in the periphery as a combinatorial code with relatively simple temporal dynamics. Subsequently, in the antennal lobe this representation is transformed into a dense and complex spatiotemporal activity pattern. Next, in the mushroom body Kenyon cells (KCs), the representation is dramatically sparsened. Finally, in mushroom body output neurons (MBONs), the representation takes on a new dense spatiotemporal format. Here, we develop a computational model to simulate this chain of olfactory processing from the receptor neurons to MBONs. We demonstrate that representations of similar odorants are maximally separated, measured by the distance between the corresponding MBON activity vectors, when KC responses are sparse. Sparseness is maintained across variations in odor concentration by adjusting the feedback inhibition that KCs receive from an inhibitory neuron, the Giant GABAergic neuron. Different odor concentrations require different strength and timing of feedback inhibition for optimal processing. Importantly, as observed *in vivo*, the KC–MBON synapse is highly plastic, and, therefore, changes in synaptic strength after learning can change the balance of excitation and inhibition, potentially leading to changes in the distance between MBON activity vectors of two odorants for the same level of KC population sparseness. Thus, what is an optimal degree of sparseness before odor learning, could be rendered sub–optimal post learning. Here, we show, however, that synaptic weight changes caused by spike timing dependent plasticity increase the distance between the odor representations from the perspective of MBONs. A level of sparseness that was optimal before learning remains optimal post-learning.

Author summary

Kenyon cells (KCs) of the mushroom body represent odors as a sparse code. When viewed from the perspective of follower neurons, mushroom body output neurons (MBONs) reveal an optimal level of coding sparseness that maximally separates the representations of odors. However, the KC–MBON synapse is highly plastic and may be potentiated or depressed by odor–driven experience that could, in turn, disrupt the optimality formed by pre–synaptic circuits. Contrary to this expectation, we show that synaptic plasticity based on spike timing of pre- and postsynaptic neurons improves the ability of the system to distinguish between the representations of similar odors while preserving the optimality determined by pre–synaptic circuits.

Introduction

The neural representation of an odor is transformed repeatedly as it traverses different layers of the olfactory system [1]. Some transformations separate the representations of odorants to enable easy discrimination [2][3]. Other transformations prepare an odor representation for eliciting behaviors by associating it with other sensory inputs and providing the context necessary for action and memory [4]. In the American desert locust, *Schistocerca americana*, neural networks peripheral to the KC–MBON synapse appear to work best as pattern decorrelators while downstream circuits appear to be specially structured to encode associative memories and organize behaviors elicited by stimuli. The olfactory network, from the receptor neurons through the antennal lobe and on to the mushroom body, is largely feedforward, and odor representations are progressively decorrelated and optimized in several ways as they traverse these layers. Odor representations arrive at the MBONs via synapses that are highly plastic and may change with the dynamic olfactory milieu of the animal [5]. Here, we ask, how does the olfactory network preserve an optimal odor representation despite activity–driven changes in the synaptic weights of the networks?

In the locust, olfactory processing in the nervous system begins when odorant molecules bind to receptors on neurons in the antennae. This leads to the opening of the receptor neuron's ion channels and a cascade of events that can lead to spiking, the suppression of spontaneous firing, or simple sequences of excitation and inhibition. Olfactory receptor neurons can be tuned narrowly or broadly, firing vigorously for some odors and less so or not at all for others [6,7]; thus, the identity of responsive receptor neurons helps encode the stimulus. Temporal features of receptor neuron spiking, including simple sequences of excitation and inhibition, also contribute to encoding the identity of the odor [8]. Olfactory receptor neurons provide input to excitatory PNs and local inhibitory (and likely some excitatory) interneurons in the antennal lobe. This dense network, with recurrent connections between excitatory and inhibitory neurons, transforms the odor representation arising in receptor neurons into a more elaborate spatiotemporal pattern [1,9,10] where the identity, concentration, and timing of the odor are represented by the identity of responsive PNs, the temporal structure of their spiking, and correlations across the PN population. Most PNs respond in some way to most odors [1,11], collectively providing a dense spatiotemporal representation of an odor. KCs in the mushroom body receive inputs from PNs and transform this dense representation into a sparse code [12] in which rare spikes occur with millisecond precision and great specificity, together describing the attributes of the eliciting odor. The sparseness of KC spiking is orchestrated by a combination of membrane conductances that ensure a high spike threshold, and feedback inhibition from a giant GABAergic neuron (GGN) proportional to the drive it

receives from the full population of KCs [13]. Thus, GGN adaptively regulates the output of KCs, maintaining the sparseness of their code over a large range of odor concentrations. The successive transformations that odor representations undergo, from dense spatiotemporal to sparse activity patterns, are thought to progressively decorrelate and distinguish odor representations and prepare them for valence and motor decisions, and storage as memories [14]. A transformation of the representation from dense to sparse is also accompanied by an expansion in the dimension of the neuronal representations (such as from 830 PNs in the AL to nearly 50,000 KCs in the MB of the locust). A similar transformation from a dense to a sparse representation is seen in different species and a number of brain areas [15–17]. The expansion maps similar inputs to widely separated outputs in a high dimensional space. If the dimensionality of the representation is sufficiently high, similar odors can be distinctly classified even when the number of inputs to individual KCs is low [18]. This mapping format has the risk of amplifying noisy representations of the same odor. However, structured connectivity, where synaptic weights reflect the correlations of the inputs, make the sparse representation resilient to noise[19]. In locusts, convergent KC activity is read out by a relatively small number of MBONs. The KC–MBON synapse undergoes experience dependent plasticity [5] (see [20] for a similar circuit in Drosophila) in a form that can be modified by associating an olfactory stimulus with a reward [5,20] mediated by octopamine. Together these features mark the KC–MBON pathway as one where sparse, decorrelated odor representations are combined with input from a reward pathway [21].

Using a model network that simulates olfactory processing in the locust from receptor neurons to MBONs, we show that the distance between the representations of different odorants, measured as the distance between MBON activity vectors, is maximized for a particular level of KC response sparseness. The degree of sparseness is determined by transformations of the odor representation in circuits before the KC–MBON synapse. However, what level of sparseness is optimal for odor discrimination by MBONs is determined by the weights of KC–MBON synapses. KC–MBON synaptic weights are, in turn, subject to the animal's experiences, mediated by octopamine reward. This gives rise to a potential conundrum: the degree of sparseness determined by the circuits prior to the KC–MBON synapse could be rendered sub–optimal by modulations to the weight of that KC–MBON synapse by associative learning of particular odors. Here, we explore how the olfactory system guards against this loss of optimal sparseness. We show that the spike timing dependent plasticity operating on the strength of KC–MBON synapse not only retains the value of optimal sparseness despite learning-dependent changes in synaptic strength, but further improves the ability of the olfactory system to differentiate between odors.

Though we focus on odor discrimination here, this need not be the sole optimizing principle operating in the KC–MBON circuit. In addition, MBONs likely play an important role in generalizing the representations of learned odors and associate it other sensory inputs.

Results

In this study we sought to address two questions. First, from the perspective of MBONs, is there an optimal value of coding sparseness to maximally separate odor representations? Second, if an optimal sparseness exists, does plasticity at the KC–MBON synapse alter it, making post–learning odor representations sub–optimal? To address these questions, we constructed a computational model of the locust olfactory system consisting of the olfactory receptor neurons, the antennal lobe network of PNs and local inhibitory interneurons, the KCs of the mushroom body, and a layer of MBONs (Fig 1A). The model antennal lobe network generates many of the key responses previously documented in the locust [22,23]. The output from the

a)

b)
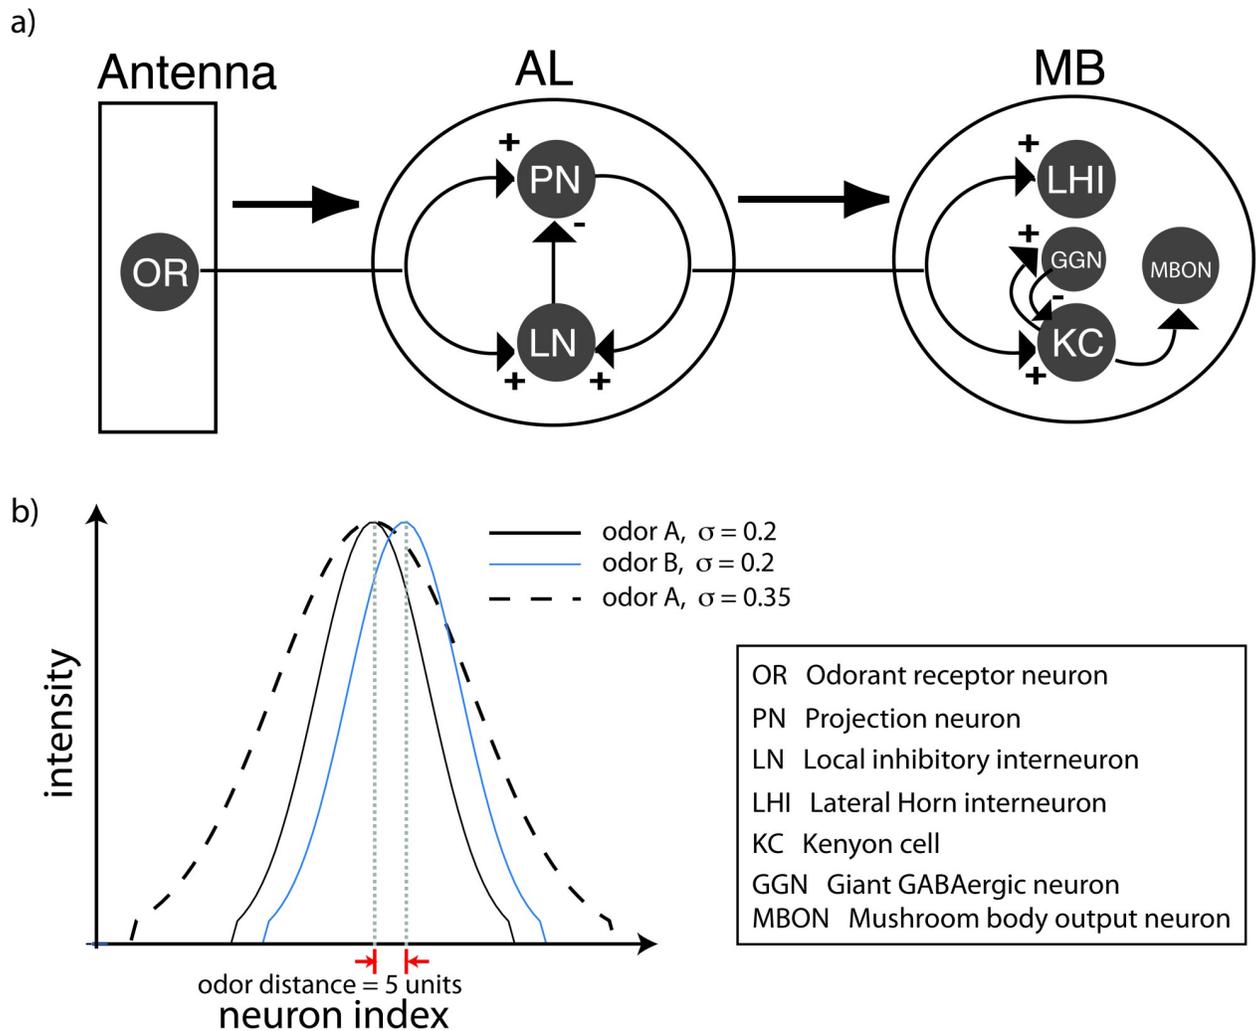
Fig 1. a) Schematic of the locust olfactory system. The first layer consists of olfactory receptor neurons (OR) in the antenna that provide input to PNs and local inhibitory interneurons (LNs) in the antennal lobe (AL). PNs project to KCs of the mushroom body (MB). In addition to excitatory input from the PNs the KCs also receive inhibitory input from the Giant GABAergic neuron (feed–back inhibition). b) Input to the AL. Each PN (and LN) indexed on the x—axis receives a dc input of intensity shown along the y–axis. The concentration of the odor is characterized by the width of the function describing the input intensities while the identity of the odor is characterized by the position of its the peak.

antennal lobe diverges widely to an array of 15,000 model KCs. This pattern of connectivity has been hypothesized to help decrease the overlap between odor representations [24,25]. KC output then converges onto a small group of MBONs. To establish whether an optimal value of sparseness exists, we systematically varied the sparseness of KC responses and checked the ability of MBONs to differentiate between two similar odorants. We then introduced spike timing dependent plasticity [5] in the KC–MBON synapse and simulated the network using multiple instances of randomly interleaved odorants to map the effect of synaptic plasticity on the optimal sparseness of KC responses.

Coding sparseness determines the distance between odor representations

Odor input activates the ORNs that drive the neurons of the antennal lobe. We did not explicitly model the ORNs; rather, we simulated ORN activity as a simplified, constant supra-
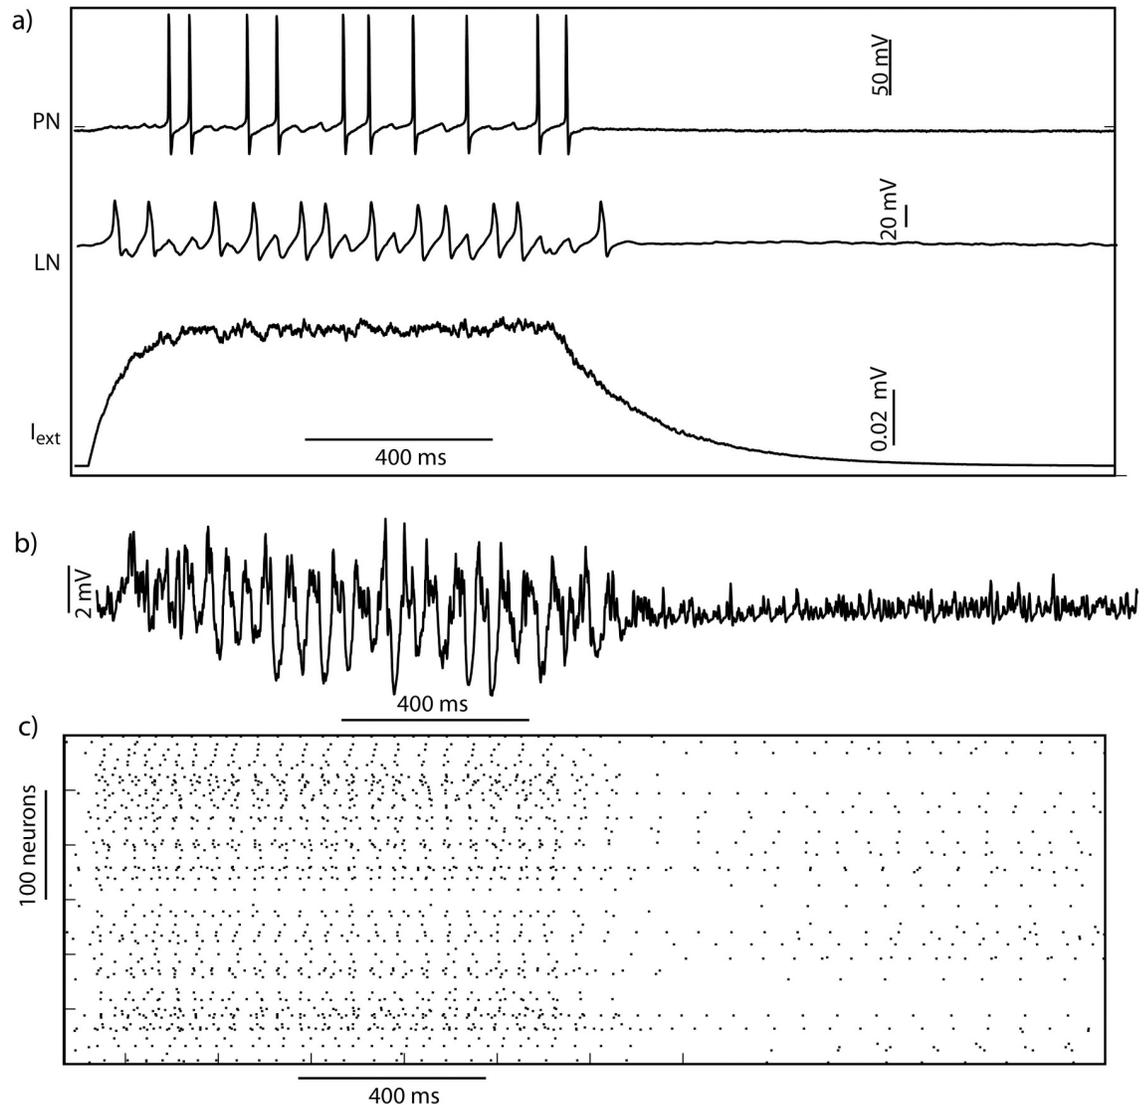
Fig 2. The top trace in (a) shows the response of a single PN to an external input (a, bottom trace) while the middle trace shows the response of an LN. The trace in (b) shows the summed projection neuron (PN) voltage response to an external input to the network. The input (I_{ext}) (a, bottom trace) was scaled by intensity (shown in Fig 1B) for each neuron. The PN raster plot for a reduced network with 300 PNs and 100 local neurons (LNs) is shown in c.

threshold depolarizing input to a subset of PNs and interneurons [26–28]. This input had an initial rise time of 60ms and a decay time of 200ms (Fig 2C bottom trace). The amplitude of the input remained constant for 1000ms except for noise that was 5–10% of the amplitude of the pulse. Each odorant was defined by the subset of antennal lobe neurons it activated and the amplitude of the depolarizing input to each. In Fig 1B the amplitude of input to the PNs of the network is shown for two odorants (solid lines). The PNs were arranged such that the amplitude profile resembled a Gaussian curve. The input curve was set to zero when the amplitude decreased below a threshold value. Note that the arrangement of PN indices according to a Gaussian activation profile does not imply any spatial structure; that is, the neuron with index *i* need not be physically adjacent to neurons with index *i-1* or *i+1* since network connections were chosen randomly. However, by defining an odor in this manner, we could conveniently

and continuously vary the identity and the concentration of odorants. The identity of an odorant could be varied by moving the location of the peak while the concentration could be increased by widening the Gaussian to recruit more PNs and interneurons (Fig 1B, dashed line) [29].

We measured the responses of antennal lobe model neurons to the odor input. As seen in earlier studies and in accordance with experiments done *in vivo*, the local field potential (measured in our model as the mean membrane potential of all the PNs) showed an odor-elicited 20Hz oscillation (Fig 2B). This global pattern was elicited by different odors and concentrations. The oscillations emerge from interactions between PNs and LNs [28]: reciprocally coupled pairs of PNs and LNs oscillate with phase shift at ~20Hz when driven by an external depolarizing input. PN spikes elicit an LN spike that in turn delays the onset of subsequent spikes in post-synaptic PNs. The extent of this delay is a function of the strength of fast-GABAergic inhibition between LNs and PNs. For sufficiently high values of inhibitory coupling, the delay is determined by the decay time-constant of fast inhibition [30]. In the network of 100 LNs and 300 PNs, rhythmically spiking LNs synchronized PNs into transiently synchronous groups. Earlier studies have shown that when GABAergic interactions are blocked by picrotoxin, synchronization of PNs and the 20Hz oscillatory LFP are lost [30],[31]. Thus, inhibition plays an important role in synchronizing the activity of PNs and generating an oscillatory local field potential. As *in vivo*, in the model during each cycle of the oscillation different groups of PNs were transiently activated. This spatiotemporal representation continuously changed over the course of the odor presentation due to mutual and transient inhibition between interneurons that, in turn, coordinated the activity of PNs [26–28].

Each odor stimulus consisted of a 1000ms constant (except for a 10% noise) depolarizing input to a subset of PNs and LNs. The temporal response of the AL neurons, in contrast, changes over the time-scales ranging from 10–100's of milliseconds. Thus, the AL response reflects the intrinsic properties of PNs and LNs and the topology of the networks they form. One of the key drivers of the spatiotemporal patterning in the locust AL is spike frequency adaptation which causes the firing rates of LNs to decrease over the duration of odor input [28]. Adaptation in LNs is caused by a Ca^{2+} dependent potassium current that builds up over time and delays the onset of subsequent spikes. As the spiking frequency of a neuron decreased, post-synaptic LNs were released from inhibition and became active. Thus, different groups of LNs were sequentially activated in response to a depolarizing input. The identity and the order of these spatiotemporal patterns were determined by the structure of the network [26]. As a result, different LNs and post-synaptic PNs are activated at different times during the odor presentation.

Each odor–concentration pair we tested elicited a different spatiotemporal pattern. An example pattern of activity is shown in Fig 2A and 2C. The amplitude of the local field potential increased with increasing concentration indicating tighter synchrony between the projection neurons that spike during each cycle, consistent with earlier studies [29,32].

The PN responses were used as input to a group of 15,000 KCs. KCs are known to respond sparsely (few neurons fire rarely) to odor stimuli [12]. However, the increased PN synchrony that accompanies increased odor concentrations [32] alone would lead to more densely spiking responses, disrupting the sparse code. How do KCs maintain sparseness over decadal variations in the concentration of the odor? Earlier studies hypothesized that input from PNs to KCs arrives along two pathways, a direct excitatory drive from the antennal lobe and slightly delayed feedforward inhibition from lateral horn interneurons (LHIs). Thus, cyclic pairs of excitatory and inhibitory input to KCs defined short windows of time during which KCs could integrate input from the antennal lobe [12]. Notably, the duration of this window was dynamically modulated by changes in the concentration of the odor, which allowed KCs to fire

sparsely despite large changes in the concentration [29]. Recent work established that LHIs do not extend GABAergic projections to KCs, eliminating the possibility of feedforward inhibition [33]. However, the cyclic inhibition underlying the dynamically modulated integration windows is now known to be generated by feedback from a single inhibitory cell, termed the Giant GABAergic Neuron (GGN), which provides input to all KCs [13,22]. Thus, cyclic inhibition regulates the sparseness of KC responses in an adaptive, concentration dependent manner [33,34].

Given these earlier findings, we first sought to determine whether there exists an optimal value of lifetime sparseness (measured as the total number of spikes generated by all KCs during an odor presentation) to discriminate odors. To determine an optimal sparseness, if one such value existed, we needed to systematically vary the sparseness of KC responses and quantify the distance between odor representations from the perspective of downstream neurons that read KC output. KCs converge onto MBONs that generate distinct responses to odorants [35]. Therefore, we used the MBONs as a read–out of KC responses. The Hamming distances between odor representations generated by MBONs were plotted as a function of different manipulations to the upstream network that changed KC lifetime sparseness. In the locust olfactory system GGN regulates the sparseness of KC responses using feedback inhibition. Numerous studies have shown that feedback inhibition in excitatory-inhibitory circuits mainly reduces later portions of the excitatory responses in each cycle. As inhibition strengthens, its onset occurs faster, thus reducing the excitatory response [29]. This effect is self-limiting though, because excitation is needed to drive inhibition [34]. If we explicitly modeled the GGN, we would obtain a single window of integration for each concentration [34,36]. However, we were interested in understanding how MBON responses varied as a function of varying KC sparseness that is, in turn, dependent upon the width of the window of integration. Thus, rather than modeling the GGN, we modeled the effect of feedback inhibition by selectively eliminating PN spikes that occurred after a threshold phase of the LFP (Fig 3A). To do so we first filtered the LFP (40Hz) (red trace in Fig 3A) and calculated the instantaneous phase of the resulting 20 Hz oscillation using a Hilbert transform. Then we removed those spikes occurring beyond a threshold phase, denoted by ϕ, in each cycle of the LFP, shown by the shaded regions in Fig 3A. In this way we could directly control this threshold phase, and therefore artificially vary the window of integration and the sparseness of KC activity. As expected, as we widened the window of integration, KCs received more input from PNs and generated progressively denser spiking. We plotted the density of KC spiking as a function of ϕ for different values of odor concentration (σ) (Fig 3C). Increasing odor concentrations led to denser KC responses for a given value of ϕ. To maintain a specific value of sparseness that maximally separated odor representations across a range of odor concentrations, the window of integration had to be shifted to lower values as the concentration increased. In the locust olfactory system a leftward shift in phase is achieved by a mechanism that arises naturally from concentration–dependent changes in odor–response latencies in PNs [29,33].

To determine the separation between odor representations from the perspective of downstream targets, we used odor–elicited KC responses to drive a group of 100 MBONs (Fig 3B). Here the MBONs were modeled using a two–dimensional map that integrates pre–synaptic input and generates spikes in response to it (see Methods). We binned the output of these neurons into temporal blocks, each block demarcated by the troughs of an LFP cycle. The response of MBONs during each cycle of an LFP oscillation was assigned one or zero to indicate whether it had spiked or not (Fig 3B).

We then calculated the Hamming distance between the representations of two odorants as a function of different windows of integration ϕ (Fig 4). Since the density of KC spikes was a monotonically increasing function of ϕ, we used ϕ as a proxy for KC sparseness. We found
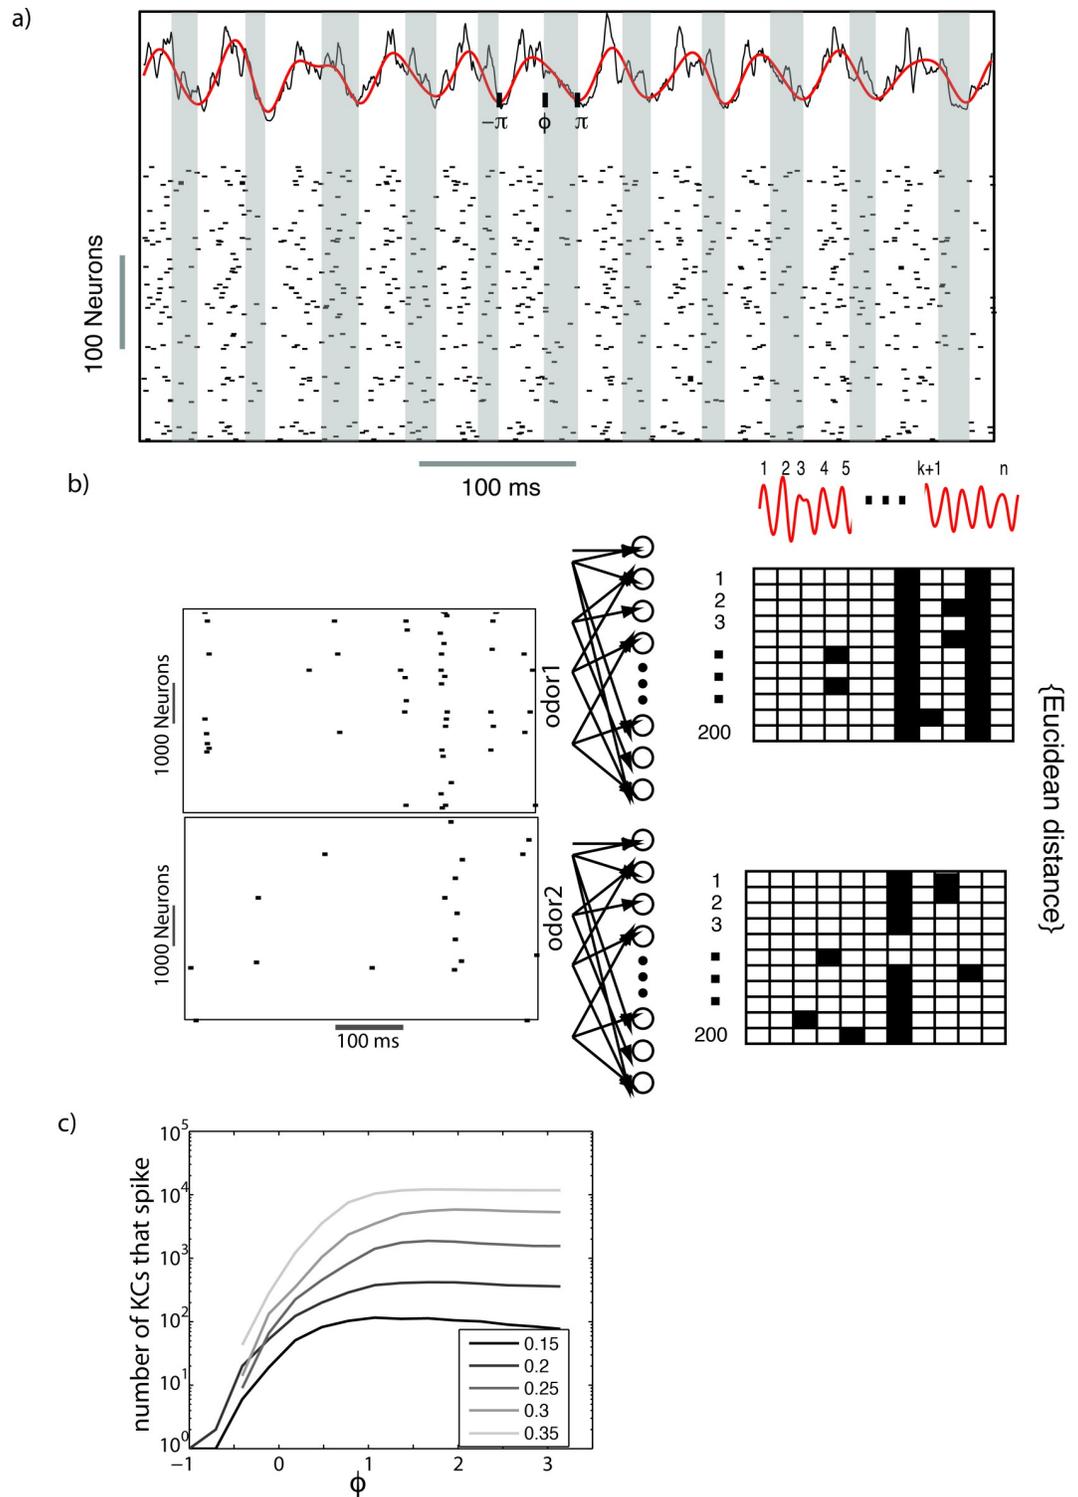
Fig 3. Individual KCs receive oscillatory input from the PNs in the AL (Top trace in (a)). The sparseness of the KC responses was modulated by choosing a window of time over which each KC integrated input from PNs. That is, spikes occurring after a specified phase (ϕ), in the shaded region in (a), were ignored and did not affect the KC responses. The raster plots in (b) show the responses of a subset of KCs to two different odors. These spikes were then fed to a layer of 100 beta lobe neurons. The response of the beta lobe neurons was converted into a binary spatiotemporal pattern. For each neuron, a single cycle of the PN-LFP was marked either 1 (dark) or 0 (blank box), depending on whether that neuron fired a spike during the cycle (Panels on the right in b). The Euclidean distance between these binary spatiotemporal patterns was used to calculate the distance between odors. The number of KCs that spike in response to an external input is plotted in (c).
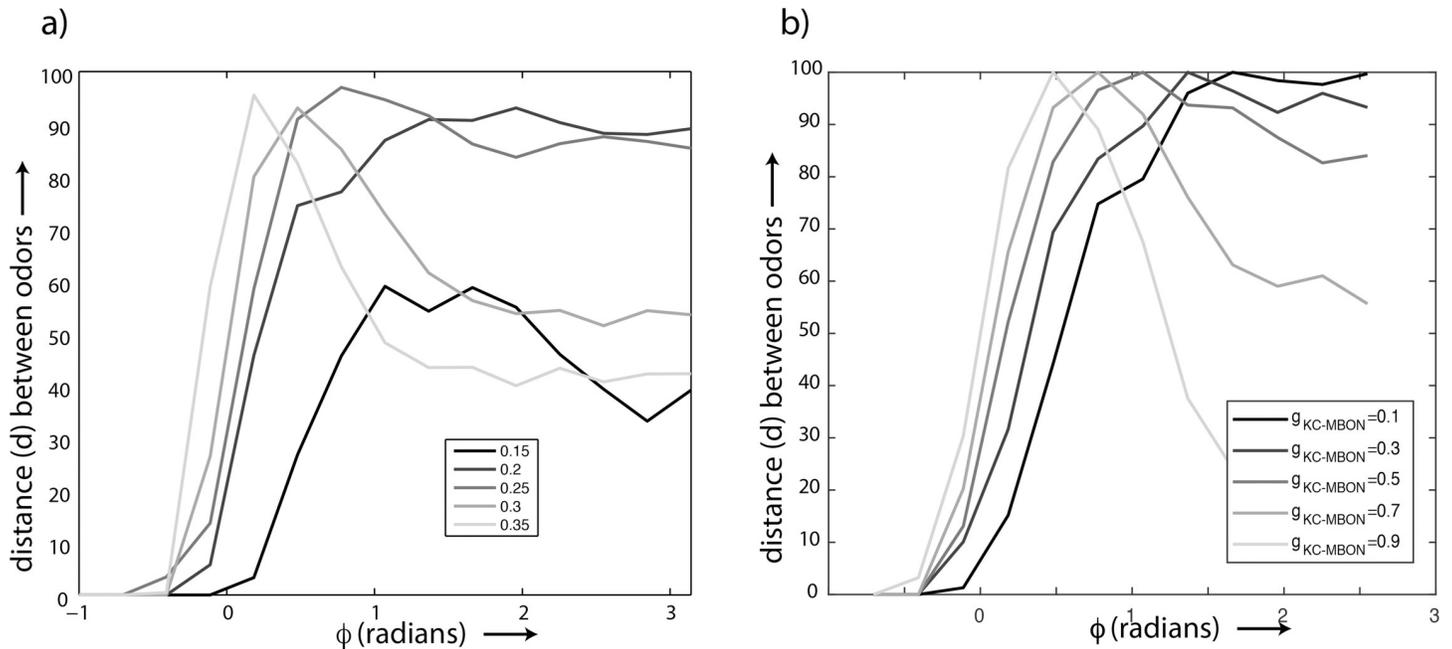
Fig 4. (a) Distance between odors plotted as a function of the window of integration for different values of the concentration (σ). b) Distance between odors for different KC–MBON synaptic weights(ϕ).

that, for low odor concentration values ($\sigma < 0.25$), the peak distance between odor representations occurred when more PN spikes were allowed to affect KC responses throughout each LFP cycle (Fig 4A). Indeed, for low odor concentrations, the synchronization and thus density of input PN spiking remained low and the responses of KCs remained sparse throughout the range of integration windows we simulated. With increasing odor concentrations, we found that the peak distance between odor representations shifted to lower values of ϕ. Furthermore, we observed a decrease in discrimination performance beyond a certain ϕ threshold. This occurred because KC responses became denser when the integration window expanded. Thus, for higher odor concentrations, we found a prominent single peak suggesting the existence of an optimal value of KC sparseness to maximize the distance between odor representations from the perspective of MBONs. Experimental recordings of lateral horn interneurons that receive convergent input from PNs show a phase shift [37] similar to shifts observed in models [29,34]. This is expected since increasing concentration also increases the amplitude of the LFP due to increased PN synchrony. GGN responds in a graded manner to PN inputs. The phase of the peak response of GGN with respect to the LFP remains invariant to changes in the concentration. GGNs respond to increased PN synchrony with stronger IPSPs that rise faster than weaker ones and regulate the window over which PN spikes are integrated [37].

Within a given animal the impact of KC spiking on MBONs can vary over time because the synapses linking them are plastic, changing in strength with experience [5]. By amplifying or decreasing the impact of KC spiking, this synaptic plasticity has the potential to degrade the effective, optimized sparseness of the KC output, potentially affecting the distance between odor representations from the perspective of MBONs. To investigate this possibility, we systematically varied the weight of the input synapses to MBONs to determine how plasticity affects the distance between odor representations. We then simulated delivery of two similar odors of the same concentration by shifting the peaks of the distributions that characterized the two odors by 5 units with respect to each other, and for two different concentrations by
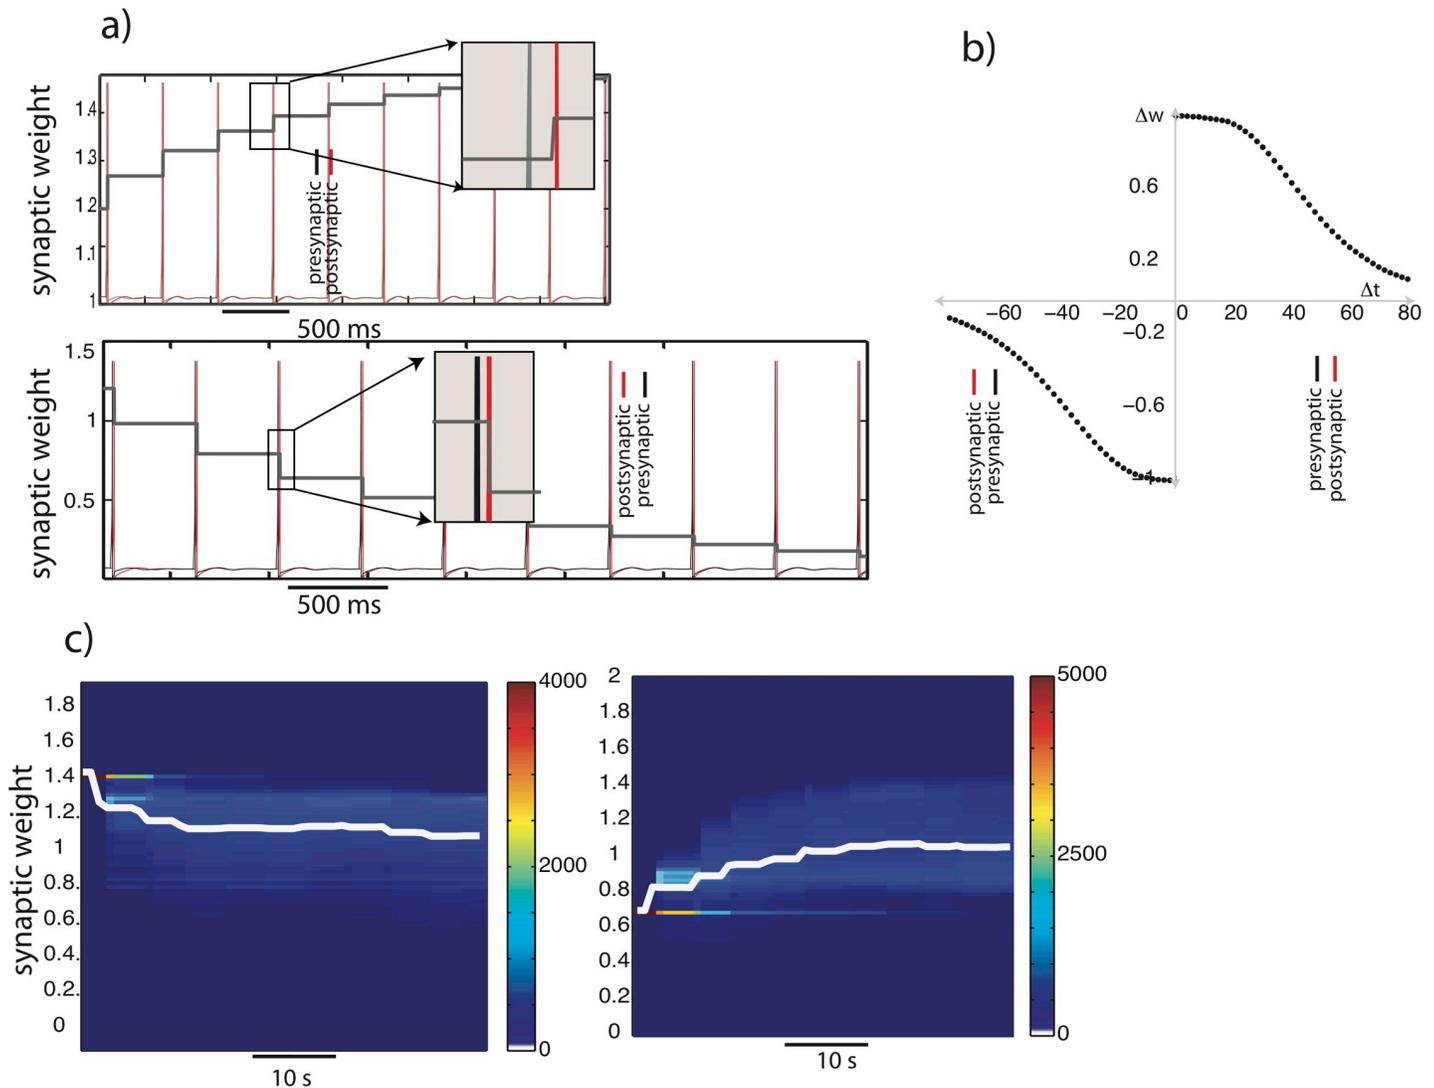
Fig 5. Spike timing dependent plasticity. The post–synaptic neuron (red) spikes follow that of the pre–synaptic neuron (black), leading to an increase in the synaptic weight (facilitation) (Fig 5A top panel). The opposite temporal order (post–synaptic spikes occur before pre–synaptic spikes) leads to a decrease in synaptic weights (depression) (Fig 5A bottom panel). The increase/decrease in synaptic weight (Δw) is shown as a function of the time difference between the pre–and the post–synaptic spike (5b). When the presynaptic spike occurs before the post–synaptic spike is positive and otherwise negative. The distribution of synaptic weights of all KC-MBONs pairs evolves over time (c). In the left panel of (c) all the initial weights were set to 1.4. The system was then stimulated with different odors of varying concentrations. The weights were sampled at fixed intervals of time. The distribution of weights was plotted using a color map (see color bar for the frequency values). The mean synaptic weight was overlaid on the distribution (white trace). The right panel shows the temporal evolution of the synaptic weights when a low initial weight (0.6) was used.

adjusting the widths of the distributions (Fig 1B). As before, patterns of antennal lobe activity served as input to KCs that, in turn, drove MBONs. Here, we used the output of MBONs to measure the distance between odor representations for different values of KC sparseness. The MBONs were modeled as simple map–based neurons, summing the input they received from KCs and generating a spike in response to supra–threshold inputs. We then systematically varied the weights of the synapse from 0.1 to 0.9. This manipulation led to a shift of the peak towards lower values of ϕ (Fig 4B). These simulations confirmed that the effective sparseness of KC output could change when the animal is exposed to different sets of odors that trigger plasticity in the KC–MBON pathway. This departure from optimal sparseness could be

detrimental to subsequent circuits that, to appropriately inform behaviors, depend on an accurate distinction between odors.

Optimal sparseness persists despite spike timing dependent plasticity

KC–MBON synapses appear to be powerful: *in vivo*, a KC spike generates an EPSP in MBONs that is, on average, nearly an order of magnitude larger than EPSPs generated in KCs by PN spikes [5]. Previous work established that the KC–MBON synapse undergoes spike timing dependent plasticity (STDP): potentiation when the presynaptic neuron fires before the post synaptic one, and depression when the presynaptic neuron fires after the post synaptic one [5]. This plasticity has been shown to maintain the oscillatory parcellation of information that begins at the antennal lobe and cascades all the way down to the MBONs. How does this plasticity affect the distance between odor representations when viewed from the perspective of MBONs?

To address this question, we modeled STDP in the KC–MBON synapse using a simple phenomenological model (Fig 5B) [38,39]. Following STDP rules, the model effectively modified the weight of the synapse depending on the time of occurrence of the presynaptic and the postsynaptic spikes such that each occurrence of a presynaptic spike before the postsynaptic spike led to an increase in the synaptic weight, and a presynaptic spike after the postsynaptic spike led to a decrease in the synaptic weight (see Methods for implementation details). We modeled MBONs parsimoniously as reduced spiking neuron models represented by two dimensional maps [40,41] (see Methods for details of implementation). The increase/decrease in weights with each pre–post pair of spikes is shown in Fig 5A. We specified a minimum and a maximum value for the synaptic weights so that the response of MBONs extended over a wide range of sparseness values. The change in synaptic weights (dw) depended on how close the current weight of the synapse was to the maximum allowed synaptic weight. The STDP equations were modeled so that for large synaptic weights (w_{max}) synaptic depression dominates over potentiation and vice versa for small synaptic weights (w_0) [42]. Using this form of STDP, we then wired the 15,000 KCs to a layer of 100 MBONs. Each MBON received input from a randomly selected group consisting of 60% of the KCs. For simplicity, we did not implement lateral inhibitory connections between MBONs that are thought to enhance the contrast of input received from KCs [5].

To model odor stimulation, we randomly interleaved multiple instances of two similar odors (peak shifted by 5 units) and an odor that was different from these odors (peak shifted by 20 units) as input to the PNs. This simulated odor input evoked spatiotemporal patterns of activity in the antennal lobe that drove the KCs and the MBONs. Initially, the narrow distribution of synaptic weights of the KC–MBON synapses was, in separate simulations, centered around two different values (Fig 5C, left vs right). Over successive odor presentations these synaptic weights changed. The median value of the synaptic weight is shown by the white lines in Fig 5C. Since KC responses are very sparse, most of the weights did not change at all. Therefore, in subsequent analyses we chose a subset of weights that changed during the course of multiple odor presentations. We found that, over odor presentations, the distribution of synaptic weights evolved in a manner such that the median synaptic weight changed monotonically toward new value (approximately 1 in Fig 5C).

In our simulations we used two different initial weights distributions. Regardless of the specific initial weight distribution, we found that the median synaptic weight evolved towards the same value over multiple odor presentations. Thus, runaway excitation of the small population of MBONs would degrade the representation of an odor by increasing the overlap between nearby odor representations. However, STDP acts as a homeostatic mechanism that maintains the level of activity of MBONs.

Next, we investigated the evolution of the odor representation exhibited by MBONs concomitant with the STDP-dependent evolution of the network weights. Fig 6 shows the distance between odor representations as a function of φ as the network weights evolved (Fig 6A; lighter colored curves correspond to the weights later in training). We found that the location of the peak (optimal value of φ) remained the same despite changes in weights. At a low odor concentration ($\sigma = 0.2$) the changes in the distance curve (marked in progressively lighter shades) were small compared to the changes at higher concentrations ($\sigma = 0.35$, red curves in Fig 6A and 6B). In all cases, however, the optimal degree of sparseness provided by the circuits before the KC–MBON synapse remained optimal after STDP-mediated changes of KC–MBON projections. However, there was a small, but significant change in the peak distance between odor representations. The peak distance between the representations of two similar odors increased (Fig 6B) while the weight distribution settled to its asymptotic values (Fig 5C). To test whether that the increase in the distance was significant, we simulated the network with three pairs of odorants over 10 trials. Thus, for each odor pair we obtained 40 values of the distance between MBON odor representations. A paired t-test showed the mean distance between odor pairs post-STDP was significantly different from that before the network was trained with a sequence of odors (Fig 6C). Next, we simulated the network with three different initial weight distributions and randomly shuffled training inputs. The training inputs, as before, consisted of different trials of two similar odors (peaks shifted by 5 units) and one odor that was different from these (peak shifted by 20 units). In all cases the weight distribution evolved such that the median weights asymptotically approached each other (Fig 5C; two initial weight distributions are shown). Here too, the peak distance between odor representations increased post-STDP compared to the response of the pre-trained network (Fig 6D). Next, we classified the odors using a linear discriminant model. We first simulated twenty trials each of two odors and calculated the pairwise Hamming distance between each point. Using multidimensional scaling and the Hamming distance matrix as a measure of similarity between the representations we mapped the odor representations onto a two-dimensional plane. We then used a linear discriminant model to classify the data into two classes. The figures below show that classification boundaries determined by the model before and after the weights were modified by STDP. The points plotted in red indicate the trials that were misclassified while the blue dots show the trials that were correctly classified (Fig 6E and 6F). The simple linear classifier used here shows an improvement in the classification accuracy after learning (Fig 6F). Thus, we conclude, an effect of STDP was to improve the ability of the olfactory system to differentiate between odors (Fig 6B). Further, we show the optimal sparseness does not change despite activity dependent changes in the synaptic weights (Fig 6C).

Discussion

In *Drosophila* and in locust, sparseness in KC firing is achieved by intrinsic high firing thresholds and feedback inhibition from a single neuron in each lobe, the anterior paired lateral (APL) neuron in *Drosophila* [43,44] and GGN in locust [13,22,33,35]. This simple architecture, with a single neuron exerting outsized influence over the olfactory system, allows relatively simple experimental perturbations that selectively change the sparseness of KC responses. Indeed, as KC spiking increases in density, the ability of insects to differentiate between similar odors decreases, but the ability to differentiate dissimilar odors is not affected [44]. This observation suggests that decreased sparseness increases the overlap between representations, but that KC representations of dissimilar odors are sufficiently distant and continue to be separable even when sparseness is compromised.
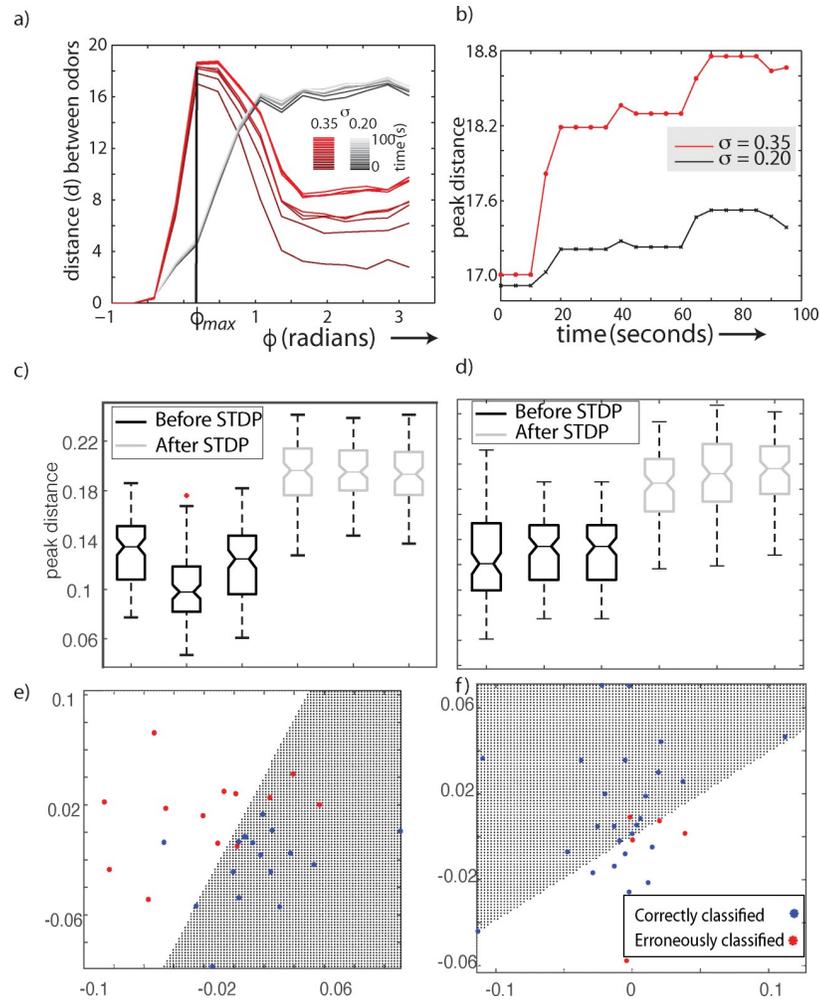
Fig 6. Role of STDP in odor discrimination. Distance between odors plotted as a function of the window of integration of KCs spikes (φ) for high (σ = 0.35 red) and low (σ = 0.2 gray) concentration values (a). Spike timing dependent plasticity reshapes the weights of the KC-beta lobe connections over time. Each line in the plot shows the odor distance–relationship at different snapshots in time ranging from 0 to 100 seconds. Darker shades indicate earlier times. (b) shows the maximum distance between two similar odors for high (σ = 0.35 red) and low (σ = 0.2 gray) concentration values over the time that STDP re–weighted the KC–MBON connections. Note, that the φ value maximizing the distance does not change during STDP-mediated learning. c)Box plot showing the distance between pairs of odors before and after STDP. The difference between the mean distances before and after STDP was statistically significant at p = 0.02. d) Box plot showing the distance between pairs of odors before and after STDP for three different STDP protocols. Each STDP protocol had a different set of initial weights and randomly shuffled learning trials. The difference between the mean distances before and after STDP was statistically significant at p = 0.02. (e-f) The response of the MBONs was projected onto a 2D Euclidean plane while preserving the Hamming distance between points. Points in the figures corresponds to 40 odor trials. Each trial may belong to one of two similar odors. We used a linear discriminant model to classify odors as belonging to one of two classes depending on where it was mapped on the plane. Misclassified odors are shown in red. The number of misclassified odors after learning (f) were less that the number before learning (e).

Here, we developed a model that couples multiple layers of the locust olfactory system. Our results demonstrate that a specific value of the window of integration (φ) of the PN inputs to KCs maximally separates the KC representations of similar odors from the perspective of their follower MBONs. As the concentration increased we found that the optimal value of φ decreased. In our model we artificially varied φ across a range of values. In contrast, in the

locust AL, ϕ is adaptively modulated by feedback inhibition and decreases with increasing concentration. Thus, the system seems wired to naturally move towards this optimum.

We hypothesized that changes in synaptic weights caused by experience-dependent plasticity could degrade what had been an optimal representation. However, our simulations show that, despite STDP-induced changes to the strength of the KC–MBON synapse, the value of optimal sparseness was maintained. This finding is particularly notable because the overall feedforward architecture of the insect olfactory system, featuring a near absence of feedback across layers, implies that downstream layers cannot 'error-correct' upstream representations. Thus, the connectivity between layers must assure that an optimal representation constructed in one layer continues to be optimal from the perspective of subsequent layers. Our model allowed us to explore the mechanism underlying the maintenance of optimal sparseness across circuit layers and despite neural plasticity.

In the locust olfactory system, odor representations are parceled into cyclic 50ms packets of information, a process that begins in the antennal lobe and cascades at least two synapses forward to the MBONs. This parcellation is maintained and stabilized against noise and other corruption by STDP that adjusts the strength of synapses when pre–synaptic input leads or lags post–synaptic output by tens of milliseconds (within an oscillatory cycle). Lateral inhibition across MBONs may further sharpen the odor representation.

We show that odor representations are optimally separated despite STDP dependent weight changes to the KC-MBON synapse. However, the goal of MBONs is not solely to separate odor representations. Presynaptic circuits already seem wired to achieve this goal. KC-MBON synapses are a locus for associative conditioning in insects [45] where octopamine mediates appetitive while dopamine mediates aversive conditioning. How are activity patterns (KC spike patterns that evoke an MBON response) associated with a reward signal? A causal interaction between a KC and an MBON can be reinforced or degraded by STDP that increases or decreases the conductance of the synapse. However, these dynamic weight changes alone cannot associate neural responses to a reward because the reward signal arrives long after the activity pattern to be rewarded has subsided. Associative conditioning requires two elements to be in order. First, the system must reliably encode the odor. Second, the reward signal must be paired with the right pattern of activity. The intervening time between the pattern to be rewarded and the reward itself is likely corrupted with random spikes. Theoretical studies posit that STDP evokes a synaptic "tag" that decays on a slow time scale and persists when a diffuse neuromodulatory signal like Octopamine or Dopamine is initiated by a reward [46]. The reward signal affects only those synapses that were potentiated by STDP and continue to exhibit traces of the "tag".

Thus, STDP plays two roles. One, it homeostatically regulates the MBON response. In the absence of such regulation, the density of MBON spikes would increase, potentially obscuring the differences between odor representations. A consequence of maintaining this homeostasis is, optimal odor representations remain optimal despite activity-dependent modulation of synaptic weights. Our paper examines this particular aspect of STDP in the absence of a reward signal. The second role of STDP becomes evident only when a reward is present. A sparse set of synapses are "tagged" by STDP, and the functional form of STDP for those synapses is modified. Associative conditioning is encoded as changes in a sparse set of KC-MBON synaptic weights. Therefore, though the sole purpose of MBONs is not pattern separation, its activity must not obscure optimality arrived at by presynaptic circuits.

The KC–MBON junction may be the location where the imperative of insect olfactory system changes from identifying the odor to associating the odor with other sensory and reward inputs. If so, MBONs may not require a precise representation of the odor. In fact, in *Drosophila*, the MBONs have been shown to be broadly tuned, and thus instantiate a representation

more redundant than that of the population of narrowly tuned KCs. However, studies in locust have shown that the odor-elicited responses of MBONs, though densely spiking, are sensitive to the temporal ordering of KC input [35], and contain information about odor identity [47]. In our analysis we included the dynamics of MBONs throughout the odor stimulus, parsing its spike trains into 50 millisecond bins (equivalent to one oscillatory cycle in locusts) and calculating the ability of the system to discriminate odorants over the entire duration of the odor. Our study suggests that odor representations are maximally separated when the neural representation of the odor in the mushroom body is optimally sparse. Despite challenges, the olfactory circuit of insects maintains this optimal sparseness over variations in the concentration and experience dependent plasticity.

Methods

The model antennal lobe consisted of a scaled–down network of 350 PNs and 100 inhibitory interneurons (the locust antennal lobe contains roughly 830 PNs and 350 local neurons). Each neuron was modeled as a single compartment with voltage and calcium dependent currents with Hodgkin-Huxley kinetics. PNs generated Na^+ spikes while inhibitory interneurons generated Ca^{2+} spikelets, as seen in the locust olfactory system [47]. The model of inhibitory interneurons included a Ca^{2+} current (I_{Ca}) and a Ca^{2+} dependent K current that caused spike rate adaptation. Model PNs included a fast sodium current I_{Ca}, a fast potassium current I_K [48], a transient potassium A-current I_A [49], and a potassium leak current I_{KL}. The equations governing the dynamics of the neurons are as follows,

$$C_m \frac{dV_{PN}}{dt} = -g_L(V_{PN} - E_L) - I_{Na} - I_K - I_A - g_{KL}(V_{PN} - E_{KL}) - I_{GABA_A} - I_{nACh} - I_{ext1} \quad (1.1)$$

$$C_m \frac{dV_{LN}}{dt} = -g_L(V_{LN} - E_L) - I_{Ca} - I_{K(Ca)} - I_K - g_{KL}(V_{LN} - E_{KL}) - I_{GABA_A} - I_{nACh} - I_{ext2} \quad (1.2)$$

The passive parameters of the model were set as follows. $C_M = 1.43 \times 10^{-4} \mu S$, $g_L = 0.15 \mu S$ and $g_{KL} = 0.05 \mu S$. $E_L = -55 mV$ and $E_{KL} = -95 mV$. The passive parameters were set the same for both the PNs (subscript PN in all the equations) and the inhibitory local interneurons (subscript LN in the equations). The intrinsic currents governing the dynamics of each neuron is given below.

Sodium current I_{Na} is given by,

$$I_{Na} = g_{Na} m^3 h (V - E_{Na}) \quad (1.3)$$

where, the Na conductance, $g_{Na} = 50 \mu S$ and the reversal potential, $E_{Na} = 50 mV$. m and h are the activation and inactivation variables that are given by,

$$\frac{dm}{dt} = -\frac{1}{\tau_m}(m - m_\infty(V_{PN})) \quad (1.4)$$

$$\frac{dh}{dt} = -\frac{1}{\tau_h}(h - h_\infty(V_{PN})) \quad (1.5)$$

where, $m_\infty = \frac{\alpha_1}{\alpha_1 + \beta_1}$ and $\tau_m = \left(\frac{1}{\alpha_1 + \beta_1}\right)$ with $\alpha_1 = \frac{-0.32(13 - V_{PN})}{\exp\left(\frac{13 - V_{PN}}{4}\right) - 1}$ and $\beta_1 = \frac{0.28(V_{PN} + 15)}{\exp\left(\frac{V_{PN} + 15}{5}\right) - 1}$. The steady state values of the inactivation variable h and the time constant τ_h are given by $h_\infty = \frac{\alpha_2}{\alpha_2 + \beta_2}$ and $\tau_h = \left(\frac{1}{\alpha_2 + \beta_2}\right)$, where, $\alpha_2 = -0.128\left(\frac{38 + V_{PN}}{18}\right)$ and $\beta_2 = \frac{4}{\exp\left(-\frac{15 + V_{PN}}{5}\right) + 1}$

The equations describing the potassium current I_K for both PNs and inhibitory interneurons are as follows,

$$I_k = g_K n^4 (V - E_K) \qquad (1.6)$$

where, $g_K = 10$ and $E_K = -95$. The activation variable of the K current is given by,

$$\frac{dn}{dt} = -\frac{1}{\tau_n}(n - n_\infty) \qquad (1.7)$$

where, $n_\infty = \frac{\alpha_3}{\alpha_3 + \beta_3}$ and $\tau_n = \frac{1}{\alpha_3 + \beta_3}$. $\alpha_3 = -0\frac{0.02(30+V)}{\exp\left(-\frac{30+V}{5}\right)-1}$ and $\beta_3 = 0.5\exp\left(-\frac{35+V}{40}\right)$

A transient potassium current, I_A, in PNs was described by the following equation,

$$I_A = g_A m^4 h(V_{PN} - E_A) \qquad (1.8)$$

where, $g_A = 10\mu S$. The steady state values of the activation and the inactivation variables are given by $m_\infty = \frac{1}{1+\exp\left(-\frac{V_{PN}+60}{8.5}\right)}$ and $h_\infty = \frac{1}{1+\exp\left(-\frac{V_{PN}+78}{6}\right)}$. The time constants are given by $\tau_m = \left[\exp\left(\frac{V+35.8}{19.7}\right) + \exp\left(-\frac{V+79.7}{12.7}\right)\right]^{-1}$ and $\tau_h = \left[\exp\left(\frac{V+46}{5}\right) + \exp\left(-\frac{V+238.4}{37.45}\right)\right]^{-1}$. The inhibitory interneurons showed spike frequency adaptation due to a calcium dependent potassium current. The equations and parameter values are as follows,

$$I_{K(Ca)} = g_{K(Ca)} m(V_{LN} - E_{K(Ca)}) \qquad (1.9)$$

where, $g_{K(Ca)} = 0.3\mu S$, and $E_{K(Ca)} = -90mV$. The steady state value of the activation variable, m, is $m_\infty = \frac{[Ca_{2+}]}{[Ca_{2+}]+2}$ and the time constant is $\tau_m = \frac{100}{[Ca_{2+}]+2}$. The calcium concentration, $[Ca_{2+}]$, dynamics is governed by the following equation,

$$\frac{d[Ca_{2+}]}{dt} = -AI_{Ca} - \frac{[Ca_{2+}] - [Ca_{2+}]_\infty}{\tau} \qquad (1.10)$$

where the equilibrium concentration of Calcium, $[Ca_{2+}]_\infty = 2.4\times10^{-4}mM$, and the time constant $\tau = 5ms$. The constant $A = 5.2 \times 10^{-4} \frac{mMcm^2}{ms\mu A}$. The calcium current in the inhibitory neurons is given by,

$$I_{Ca} = g_{Ca} m^2 h(V - E_{Ca}) \qquad (1.11)$$

where, $g_{Ca} = 2\mu S$ and $E_{Ca} = 140mV$. $m_\infty = \frac{1}{1+\exp\left(-\frac{V_{LN}+20}{6.5}\right)}$ and $h_\infty = \frac{1}{1+\exp\left(\frac{V_{LN}+25}{12}\right)}$. The time constants are $\tau_m = 1.5$ and $\tau_h = 0.3\exp\left(\frac{V-40}{13}\right) + 0.002\exp\left(-\frac{V-60}{29}\right)$

Fast GABAergic synapses between interneurons and between PNs and inhibitory interneurons were modelled using first order activation schemes. Similarly, nicotinic cholinergic input from PNs was used to drive the inhibitory interneurons. 50 of the 350 PNs extended excitatory input to the other PNs. All other PNs did not extend direct connections to each other. GABAergic and cholinergic synapses were both described by the following equations,

$$I_{syn} = g_{syn}[O](V - E_{syn}) \qquad (1.12)$$

where the reversal potential is $E_{nAch} = 0mV$ for cholinergic receptors and $E_{GABA_A} = -70mV$ for fast GABA receptors. $[O]$ is the fraction of open channels that is calculated according to,

$$\frac{d[O]}{dt} = \alpha(1 - [O])[T] - \beta[O] \qquad (1.13)$$

The rate constants, $\alpha = 10ms^{-1}$ and $\beta = 0.16ms^{-1}$ for GABAergic synapses and $\alpha = 10ms^{-1}$ and $\beta = 0.2ms^{-1}$ for cholinergic synapses. When the receptors are activated following a spike, the term $[T]$ becomes non-zero. For cholinergic neurons this was modelled as the product of Heaviside functions in the following form,

$$[T] = A\Theta(t_{max} - t_0 - t)\Theta(t - t_0) \tag{1.14}$$

where, t_0 is the time of receptor activation, $A = 0.5$ and $t_{max} = 0.3ms$.

For GABAergic synapses,

$$[T] = \frac{1}{1 + \exp\left(-\frac{V(t) - V_0}{\sigma}\right)} \tag{1.15}$$

Kenyon cells and MBONS

We modeled a large array (15000) of KCs and 100 MBONs. Given the large number of KCs, we modeled each as a two-dimensional map that can replicate in a computationally efficient way the dynamics of a variety of conductance-based neurons and networks of these neurons, but is computationally efficient [40,41]. KCs and MBONs were modeled as regular spiking cells governed by the following equations,

$$x_{n+1} = f(x_n, x_{n-1}, y_n + \beta_n) \tag{1.16}$$

where the function, $f(x_n, x_{n-1}, y_n + \beta_n)$ is defined as,

$$f(x_n, x_{n-1}, y_n + \beta_n) = \begin{aligned} &= \alpha/(1 - x_n) + u \text{ for } x_n \leq 0 \\ &= \alpha + u \text{ for } 0 < x_n < \alpha + u \\ &= -1 \text{ for } x_n \geq \alpha + u \text{ or } x_{n-1} > 0 \end{aligned} \tag{1.17}$$

where $u = y_n + \beta_n$ $\mu = 0.0005$. Both KCs and MBONs received feedforward excitatory input. We did not model any lateral inhibition in these layers. Synaptic input is described by the following equations,

$$\begin{aligned} I_{n+1}^{syn} &= \gamma I_{n+1}^{syn} + g_{syn}(x_n^{post} - x_{rp}) \\ &= 0 \text{ otherwise} \end{aligned} \tag{1.18}$$

The KC-MBON synapse showed spike timing dependent plasticity. We modeled STDP using an online update rule. Each pre-synaptic KC spike activated a variable that decayed exponentially post activation in the absence of other spikes. The dynamics followed the equation,

$$\frac{dx}{dt} = -x + a_+ \sum_f \delta(t - t^f) \tag{1.19}$$

where, t^f is the time at which a spike occurs. The effect of the spike on the weight of the KC-MBON synapse is given by the factor a_+. In the absence of a spike the variable x decays exponentially to zero. A similar synaptic trace was defined to respond to postsynaptic spikes given by the following equation,

$$\frac{dy}{dt} = -y + a_- \sum_f \delta(t - t^f) \tag{1.20}$$

The weight of the synapse evolved in response to the timing of the pre- and post-synaptic spikes according to the following equation,

$$\frac{dw}{dt} = A_+ x(t) \sum_n \delta(t - t^f) - A_- y(t) \sum_n \delta(t - t^f) \tag{1.21}$$

Further, the factors A_+ and A_- changed in a manner that depended on the current weight $w(t)$ of the synapse. Increases in weight when the synapse was close to a maximum weight were lower in magnitude than when it was further away from W_{max}. This was achieved by introducing soft bounds to the weight by setting $A_+ = (W_{max} - w)\eta_+$ and $A_- = w\eta_-$

Author Contributions

Conceptualization: Collins Assisi, Mark Stopfer, Maxim Bazhenov.

Data curation: Collins Assisi, Maxim Bazhenov.

Formal analysis: Collins Assisi.

Funding acquisition: Collins Assisi, Mark Stopfer, Maxim Bazhenov.

Investigation: Collins Assisi, Maxim Bazhenov.

Methodology: Collins Assisi, Maxim Bazhenov.

Project administration: Collins Assisi, Mark Stopfer.

Resources: Collins Assisi, Mark Stopfer, Maxim Bazhenov.

Software: Collins Assisi, Maxim Bazhenov.

Supervision: Collins Assisi, Mark Stopfer, Maxim Bazhenov.

Validation: Collins Assisi.

Visualization: Collins Assisi.

Writing – original draft: Collins Assisi.

Writing – review & editing: Collins Assisi, Mark Stopfer, Maxim Bazhenov.

References

1. Laurent G. Olfactory network dynamics and the coding of multidimensional signals. Nat Rev Neurosci. 2002; 3: 884–895. https://doi.org/10.1038/nrn964 PMID: 12415296

2. Friedrich RW, Habermann CJ, Laurent G. Multiplexing using synchrony in the zebrafish olfactory bulb. Nat Neurosci. 2004; 7: 862–871. https://doi.org/10.1038/nn1292 PMID: 15273692

3. Wiechert MT, Judkewitz B, Riecke H, Friedrich RW. Mechanisms of pattern decorrelation by recurrent neuronal circuits. Nat Neurosci. 2010; 13: 1003–1010. https://doi.org/10.1038/nn.2591 PMID: 20581841

4. Hige T, Aso Y, Rubin GM, Turner GC. Plasticity-driven individualization of olfactory coding in mushroom body output neurons. Nature. 2015; 526: 258–262. https://doi.org/10.1038/nature15396 PMID: 26416731

5. Cassenaer S, Laurent G. Hebbian STDP in mushroom bodies facilitates the synchronous flow of olfactory information in locusts. Nature. 2007; 448: 709–713. https://doi.org/10.1038/nature05973 PMID: 17581587

6. Hallem EA, Carlson JR. Coding of Odors by a Receptor Repertoire. Cell. 2006; 125: 143–160. https://doi.org/10.1016/j.cell.2006.01.050 PMID: 16615896

7. Kreher SA, Kwon JY, Carlson JR. The molecular basis of odor coding in the Drosophila larva. Neuron. 2005; 46: 445–456. https://doi.org/10.1016/j.neuron.2005.04.007 PMID: 15882644

8. Raman B, Joseph J, Tang J, Stopfer M. Temporally Diverse Firing Patterns in Olfactory Receptor Neurons Underlie Spatiotemporal Neural Codes for Odors. J Neurosci. 2010; 30: 1994–2006. https://doi.org/10.1523/JNEUROSCI.5639-09.2010 PMID: 20147528

9. Laurent G, Wehr M, Davidowitz H. Temporal representations of odors in an olfactory network. J Neurosci. 1996; 16: 3837–3847. https://doi.org/10.1523/JNEUROSCI.16-12-03837.1996 PMID: 8656278

10. Wehr M, Laurent G. Odour encoding by temporal sequences of firing in oscillating neural assemblies. Nature. 1996; 384: 162–166. https://doi.org/10.1038/384162a0 PMID: 8906790

11. Mazor O, Laurent G. Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. Neuron. 2005; 48: 661–673. https://doi.org/10.1016/j.neuron.2005.09.032 PMID: 16301181

12. Perez-Orive J, Mazor O, Turner GC, Cassenaer S, Wilson RI, Laurent G. Oscillations and sparsening of odor representations in the mushroom body. Science (80-). 2002; 297: 359–365.

13. Papadopoulou M, Cassenaer S, Nowotny T, Laurent G. Normalization for Sparse Encoding of Odors by a Wide-Field Interneuron. Science (80-). 2011; 332: 721–725.

14. Kanerva P. Sparse distributed memory and related models. Associative Neural Memories: Theory and Implementation. 1993. pp. 50–76.

15. Marr D. A theory of cerebellar cortex. J Physiol. 1969; 202: 437–470. https://doi.org/10.1113/jphysiol.1969.sp008820 PMID: 5784296

16. Albus JS. A theory of cerebellar function. Math Biosci. 1971; 10: 25–61.

17. Olshausen BA, Field DJ. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. Nature. 1996; 381: 607–609. https://doi.org/10.1038/381607a0 PMID: 8637596

18. Litwin-Kumar A, Harris KD, Axel R, Sompolinsky H, Abbott LF. Optimal Degrees of Synaptic Connectivity. Neuron. 2017; 93: 1153–1164.e7. https://doi.org/10.1016/j.neuron.2017.01.030 PMID: 28215558

19. Babadi B, Sompolinsky H. Sparseness and Expansion in Sensory Representations. Neuron. 2014; 83: 1213–1226. https://doi.org/10.1016/j.neuron.2014.07.035 PMID: 25155954

20. Aso Y, Hattori D, Yu Y, Johnston RM, Iyer NA, Ngo TTB, et al. The neuronal architecture of the mushroom body provides a logic for associative learning. Elife. 2014; 3: e04577. https://doi.org/10.7554/eLife.04577 PMID: 25535793

21. Strube-Bloss MF, Rössler W. Multimodal integration and stimulus categorization in putative mushroom body output neurons of the honeybee. R Soc Open Sci. 2018; 5.

22. Leitch B, Laurent G. GABAergic synapses in the antennal lobe and mushroom body of the locust olfactory system. J Comp Neurol. 1996; 372: 487–514. https://doi.org/10.1002/(SICI)1096-9861(19960902)372:4<487::AID-CNE1>3.0.CO;2-0 PMID: 8876449

23. Stopfer M, Jayaraman V, Laurent G. Intensity versus Identity Coding in an Olfactory System. Neuron. 2003; 39: 991–1004. https://doi.org/10.1016/j.neuron.2003.08.011 PMID: 12971898

24. Jortner RA. Network architecture underlying maximal separation of neuronal representations. Front Neuroeng. 2013; 5. https://doi.org/10.3389/fneng.2012.00019 PMID: 23316159

25. Jortner RA, Farivar SS, Laurent G. A simple connectivity scheme for sparse coding in an olfactory system. J Neurosci. 2007; 27: 1659–1669. https://doi.org/10.1523/JNEUROSCI.4171-06.2007 PMID: 17301174

26. Assisi C, Stopfer M, Bazhenov M. Using the structure of inhibitory networks to unravel mechanisms of spatiotemporal patterning. Neuron. 2011; 69: 373–386. https://doi.org/10.1016/j.neuron.2010.12.019 PMID: 21262473

27. Assisi C, Bazhenov M. Synaptic inhibition controls transient oscillatory synchronization in a model of the insect olfactory system. Front Neuroeng. 2012; 5: 7. https://doi.org/10.3389/fneng.2012.00007 PMID: 22529800

28. Bazhenov M, Stopfer M, Rabinovich M, Huerta R, Abarbanel HD, Sejnowski TJ, et al. Model of transient oscillatory synchronization in the locust antennal lobe. Neuron. 2001; 30: 553–567. https://doi.org/10.1016/s0896-6273(01)00284-7 PMID: 11395014

29. Assisi C, Stopfer M, Laurent G, Bazhenov M. Adaptive regulation of sparseness by feedforward inhibition. Nat Neurosci. 2007; 10: 1176–1184. https://doi.org/10.1038/nn1947 PMID: 17660812

30. Bazhenov M, Stopfer M, Rabinovich M, Huerta R, Abarbanel HDI, Sejnowski TJ, et al. Model of transient oscillatory synchronization in the locust antennal lobe. Neuron. 2001; 30: 553–567. https://doi.org/10.1016/s0896-6273(01)00284-7 PMID: 11395014

31. Stopfer M, Bhagavan S, Smith BH, Laurent G. Impaired odour discrimination on desynchronization of odour-encoding neural assemblies. Nature. 1997; 390: 70–74. https://doi.org/10.1038/36335 PMID: 9363891

32. Stopfer M, Laurent G. Short-term memory in olfactory network dynamics. Nature. 1999; 402: 664–668. https://doi.org/10.1038/45244 PMID: 10604472

33. Gupta N, Stopfer M. Functional Analysis of a Higher Olfactory Center, the Lateral Horn. J Neurosci. 2012; 32: 8138–8148. https://doi.org/10.1523/JNEUROSCI.1066-12.2012 PMID: 22699895

34. Kee T, Sanda P, Gupta N, Stopfer M, Bazhenov M. Feed-Forward versus Feedback Inhibition in a Basic Olfactory Circuit. PLoS Comput Biol. 2015; 11: 1–24.

35. Gupta N, Stopfer M. A temporal channel for information in sparse sensory coding. Curr Biol. 2014; 24: 2247–2256. https://doi.org/10.1016/j.cub.2014.08.021 PMID: 25264257

36. Assisi C, Stopfer M, Bazhenov M. Excitatory local interneurons enhance tuning of sensory information. PLoS Comput Biol. 2012; 8: e1002563. https://doi.org/10.1371/journal.pcbi.1002563 PMID: 22807661

37. Gupta N, Stopfer M. Functional Analysis of a Higher Olfactory Center, the Lateral Horn. J Neurosci. 2012; 32: 8138–8148. https://doi.org/10.1523/JNEUROSCI.1066-12.2012 PMID: 22699895

38. Gerstner W, Kempter R, Van Hemmen JL, Wagner H. A neuronal learning rule for sub-millisecond temporal coding. Nature. 1996; 383: 76–78. https://doi.org/10.1038/383076a0 PMID: 8779718

39. Sjöström J, Gerstner W, Markram H. Spike-timing dependent plasticity. Scholarpedia. 2010.

40. Rulkov NF. Modeling of Spiking-Bursting Neural Behavior Using Two-Dimensional Map. Phys Rev E Stat Nonlin Soft Matter Phys. 2002; 65: 41922.

41. Rulkov NF, Bazhenov M. Oscillations and synchrony in large-scale models of cortical network. Biophys J. 2008; 34: 279–299.

42. Kistler WM, Hemmen JL Van. Modeling Synaptic Plasticity in Conjunction with the Timing of Pre- and Postsynaptic Action Potentials. 2000; 405: 385–405.

43. Liu X, Davis RL. The GABAergic anterior paired lateral neuron suppresses and is suppressed by olfactory learning. Nat Neurosci. 2009; 12: 53–59. https://doi.org/10.1038/nn.2235 PMID: 19043409

44. Lin AC, Bygrave AM, Calignon A De, Lee T, Miesenböck G. Sparse, decorrelated odor coding in the mushroom body enhances learned odor discrimination. 2014; 17.

45. Cassenaer S, Laurent G. Conditional modulation of spike-timing-dependent plasticity for olfactory learning. Nature.

46. Izhikevich EM. Solving the distal reward problem through linkage of STDP and dopamine signaling. Cereb Cortex. 2007; 17: 2443–2452. https://doi.org/10.1093/cercor/bhl152 PMID: 17220510

47. MacLeod K, Bäcker A, Laurent G. Who reads temporal information contained across synchronized and oscillatory spike trains? Nature. 1998; 395: 693–698. https://doi.org/10.1038/27201 PMID: 9790189

48. Traub RD, Wong RK, Miles R, Michelson H. A model of a CA3 hippocampal pyramidal neuron incorporating voltage-clamp data on intrinsic conductances. J Neurophysiol. 1991; 66: 635–650. https://doi.org/10.1152/jn.1991.66.2.635 PMID: 1663538

49. Huguenard JR, Coulter DA, Prince DA. A fast transient potassium current in thalamic relay neurons: kinetics of activation and inactivation. J Neurophysiol. 1991; 66: 1304–1315. https://doi.org/10.1152/jn.1991.66.4.1304 PMID: 1662262