Plasticity in inhibitory networks improves pattern separation in early olfactory processing

Shruti Joshi[1,2*], Seth Haney[2], Zhenyu Wang[3], Fernando Locatelli[5], Hong Lei[4], Yu Cao[6], Brian Smith[4], Maxim Bazhenov[2*]

[1]Department of Electrical and Computer Engineering, University of California San Diego, USA.
[2]Department of Medicine, University of California San Diego, USA.
[3]Department of Electrical, Computer and Energy Engineering, Arizona State University, USA.
[4]School of Life Science, Arizona State University, USA.
[5]Facultad de Ciencias Exactas y Naturales, Universidad de Buenos Aires, Instituto de Fisiología, Biología Molecular y Neurociencias, CONICET, Buenos Aires, Argentina.
[6]Department of Electrical and Computer Engineering, University of Minnesota, USA.

*Corresponding author(s). E-mail(s): s4joshi@ucsd.edu; mbazhenov@health.ucsd.edu;
Contributing authors: sethdhaney@gmail.com; zwang586@asu.edu; fflocatelli@gmail.com; hong.lei@asu.edu; yucao@umn.edu; brianhsmith@asu.edu;

Abstract

Distinguishing between nectar and non-nectar odors is challenging for animals due to shared compounds and varying ratios in complex mixtures. Changes in nectar production throughout the day - and potentially many times within a forager's lifetime - add to the complexity. The honeybee olfactory system, containing fewer than 1,000 principal neurons in the early olfactory relay, the antennal lobe (AL), must learn to associate diverse volatile blends with rewards. Previous studies identified plasticity in the AL circuits, but its role in odor learning remains poorly understood. Using a biophysical computational network model, tuned by in vivo electrophysiological data, and live imaging of the honeybee's AL, we explored the neural mechanisms and functions of plasticity in the early olfactory system. Our findings revealed that when trained with a set of rewarded and unrewarded odors, the AL inhibitory network suppresses shared chemical compounds while enhancing responses to distinct compounds. This results in improved pattern separation and a more concise neural code. Our calcium imaging data support these predictions. Analysis of a graph convolutional neural network performing an odor categorization task revealed a similar mechanism for contrast enhancement. Our study provides insights into how inhibitory plasticity in the early olfactory network reshapes the coding for efficient learning of complex odors.

Significance Statement

By combining biophysical modeling, machine learning, electrophysiology and analysis of calcium imaging data, we demonstrate that associative and nonassociative plasticity in the honeybee antennal lobe (AL) - the first relay of the insect olfactory system - work together to enhance the contrast between rewarded and unrewarded odors. Training the AL's inhibitory network within specific odor environments enables the suppression of neural responses to common odor components while amplifying responses to distinctive ones. This study sheds light on the olfactory system's ability to adapt and efficiently learn new odor-reward associations across varying environments, and it proposes innovative, energy-efficient principles applicable to artificial intelligence.

1 Introduction

A honeybee can locate nectar-producing flowers by detecting floral aromas composed of many volatile compounds (Pichersky and Gershenzon, 2002; Negre et al., 2003). However, nectar-producing and non-producing flowers contain many of the same compounds, making it difficult for the honeybee to determine the map between odor sensing and reward prediction (Smith et al., 2006). This task is further complicated by the fact that nectar production may change within days and across locations, potentially many times within the foraging lifetime of a honey bee. Although the olfactory coding space is large (close to 1000 projection neurons in the honeybee antennal lobe, (Paoli and Galizia, 2021; Galizia, 2014)), mapping the sensory environment requires learning and relearning the association of reward with variable blends of volatile compounds. Nevertheless, honeybees can adapt to the variances in their environment owing to their keen ability to discriminate between a wide range of olfactory stimuli (Menzel, 1990; Farooqui et al., 2003; Fernandez et al., 2009; Locatelli et al., 2013).

Different forms of plasticity are a ubiquitous feature of the early olfactory processing in the brains of both mammals and insects. In the neural networks of the mammalian olfactory bulb (OB) and insect antennal lobe (AL), both nonassociative (unsupervised) and associative (supervised) forms of plasticity have been described (Linster et al., 2009; Wilson and Linster, 2008; Fernandez et al., 2009; Locatelli et al., 2016). Together with the now well-established similarities in anatomical connectivity within the AL and OB networks (Hildebrand and Shepherd, 1997; Fulton et al., 2024), it is clear that early olfactory processing in these phylogenetically very different groups of animals works in much the same way to augment olfactory processing.

Studies across different insect species have revealed conserved mechanisms of plasticity in early olfactory processing. In Drosophila, prolonged odor exposure leads to selective potentiation of inhibitory synapses between local interneurons (LNs) and projection neurons (PNs), resulting in habituation (Das et al., 2011). In moths, associative conditioning modifies both excitatory and inhibitory connections within the AL network, enhancing responses to behaviorally relevant odors (Daly et al., 2004). Work in locusts has demonstrated that plasticity in the antennal lobe can shape representations to improve stimulus

discrimination (Stopfer and Laurent, 1999; Bazhenov et al., 2005). In honeybees, both associative and non-associative forms of plasticity have been shown to modify AL network dynamics, with computational models suggesting that such modifications can enhance discrimination between similar odors (MaBouDi et al., 2017; Marachlian et al., 2021; Locatelli et al., 2013, 2016; Chen et al., 2015). Together, these studies highlight inhibitory synaptic plasticity as a conserved mechanism for optimizing early olfactory processing across insect species.

In this new study, we combined Ca^{2+} imaging from the honeybee AL with a biophysically realistic computational modeling to characterize the role of inhibitory synaptic plasticity in the AL in improving odor discrimination after olfactory learning. We found that representation of chemical compounds shared between rewarded and habituated odors are suppressed while distinct ones are enhanced after learning to increase contrast and reduce overlap between odor representations. This change in representation requires both associative and nonassociative plasticity on the inhibitory synapses between AL neurons. Analysis of the Ca^{2+} imaging data revealed a change in odor representations after learning, supporting the model's prediction. Applying these ideas to the Artificial Neural Network (ANN) models trained to perform the odor categorization task revealed principles of contrast enhancement applicable to machine learning. In sum, we demonstrate the role of inhibitory synaptic plasticity in the AL for effective odor discrimination and propose a novel computationally efficient mechanism for performing categorization tasks.

2 Results

Differential conditioning created distinct representations of odors *in vivo*

In our previous work (Locatelli et al., 2016) we used Ca^{2+} imaging to test how the representation of synthetic odor blends changes in the honeybee AL after differential conditioning (Figure 1A). Two synthetic classes of odors were based on varieties of the common snapdragon flower (*A. majus*): Potomac Pink (PP) and Pale Hybrid (PH) (Wright et al., 2005). Each class was a blend of volatile chemicals that recapitulated mean and variance of naturally occurring PH and PP flowers. These two classes contained the same volatile chemicals at different ratios that makes distinguishing them a difficult task.

We found that odors within each class created distinct representations in the AL via activation of different sets of glomeruli (Figure 1B,C). Further, these representations became more distinct after differential training. Specifically, in absolute conditioning, honeybees were rewarded after responding to PH odors; in differential conditioning honeybees were also habituated to the PP odors (Locatelli et al., 2016) (Figure 1E). We found that the correlation between the representations of PH and PP odors decreased after differential conditioning (Figure 1D).

Computational model of the mechanisms of differential conditioning

Here, we sought to investigate the mechanisms that produce differential learning between complex odors in the honeybee AL using a computational model. The model was a network of Hodgkin-Huxley type excitatory principal neurons (PNs) and inhibitory local neurons (LNs), representing honeybee AL network, and it was adapted from our previous work (Chen et al., 2015) (see Methods, Figure 2A). This model was previously tuned to capture characteristic features of the AL dynamics, e.g., odor triggered LFP oscillations (Stopfer et al., 1997), and thus it enables us to study mechanisms of plasticity without compromising realistic AL dynamics which may play a role in shaping representations.

To model reward learning - appetitive conditioning - we simulated activity-dependent presynaptic facilitation at both the LN-PN and LN-LN synapses (see Methods) inspired by the finding that octopamine receptors activated by reward signal are localized on the inhibitory LNs in the AL (Sinakevitch et al., 2013). Octopamine is needed for appetitive olfactory learning (Hammer, 1993; Farooqui et al., 2003) and it has been shown to regulate inhibitory connections within the AL, affecting odor representations (Rein et al., 2013). Here we focus on AmOA1 octopamine receptor (Sinakevitch et al., 2011) expressed alongside GABA receptors in the honeybee AL (Sinakevitch et al., 2013). In honeybees (Grohmann et al., 2003) and fruit flies (Lee et al., 2003), octopamine binding to this receptor leads to an increase in intracellular Ca^{2+} that primes adenylyl cyclase to enhance cAMP and cAMP-dependent protein kinase activation that further regulates synaptic plasticity (Capogna et al., 1995; Han et al., 1998; Müller, 2002; Antonov et al., 2003). Therefore, we simulated the appetitive conditioning by increasing the synaptic weights of LN-PN and LN-LN synapses as the presynaptic LN was activated by the conditioned odor. Habituation was modeled as post-synaptic facilitation of LN-PN and LN-LN synapses, i.e., it depended on activity of

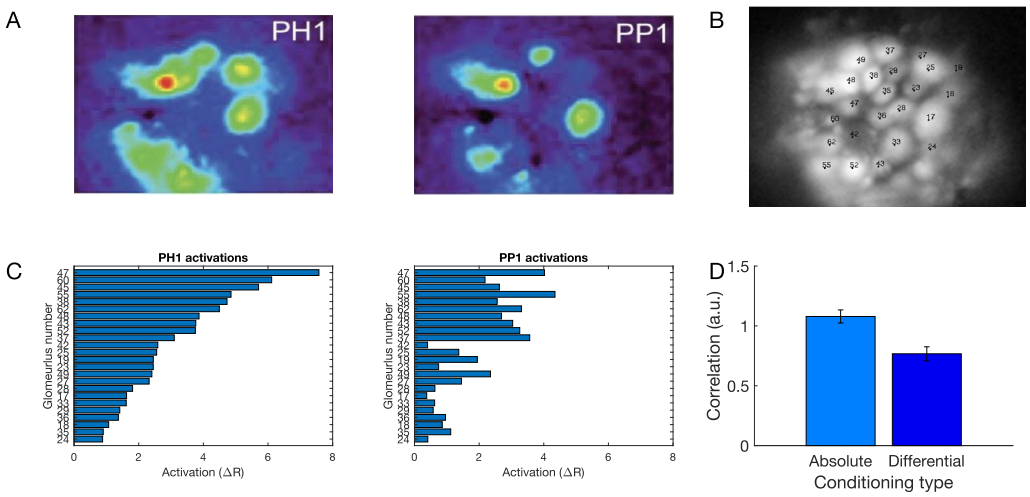

E

Conditioning	Training	Testing/Imaging (10 trials)
Differential Conditioning	PH_3+ I PP_2- I PP_4- … PP_3- I PH_3+ I PP_3-	$PH_{1\text{-}6}$ I $PP_{1\text{-}6}$
Absolute Conditioning	PH_2+ I MO- I PH_4+ … PH_3+ I MO- I PH_2+	$PH_{1\text{-}6}$ I $PP_{1\text{-}6}$

Fig. 1 Olfactory representations in the honeybee AL. A) Calcium imaging responses elicited by two different odors (PH1 and PP1) in a representative bee. Here, PH1 and PP1 belong to different classes of odors, with PH odors being presented along with a reward (rewarded odors) and PP odors presented without any reward (habituated odors). B) Position of the glomeruli within the antennal lobe, for which calcium imaging responses were measured. C) Mean activation across the different glomeruli elicited by the PH1 and PP1 odors. Glomeruli are ordered such that the glomeruli most activated by PH1 at the top and the glomeruli least activated by PH1 at the bottom. The same ordering is used for PP1 odor as well to highlight the difference in the activation patterns between the two odors. D) Correlation of glomeruli activation between different odor classes (PH and PP) decreases after differential conditioning compared to absolute conditioning (Wilcoxon signed rank test, p=4.8e-8, n=7 bees with absolute conditioning and 10 bees with differential conditioning). This indicates that the glomerular representation gets more separable after differential conditioning. The error bars here show standard error of mean. E) Honeybee training protocol for differential and absolute conditioning. In differential conditioning, the bee was trained with rewarded (PH) and habituated (PP) odors presented in a randomized order. For absolute conditioning, the bee was trained with only rewarded odors. Mineral Oil (MO) was presented to the bee without any reward during absolute conditioning. Each trial lasts for 4s, with 500ms of odor presentation (Data from (Locatelli et al., 2016)). The same training protocol was used for training the computational model.

postsynaptic LNs or PNs, based on the finding that habituation causes LN-PN facilitation in the fruit fly AL (Sudhakaran et al., 2012; Larkin et al., 2010).

The spiking activity of a representative PN and LN are shown in Figure 2C, D. Figure 2E shows the local field potential (LFP) obtained by averaging membrane voltages across the PN population and Figure 2F shows temporal structure of the input current pulse to the neurons during odor stimulation. The amplitude of this current pulse to each neuron depends on the odor currently being presented.

To account for the fact that natural odors are composed of a blend of distinct volatile chemicals, we created inputs to the model combining activation of multiple distinct odor 'percepts'. A percept can be thought of as AL activation induced by specific chemical component of the odor (see (Locatelli et al., 2016) and Figure 5A below). Each percept in the model was represented by a distinct population of "virtual" ORNs making connections to simulated LNs and PNs. This effectively divided all the LNs and PNs into several percept groups (Figure 2B). A percept in the model could include more than one glomerulus consistent with an observation that a chemical can activate more than one glomerulus, each at a different intensity, in vivo (Paoli and Galizia, 2021). As with real floral odors (Locatelli et al., 2016), model odors were composed of multiple overlapping percepts creating a difficult discrimination task.

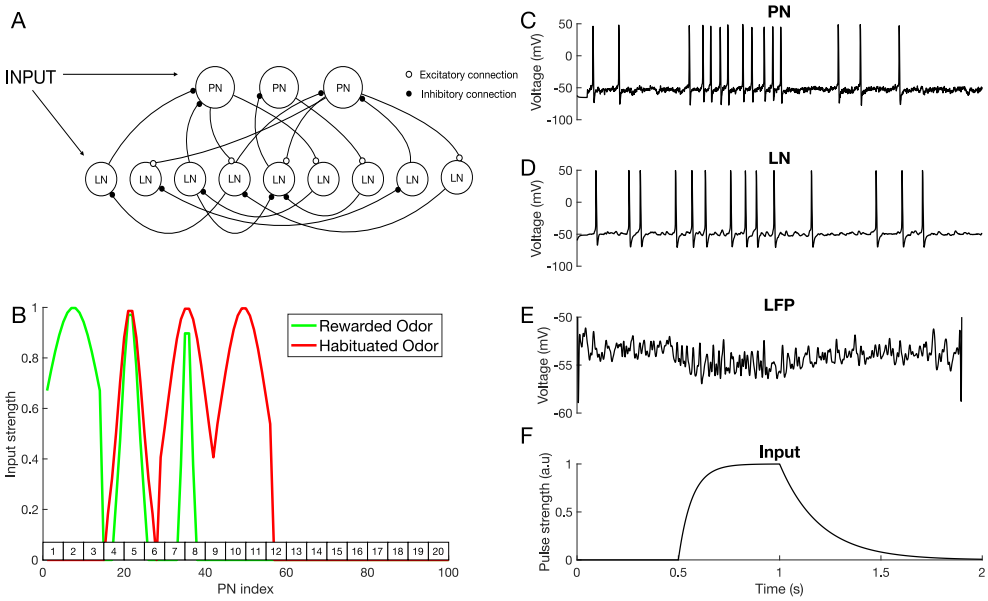

Fig. 2 Computational model of the antennal lobe. A) Architecture of the AL network model, including excitatory PNs and inhibitory LNs (see Methods). B) Spatial pattern of odor stimuli for rewarded and habituated odors. The PNs are split into 7 groups, each representing one odor percept: Percept 1 = [PN1 – PN14]; Percept 2 = [PN15 – PN28]; . . . . Each simulated odor activates three selected percepts. In this experiment, percepts 1, 2 and 3 are activated by the rewarded odor (green), while percepts 2, 3 and 4 are activated by the habituated odor (red). The rectangles on top of the X-axis show the glomerulus number, which are used to set up anatomical connectivity between LNs and PNs (see Methods). C-D) Responses of a representative PN (C) and LN (D) to an odor input in the untrained network. E) Averaged PN activity (Local field potential). F) The 500ms time modulated current pulse which was provided as an inputs to individual neurons during odor stimulation (Gaussian noise was added on top of this input - see Methods).

While our model dynamics reasonably match *in vivo* data (see Supplementary Figure1), PN firing rates in the model are slightly lower than those observed *in vivo*, particularly during odor stimulation and post-odor response. It is important to note, however, that *in vivo* recordings are naturally biased towards more active neurons. Additionally, model PN responses exhibit somewhat simplified temporal patterns, which is related to the lack of access to the simultaneous ORN recordings. Therefore, we used simplified approximations of the odor input, similar to our previous work in locusts (Bazhenov et al., 2001b,a), moths (Ito et al., 2009), and honeybees (Chen et al., 2015).

Differential training places stimuli further apart in the coding space

We trained our model using differential conditioning, where one class of odors $A = (A1, A2, ...)$ was rewarded and another $B = (B1, B2, ...)$ was habituated. Within each class, different odors were simulated by changing the width of Gaussian for individual percepts (Figure 2B). The model network was exposed to a sequence of $N = 30$ odors total presented in a randomized order, for a total of 60s (each odor presentation lasted 2s). The plasticity rule was selected based on which type of odor (i.e., rewarded or habituated) was presented. We found that the training led to significant changes in the PNs' population response for the rewarded odors (Figure 3A). Specifically, neurons associated with percepts that were unique to the rewarded odors showed enhanced firing rates after training, whereas neurons associated with percepts in the overlap between rewarded and habituated odors were strongly suppressed. For habituated odors, plastic changes led to rather minor and non-specific reduction of the PN responses. These combined effects effectively placed stimuli further apart in the AL coding space enhancing discrimination between rewarded and habituated odors.

To evaluate the robustness of our findings, we conducted a sensitivity analysis of the key model parameters. Specifically, we systematically varied the following: fast and slow GABAergic and ACh synaptic conductances between LNs and PNs, input synaptic conductances to LNs and PNs, and synaptic plasticity parameters (facilitation rate and decay time constant). We found that the model's core behavior

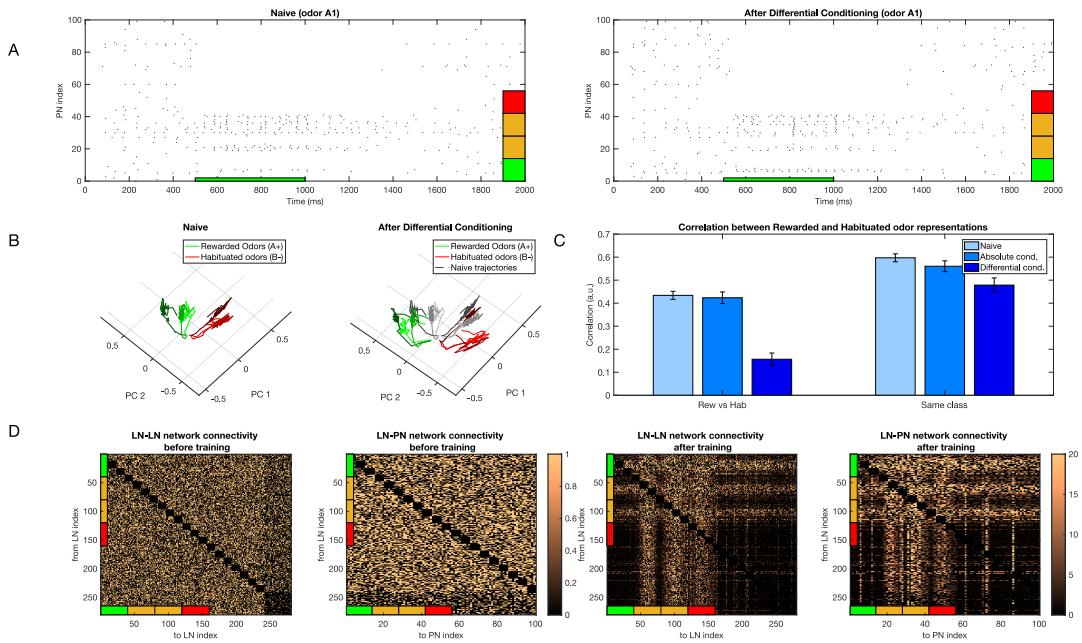

Fig. 3 Change in representation of simple odors after learning. A) Raster plot of the PN population response to rewarded odor (odor $A1$) before (left) and after (right) differential conditioning. Note, that activity of PNs activated by both rewarded and habituated odors (denoted by orange percepts) decreased while the activity of PNs activated uniquely by the rewarded odor (denoted by the green percept) increased. Red indicates percepts uniquely activated by the habituated odor. B) Dynamical trajectory of the spatiotemporal odor responses in two-dimensional (2D) PCA space. The green lines show trajectories for 3 Rewarded Odors and the red ones for 3 Habituated Odors. After training, the trajectories of the rewarded and habituated odors shift away from each other. Black lines on the right plot indicate naive trajectories. C) Left panel shows the correlation between representations of rewarded and habituated odors across three conditions: naive networks (light blue), after absolute conditioning (medium blue), and after differential conditioning (dark blue). While the correlation between odor representations from different classes remains largely unchanged after absolute conditioning compared to the naive state (p=0.32), it decreases significantly following differential conditioning (p=6e-7). The right panel illustrates correlations between representations of odors within the same class. These within-class correlations show a modest decrease after both absolute conditioning (p=0.1) and differential conditioning (p=0.07) relative to the naive condition. For both panels, correlations were computed pairwise between odors and averaged across all pairs, with error bars representing the standard deviation across trials.
D) Connectivity of the inhibitory networks (LN-LN and LN-PN) before (left) and after (right) training. Connectivity after differential conditioning shows a grid-like structure which encodes the relationship between percepts belonging to the rewarded and habituated odors. The colorbars show the final value of the weight divided by its initial value.

- enhanced separation between rewarded and habituated odor representations after learning - remains qualitatively consistent across a broad range of parameter values (Supplementary Figure 2).

We visualized effect of training using principal component analysis (PCA) (Figure 3B). We found that the trajectories of the naive response showed a relatively even distribution across PCA space. In contrast, after training, the sets of trajectories representing rewarded and habituated odors were strongly separated.

We found a significant decrease in correlation after differential training when comparing rewarded vs. habituated odors and a modest decrease when comparing odors within the same class, i.e., just rewarded or just habituated odors (Figure 3C). Note that odors from the same class are inherently more similar to each other than odors from different classes, even in the naive condition. Both these results are in agreement with experimental data (Locatelli et al., 2016). The changes to odor representation in our model were due to changes in the inhibitory network, as these are the only synapses modified in the model. Analysis of the connectivity matrix after learning revealed a characteristic structure reflecting the structure of percepts in rewarded and habituated odors (Figure 3D).

Next, we individually tested the effect of associative vs nonassociative plasticity on the ability of the trained network to distinguish between the rewarded and habituated odors by selectively applying only pre or post-synaptic plasticity and measuring the correlation between the learned odor representations. In Figure 4A, we show the correlation between rewarded and habituated representations for different

conditions - (i) Naive (pre-synaptic and post-synaptic plasticity are both off), (ii) Associative only (pre-synaptic plasticity is on, post-synaptic plasticity is off; absolute conditioning), (iii) Non-Associative only (pre-synaptic plasticity is off, post-synaptic plasticity is on), and (iv) Both (pre-synaptic and post-synaptic plasticity are both on; differential conditioning). We found that a significant decrease in correlation is seen when both associative and nonassociative plasticity are active, along with nonassociative only (Figure 4A, left) when compared with the naive case. The associative only case does not show a significant decrease in correlation when compared with naive. We also found a modest decrease in the correlation measured between representations of odors of the same class for all 3 conditions compared to naive (Figure 4A, right). The structure of LN-LN and LN-PN networks after learning for the associative only and nonassociative only cases is shown in Figure 4B and C, respectively. As expected, it reveals changes in only subset of connections (compare to Figure 3D). In Figure 4B and C, we can also see some bright vertical (Figure 4C) or horizontal (Figure 4B) lines showing increased inhibition to certain PNs which are not activated directly by the odor (e.g., vertical lines between PN 60 and 80 and 80 and 100 in Figure 4C, right panel). These lines appear because some PNs have a high background firing which seems to increase even more after training.

The training increased the total amount of inhibition in the network. To verify that the odor-specific structure of inhibitory connections, rather than the overall increase in inhibition, aids odor discrimination, we conducted a shuffling experiment. Specifically, we maintained an identical distribution of the weights in the shuffled networks but disrupted the structured connectivity patterns that emerged during learning. These shuffled networks, despite having an equivalent total inhibitory drive, failed to produce the between- and within-class odor decorrelation observed in the structured network (see Supplementary Figure 3). This supports our prediction that enhanced odor discrimination emerges from the precise restructuring of inhibitory circuits during learning, rather than from global increases in inhibition.

Training with realistic odorants also produces contrast enhancement

To ensure that the effects we have shown so far are not due to the specific construction of the model odors, we created a new set of odors where proportions of percepts were aligned with the proportions of the measured volatile chemicals present in each of the odor blend categories (i.e., PH1, PH2,.. and PP1, PP2,...) that were used in behavioral experiments in (Locatelli et al., 2016). These PH and PP odor classes were created by calculating the proportion of the chemical components present in the blends and setting the activation of each percept according to it (Figure 5A, see Methods for details). Training these odors using differential conditioning (PH was rewarded and PP was habituated) produced a separation of odor representations in the coding space (Figure 5C); both the correlations between odors from the different classes (i.e., rewarded and habituated) and between odors from the same class (e.g., PH1 vs PH2,..., PP1 vs PP2,...) decreased (Figure 5B), in line with our previous results obtained with simpler model (Figure 3B). This finding has been also seen in vivo (Locatelli et al., 2016). There, decrease in similarity of representations was found between different PH and PP odors, indicating that differential conditioning might make recognition of individual rewarded odors easier.

Our minimal model (Figure 3) predicts that neurons associated with percepts that are unique to the rewarded odors increase firing rates after training, whereas neurons associated with percepts in the overlap between rewarded and habituated odors reduce activity. To quantify this effect for more complex model (Figure 5), we calculated a uniqueness score for each percept by calculating the normalized difference between its activation by the rewarded odor and its activation by the habituated one (see Methods). Next, we calculated a change of each percept activity after training (differential conditioning). The result is plotted in Figure 5D (red dots and line). It revealed that percepts with higher uniqueness index increased activity after training (positive correlation), while those with low uniqueness index decreased, thus supporting our hypothesis.

To test the model's prediction, we sought to quantify this effect *in vivo*. The Ca^{2+} imaging data from honeybee has been published previously (Locatelli et al., 2016) and here we re-analyzed this data. Similar to the model analysis, we calculated a uniqueness score for each glomerulus. This was done by determining, for each glomerulus, the difference between its activation by the rewarded odor and its activation by the habituated one (see Methods). Here, we compare odor representations obtained after differential conditioning versus absolute conditioning, as experimental recordings for the naive case are not available. We found a strong correlation between the uniqueness score of a glomerulus and the change in its activity due to differential conditioning compared to absolute conditioning (Figure 5D, black dots and line). Specifically, glomeruli tended to increase their activity after training if they were unique in the

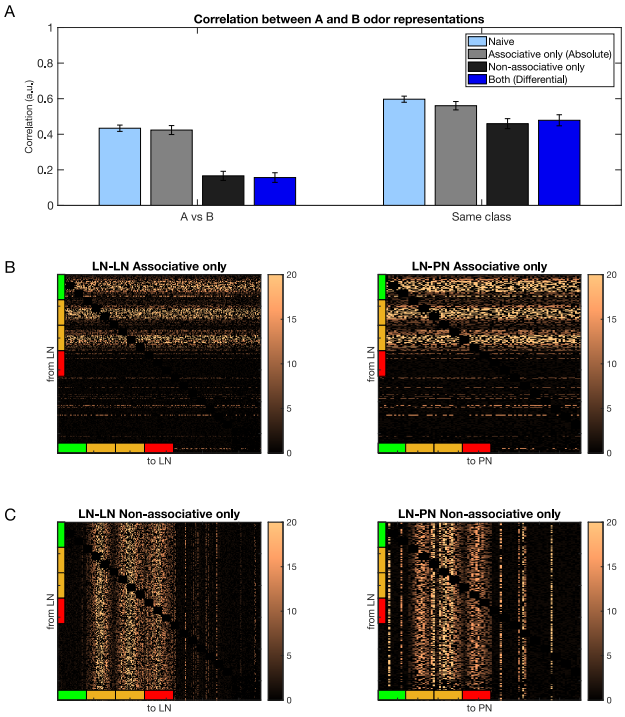

Fig. 4 Connectivity of the inhibitory network after associative only and nonassociative only conditioning. A) Left shows the correlation between rewarded and habituated odor representations does not decrease significantly after associative (gray) conditioning (p=0.325) but decreases significantly after nonassociative (black) conditioning (p=2e-7) and differential (dark blue) conditioning (p=6e-7) compared to naive (light blue). Right shows correlation calculated for odors belonging to the same class, with correlation not changing significantly after associative (gray) conditioning (p=0.1), but decreasing modestly after nonassociative (black, p=0.04) and differential (dark blue) conditioning (p=0.07). The error bars represent standard deviation across trials. B) The LN-LN and LN-PN network after associative only training. C) The LN-LN and LN-PN network after nonassociative only training. The colorbars show the value of the weights after training divided by the initial weights.

representation of the rewarded odor and vice versa, in agreement with what our model predicts. This effect of stimuli being places further apart in the coding space, that aligns with the uniqueness of the percepts, we term *contrast enhancement*.

The inhibitory AL network adapts to discriminate between odors in a new environment

As honeybees move from one environment to another, the odors associated with reward and habituation may change. The animal needs to learn the association of these new odors with the reward contexts present in the new environment. Changes in nectar production throughout the day and potentially many times within a forager's lifetime add to the complexity. The inhibitory AL network, therefore, must be flexible enough to learn representations for the new odors and successfully discriminate between them without immediately forgetting the old environment.

The odors from the new environment (we will call it Env2) may or may not overlap with the odors from the first environment (Env1) the animal has learned already. Below, we discuss four cases based on the overlap between the odors of the two environments. As before, each environment has two classes of odors (e.g., class P - rewarded, and class Q - habituated), each class containing 10 odors.

Case 1: The odor classes from Env1 (P and Q) do not overlap with the odor classes from Env2 (S and T). Case 2: Both rewarded and habituated odor classes from Env1 and Env2 overlap. Here, Env2 consists of odor class M (overlaps with class P) and odor class N (overlaps with class Q). Case 3: One odor class from Env2 (class M) overlaps with the rewarded odor class from Env1 (class P), and the second odor class from Env2 (class S) does not overlap with any odors from Env1. Case 4: One odor class from Env2

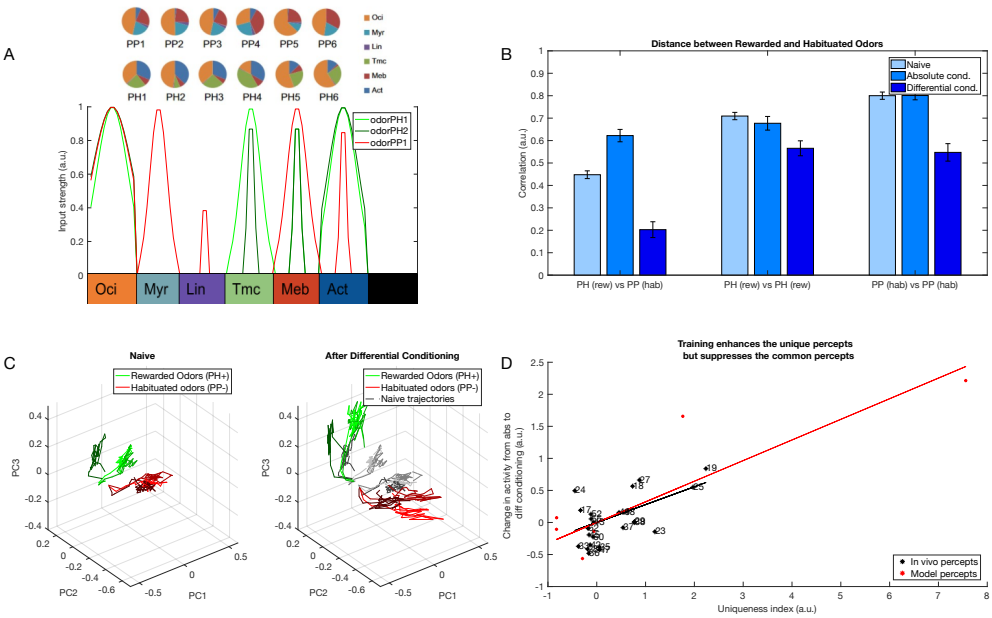

Fig. 5 Contrast enhancement with complex odors based on *in vivo* data. A) Odors obtained from in *vivo experiments* (PH,PP) converted into inputs to the model (see Methods for details). B) Correlation between rewarded (PH) and habituated (PP) odor representations decreases after differential conditioning (dark blue) compared to the naive network (light blue, p=4.2e-29) and absolute conditioned network (middle tone blue, p=8.7e-38). Correlation between odors belonging to same rewarded class (e.g., PH1 vs PH2, PH1 vs PH3, ...) increases slightly after absolute conditioning (middle tone blue, p=0.05) and significantly after differential conditioning (dark blue, p=4.7e-7) compared to the naive case. Correlation between the odors belonging to the habituated class does not decrease after absolute conditioning (middle tone blue, p=0.71) but decreases significantly after differential conditioning (dark blue, p=1.5e-12) compared to the naive case. These results match the experimental data from (Locatelli et al., 2016). C) The dynamical odor trajectory in the 3D PCA space shows the 3 rewarded (PH1, PH2, PH3) and 3 habituated odor (PP1, PP2, PP3) representations shifting in the opposite direction which indicates increased discriminability between odor representations. The error bars represent standard deviation across n=10 trials. D) Contrast enhancement based on Ca^{2+} imaging data (black) and model (red). After differential conditioning, activity in the glomeruli/percepts which have a high uniqueness index (i.e., uniquely activated by the rewarded odor) increases while activity in the glomeruli/percepts which have a low uniqueness index (i.e., activated by both rewarded and habituated odors) decreases (R^2=0.37, p=0.0015 for *in vivo* percepts and R^2=0.76, p=0.02 for model percepts). Note, similar trend *in vivo* and in the model.

(class N) overlaps with the habituated odor class from Env1 (class Q), and the second odor class from Env2 (Class S) does not overlap with any odor from Env1.

The structure of each odor class used in the experiments is shown in Figure 6, at the top of the plots. Each class was based on three percepts, and, as before, individual odors within each class were obtained by varying the activation of individual percepts. The reward structure for the first environment is always such that odor class P is rewarded and class Q is habituated (P+Q-). Each case is further divided into two sub-cases based on which odor class from the second environment is rewarded and which is habituated. Finally, we included a mechanism where synaptic weights return to baseline in the absence of reinforcement (see Methods). This "forgetting" mechanism was slow, so the network does not fully recover to its original naive state between training on Env1 and Env2. Because of this, in some cases, the effects of training on Env1 and Env2 were asymmetric, as Env2 was applied to a network that had already been pretrained on Env1.

Case 1: When there is no overlap between odors from Env1 and Env2, the correlation between representations of classes P and Q decreases significantly after training on Env1 (P+Q-), but increases again after training on Env2. This occurs because the weights associated with discriminating between P and Q odors receive no reinforcement during training on Env2 and slowly return to baseline. The correlation between odors S and T only decreases after training on Env2, regardless of the reward structure (S+T- or S-T+), indicating effective learning in the new environment without interference from prior learning (Figure 6A).

Case 2: When both rewarded and habituated odors overlap between environments (Figure 6B), we observe that if the reward structure remains consistent (M+N-) (Figure 6B, left), the correlation between

P and Q representations remains low after training on Env2. This is because training with overlapping Env2 promotes some of the same synaptic changes learned in Env1. However, when reward associations are reversed (M-N+) (Figure 6B, right), the decorrelation between P and Q increases again to the naive level, more so than in the case of non-overlapping odors, suggesting interference between the two environments.

Case 3: When only the rewarded odors overlap (class M overlaps with class P), the correlation between P and Q remains low when the reward structure of Env2 is consistent with Env1 (Figure 6C, left), but it significantly increases when the reward structure of Env2 is opposite to Env1 (Figure 6C, right). The network is able to learn new associations between M and S regardless of the reward structure of Env2 (Figure 6C).

Case 4: When odors from Env2 (class N) only overlap with the habituated odors from Env1 (class Q), we observe an increase in correlation between P and Q after training on Env2, regardless of its reward structure (Figure 6D). The network is also able to learn new associations between S and N. This suggests that changing the reward association of previously habituated components has minimal effect on existing memories compared to cases where rewarded components change their associations.

These results demonstrate that the AL network can effectively learn new odor-reward associations while maintaining existing ones, particularly when there is minimal overlap between environments or when reward structures remain consistent. The network shows more flexibility in adapting to new reward associations when they do not directly conflict with strongly learned rewarded odor components from the previous environment.

Notably, the overlap between an odor from Env2 and the rewarded odor from Env1 has a more significant impact on the odor correlations compared to an overlap with the habituated odor from Env1. This observation suggests that the network is more sensitive to changes in the reward associations of previously rewarded odors than to changes in the associations of previously habituated odors.

Our study predicts that it takes longer for the network to learn the second environment when the reward structure is reversed, which is consistent with experimental studies on reversal learning (Pavlov, 1928; Komischke et al., 2002). Latent inhibition refers to the observation that prior exposure to a stimulus without any associated reward can slow down the subsequent learning of associations involving that stimulus (Chandra et al., 2010; Bazhenov et al., 2013). In the context of this experiment, the prior learning of reward associations in Env1 can interfere with the learning of new associations in Env2 when the odors overlap but have opposite reward structures, resulting in slower learning and adaptation to the new environment.

.

Graph convolution neural networks and inhibitory learning converge on similar solutions

Can the principles learned from olfaction be exploited to improve performance of artificial neural networks (ANNs)? The encoding of odors in the AL takes place by learning the relationships between the components of the odors and the reward associated with the odors. Graph convolutional network (GCN) is a class of ANNs that take graph data as the input and aggregates information from neighboring nodes of the graph to perform a task (e.g., classification) (Duvenaud et al., 2015; Kipf and Welling, 2016). This enables the GCN to utilize the relationships between nodes to improve the classification performance (Kipf and Welling, 2016). Hence, GCNs can be considered analogous to the AL, in that they encode relationships between different components of the input and utilize these relationships to categorize the inputs.

Here, we compared results obtained using our biophysical model of the honeybee AL network to a GCN. To enable classification, GCN network was followed by simple fully connected network, that can be seen as analogous to subsequent processing centers of the olfactory system, such as mushroom bodies. We trained each network on chemical gas sensor data (Vergara et al., 2012), where the biophysical network model was trained via the inhibitory plasticity (as described before) (Figure 7) and the GCN was trained using backpropagation (see Methods) (Figure 8). The numeric features in this dataset, which included response magnitude, On and Off time constants, and others (see Methods), were used directly to train the GCN model. In the biophysical model, we used three features (response magnitude and On and Off time constants) from the data to construct input pulses to model odor stimulation (Figure 7A).

We selected half of the odors to be associated with the reward and the other half to be habituated. Learning in the biophysical model increased the discriminability between the rewarded and habituated odors as indicated by the decrease in correlation (Figure 7B). Note that this discrimination task was more

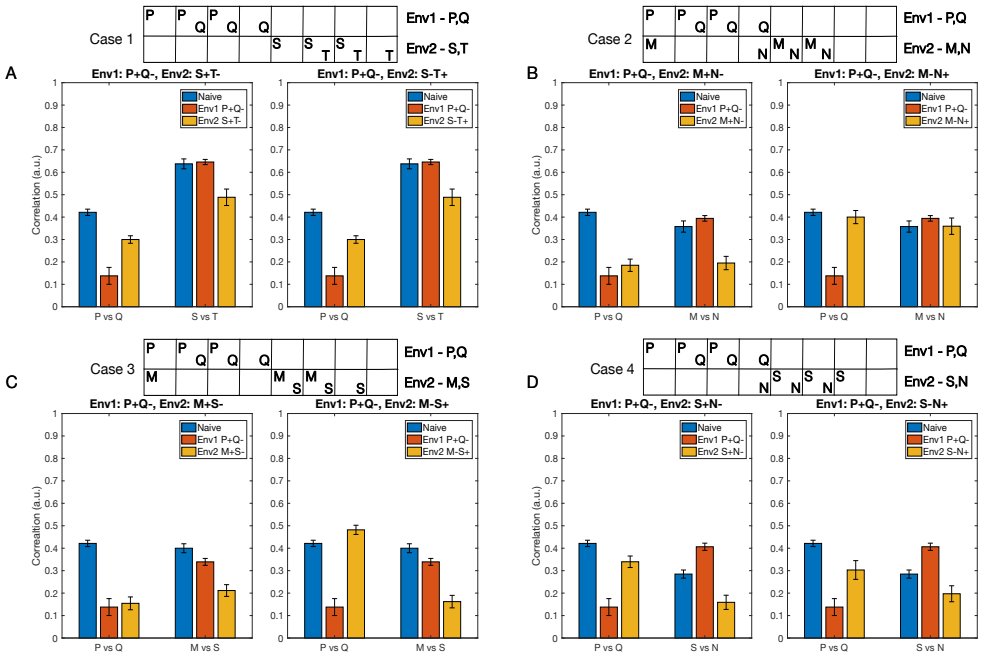

Fig. 6 Learning by the AL inhibitory network in different environments. The network is initially trained on environment 1 (Env1) containing rewarded odor class P and habituated odor class Q. Subsequently, the network is trained on a new environment 2 (Env2), containing different odor classes. Bars show the correlation between odor representations for classes P and Q (left set of bars in each subplot) and between odor classes from Env2 (right set of bars) at different stages: naive (blue), after training on Env1 (orange), and after subsequent training on Env2 (yellow). Each case (1-4) represents a different scenario of odor overlap between Env1 and Env2, with sub-cases (left and right subplots for each case) depicting different reward structures in Env2. A) Case 1: No odor overlap between environments (Env2 odors: S,T). B) Case 2: Complete overlap - both odors from Env2 (M,N) overlap with odors from Env1 (P,Q respectively). C) Case 3: Partial overlap - one odor from Env2 (M) overlaps with the rewarded odor from Env1 (P), while the second odor (S) has no overlap. D) Case 4: Partial overlap - one odor from Env2 (N) overlaps with the habituated odor from Env1 (Q), while the second odor (S) has no overlap. The boxes at the top illustrate the distinct percepts activated by the odors in each environment, with overlapping percepts indicated by shared columns. Lower correlation values indicate better separation between odor representations. Error bars show standard deviation across n=10 trials.

difficult compared to the previous simulations, as identical sets of neurons were activated by different odors. Only the amplitude and On/Off time constants of the input were varied for each percept.

The GCN network (Figure 8A) was able to identify which class an odor belonged to (i.e., rewarded or habituated) with 88% accuracy. Importantly, we found that both the GCN and the biophysical model exhibited the same characteristics in changes to odor representations. Specifically, units representing components unique to an odor were enhanced, while units representing components common to many odors were suppressed (Figure 7C and Figure 8B). Both the top-down approach, rooted in backpropagation, and the bottom-up approach, grounded in local biological learning rules, converged on the same contrast enhancement strategy. This result highlights the effectiveness of contrast enhancement as a computational strategy for accomplishing categorization tasks.

3 Discussion

In vivo recordings from the honeybee's antennal lobe (AL) have shown that separation among patterns of neural activity representing odor blends from different varieties of flowers improves after differential reinforcement for each variety (Locatelli et al., 2016). However, the mechanisms underlying this phenomenon remain poorly understood. The primary goal of this new study was to uncover the plasticity mechanisms that modify neural representations of natural odor blends in the early olfactory system - the honeybee AL - using a combination of computational modeling, machine learning, and Ca^{2+} imaging.

Natural odors are generally blends of several chemical components (Pichersky and Gershenzon, 2002). These odors are present in varying concentrations, and nectar production may change from season to

Table 1 Statistical Analysis Results for Odor Discrimination

Case	Comparison	p-value	Case	Comparison	p-value
Case 1 pos	P/Q (Naive vs Env1)	$1.44 \times 10^{-4}***$	Case 3 pos	P/Q (Naive vs Env1)	$1.44 \times 10^{-4}***$
	P/Q (Env1 vs Env2)	$0.002**$		P/Q (Env1 vs Env2)	0.423
	P/Q (Naive vs Env2)	$0.018*$		P/Q (Naive vs Env2)	$1.21 \times 10^{-4}***$
	S/T (Naive vs Env1)	0.931		M/S (Naive vs Env1)	0.557
	S/T (Env1 vs Env2)	0.174		M/S (Env1 vs Env2)	0.292
	S/T (Naive vs Env2)	0.206		M/S (Naive vs Env2)	0.089
Case 1 neg	P/Q (Naive vs Env1)	$1.44 \times 10^{-4}***$	Case 3 neg	P/Q (Naive vs Env1)	$1.44 \times 10^{-4}***$
	P/Q (Env1 vs Env2)	$8.29 \times 10^{-4}***$		P/Q (Env1 vs Env2)	$1.94 \times 10^{-5}***$
	P/Q (Naive vs Env2)	$0.022*$		P/Q (Naive vs Env2)	0.152
	S/T (Naive vs Env1)	0.931		M/S (Naive vs Env1)	0.557
	S/T (Env1 vs Env2)	0.099		M/S (Env1 vs Env2)	0.105
	S/T (Naive vs Env2)	0.128		M/S (Naive vs Env2)	$0.016*$
Case 2 pos	P/Q (Naive vs Env1)	$1.44 \times 10^{-4}***$	Case 4 pos	P/Q (Naive vs Env1)	$1.44 \times 10^{-4}***$
	P/Q (Env1 vs Env2)	0.062		P/Q (Env1 vs Env2)	$2.77 \times 10^{-4}***$
	P/Q (Naive vs Env2)	$2.76 \times 10^{-4}***$		P/Q (Naive vs Env2)	0.061
	M/N (Naive vs Env1)	0.813		S/N (Naive vs Env1)	0.415
	M/N (Env1 vs Env2)	0.259		S/N (Env1 vs Env2)	0.088
	M/N (Naive vs Env2)	0.240		S/N (Naive vs Env2)	0.334
Case 2 neg	P/Q (Naive vs Env1)	$1.44 \times 10^{-4}***$	Case 4 neg	P/Q (Naive vs Env1)	$1.44 \times 10^{-4}***$
	P/Q (Env1 vs Env2)	$1.49 \times 10^{-5}***$		P/Q (Env1 vs Env2)	$1.80 \times 10^{-4}***$
	P/Q (Naive vs Env2)	0.522		P/Q (Naive vs Env2)	$0.009**$
	M/N (Naive vs Env1)	0.813		S/N (Naive vs Env1)	0.415
	M/N (Env1 vs Env2)	0.812		S/N (Env1 vs Env2)	0.141
	M/N (Naive vs Env2)	0.988		S/N (Naive vs Env2)	0.498

$^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$. All p-values were calculated using unpaired two-tailed Student's t-tests.

season and environment to environment. Although the olfactory coding space is very large, optimally mapping the sensory environment requires learning and relearning the association of rewards with variable blends of volatile compounds (Pichersky and Gershenzon, 2002). We found that, through a combination of associative and nonassociative plasticity in the inhibitory AL network, the activity of components common to both the rewarded and habituated odors decreases, while the activity of components unique to the rewarded odors increases after differential conditioning. This leads to a decrease in overlap between the representations of the odor classes, while simultaneously making the representation of the rewarded odors more compact. This "contrast enhancement" effect was predicted by our biophysical model and the model's predictions were confirmed by the Ca^{2+} imaging data from the honeybee AL.

We also found that, when subsequently stimulated by odors from a different environment, the inhibitory AL network learned the odor-reward associations from the new environment while preserving associations from the old environment. To test the application of these principles to Artificial Intelligence, we developed and trained a machine learning model based on the GCN to perform the same task of categorizing complex odors. Analyzing the activity of the units in the GCN, we observed a similar strategy of changes in the representation of the rewarded odors as seen in the biophysical model and in vivo data.

Previous works have successfully applied computational network models to study the mechanisms of odor-induced synchronization (Bazhenov et al., 2001b; Sanda et al., 2016; Assisi et al., 2011; Ito et al., 2009; Bazhenov et al., 2001a) and neuronal plasticity (Locatelli et al., 2013; Linster and Smith, 1997; Chen et al., 2015; Bazhenov et al., 2013; Finelli et al., 2008; Bazhenov et al., 2005; Assisi et al., 2020) in the early olfactory system of insects. The study by (Locatelli et al., 2013) used both firing rate and conductance-based models to explore nonassociative plasticity in the AL. The model in (Chen et al., 2015) revealed that stimulus-specific changes in synaptic inhibition are sufficient to explain shifts in odor representations after olfactory learning.

Our new study predicts that the AL network may utilize inhibitory synaptic plasticity to optimize its responses to odors associated with reward in a given odor environment. Considering the limited neuronal resources of the insect olfactory system and the complexity of the tasks it needs to solve, this study proposes a novel strategy by which the early olfactory system facilitates odor processing at downstream levels, focusing on odors relevant for survival. The model predictions align with imaging

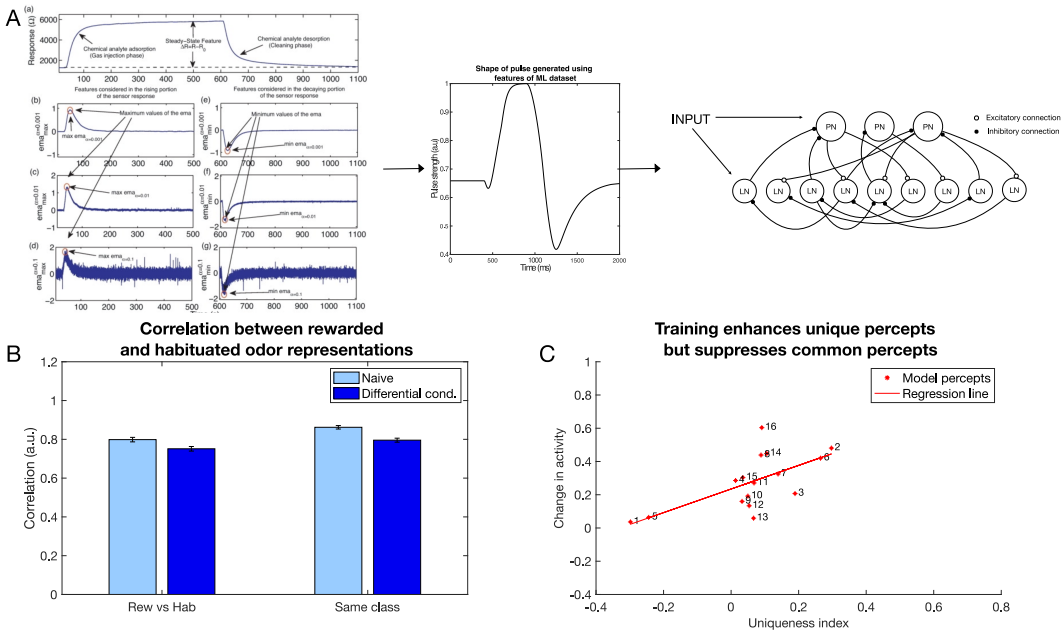

Fig. 7 Contrast enhancement in a large network model with inputs generated from chemical gas sensor data (Gas Sensor Array Drift Dataset from (Vergara et al., 2012)). A) The features extracted from the response curves are converted into a pulse that is then fed into the larger biophysical model as an input. B) Correlation between (left set of plots, p=7.4e-18) and within classes (right set of plots, p=1.6e-17) decreases after differential conditioning (dark blue) compared to naive (light blue). The error bars represent standard deviation across trials. C) Contrast enhancement analysis for the large biophysical network using chemical gas sensor data revealed similar results to the smaller biophysical model and in vivo data ($R^2 = 0.416$, p=0.007)(compare to Fig. 5D).

studies in the honeybee (Faber et al., 1999; Locatelli et al., 2016; Sachse and Galizia, 2002; Sandoz, 2003) and other insects (Wilson et al., 2004) showing that odor representations in the AL change after differential conditioning, with rewarded and unrewarded odor representations becoming less correlated. These findings support the notion that the increase in the separation of odor representations in the AL *in vivo* occurs through contrast enhancement, thereby validating the prediction made by our model. The main principles of nonassociative and associative plasticity operating in the first relay for olfactory coding appear to be similar in the brains of both insects and vertebrates (Linster et al., 2009; Wilson and Linster, 2008; Fernandez et al., 2009; Locatelli et al., 2016, 2013) and our model predictions can be generalized beyond insect olfaction.

Honeybees have a rich repertoire of learning behaviors to odors that range from nonassociative to associative and operant conditioning (Hammer and Menzel, 1995; Smith et al., 2006). Biogenic amines have been shown to play an important role in reinforcement during olfactory learning (Hammer, 1993; Hammer and Menzel, 1998; Aso et al., 2010). In the honeybee, a cluster of cells at the base of the brain called Ventral Unpaired Medial (VUM) cells receive input from sucrose (US)-sensitive taste receptors (Hammer, 1993; Laurent et al., 1998; Bicker and Kreissl, 1994; Abel et al., 2001; Sinakevitch et al., 2005). The outputs from VUM cells spread broadly throughout association areas of the brain, including the AL and the MB, where they release the biogenic amines octopamine and tyramine as the reinforcement signal. Immunological studies have shown that the octopamine receptor (AmOA1) is expressed alongside GABA receptors in the honeybee AL (Sinakevitch et al., 2013). Other studies (Rein et al., 2013) have shown that octopamine modulates odor representations in PNs by regulating inhibitory connections in the AL. While other types of octopamine receptors, such as Octβ1R and Octβ2R, are also expressed in the honeybee brain, here we focus on AmOA1 (Sinakevitch et al., 2011), which is orthologous to Drosophila's DmOAMB receptor. In honeybees (Grohmann et al., 2003) and fruit flies (Lee et al., 2003), octopamine binding to this receptor leads to an increase in intracellular calcium, presumably resulting in an excitatory effect. Importantly, this receptor is present on a subset (i.e., not all) of the inhibitory interneurons in the antennal lobe of honeybees and fruit flies (Sinakevitch et al., 2011). Thus, activation of these receptors

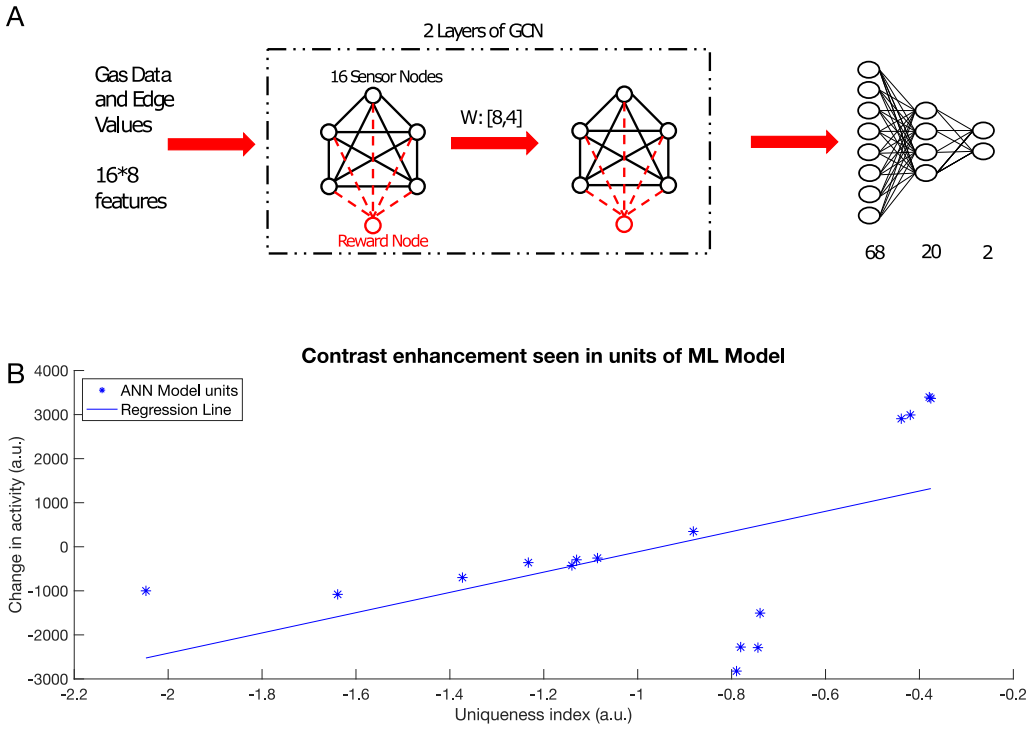

Fig. 8 Contrast enhancement in a machine learning model. A) Schematic of the GCN model. The model receives features from the Gas Sensor Array Drift Dataset and processes it through two graph convolutional layers. The graph is then unrolled into a vector and sent through two fully connected layers for categorization. Training weights between GCN layers can be thought of as analogous to the biophysical model training. B) Change in activity in the graph nodes in the first and second layer of the GCN vs uniqueness index of that node ($R^2 = 0.276$, p=0.031). Note contrast enhancement: nodes with higher uniqueness tend to increase the activity and those with lower uniqueness index tend to decrease the activity.

on a subset of LNs may inhibit other LNs and thereby lead to partial PN disinhibition, as reported in (Rein et al., 2013).

Repeated or prolonged exposure to a stimulus without any reinforcement reduces the behavioral response to those odors, a process known as habituation (Thompson and Spencer, 1966; Bazhenov et al., 2013; Chandra et al., 2010), which leads to latent inhibition (Chandra et al., 2010; Latshaw et al., 2023). Studies in the fruit fly AL have revealed that habituation arises from the facilitation of inhibitory LN to PN synapses involved in odor representation (Larkin et al., 2010; Sudhakaran et al., 2012).

The model predicts that both associative and nonassociative plasticity are required to explain *in vivo* data. The suppression of the common component in the rewarded odors was driven by an increase in the inhibitory LN to PN connection strength, caused by postsynaptic (nonassociative) plasticity. This effect resulted from the firing of PNs triggered by the habituated odor, aligning with the findings of (Locatelli et al., 2013) and (Chen et al., 2015). These studies demonstrated that nonassociative inhibitory plasticity from LNs to PNs plays a crucial role in 'filtering out' components of the habituated odor after honeybee was exposed to the odor without receiving a reward. The increase in the activity of the unique component in the rewarded odors occurred due to both presynaptic and postsynaptic facilitation, leading to an increase in inhibition of the LNs that were inhibiting PNs representing the unique components (i.e., PNs' disinhibition). A similar effect, an increase in the activity of the glomeruli for the rewarded but not the habituated odor, was also observed in (Faber et al., 1999). Together, these two mechanisms contributed to shifting the PN responses for the rewarded odor away from those of the habituated odor. Our AL model focuses on the synaptic connections between excitatory PNs and inhibitory LNs. While

those represent the canonical AL circuit, other types of connections are found in different insect species, e.g., in fly olfactory system the vast majority of LN also synapses onto ORNs (Wilson and Linster, 2008). While we acknowledge this model limitation, our focus on inhibitory plasticity (LN-to-PN and LN-to-LN synapses) reflects our central hypothesis that experience-dependent modifications of inhibitory circuits reshape odor representations in the AL.

The mushroom body (MB), the next layer in olfactory processing in insects, is considered a major center for associative learning and olfactory memory (Hammer and Menzel, 1998). Information from the PNs in the AL is relayed to the MB, where Kenyon cells (KCs) represent odors with sparse firing patterns (Stopfer et al., 2003; Perez-Orive et al., 2002). Remarkably, approximately 800 PNs synapse onto around 170,000 KCs (Mobbs and Young, 1997; Gronenberg, 1986), significantly increasing the dimensionality of the odor representation space. Experimental studies have shown that KC responses are influenced by associative reinforcement learning, which stabilizes odor representations in the MB. In contrast, odor presentations without reinforcement weaken their representation in the MB (Szyszka et al., 2008). Modeling studies (Rajagopalan and Assisi, 2020) have also found that the categorization of odors in the MB depends on changes in PN responses caused by conditioning. Our study suggests a potential mechanism for this change in KC responses. After training, the rewarded odor triggers more population-specific AL activity, leading to more reliable firing of the KCs receiving input from these PNs. This observation aligns with the findings of (Szyszka et al., 2008).

Artificial intelligence (AI) has made rapid progress in recent years in solving tasks such as image classification, speech recognition and natural language processing. Although it has now grown into a separate field, AI can trace its roots back to neuroscience with the earliest artificial neural networks designed to mimic information processing in the brain (Hassabis et al., 2017). Currently, efforts are being made to incorporate biological principles to develop a new generation of AI (Hayes et al., 2021; Kudithipudi et al., 2022). Graph Neural Networks (GNNs) have been developed to work with graph data. There have been significant advances in GNNs, increasing their capabilities and expressive power (Kipf and Welling, 2016; Reiser et al., 2022). In one study, GNNs were used to learn a generalizable perceptual representation of odors (Sanchez-Lengeling et al., 2019). Here, we developed a graph convolutional network (GCN) model as an artificial parallel to the insect antennal lobe. We trained this network to perform an odor categorization task and analyzed the activity of the artificial neurons in the first and second layers of the GCN. This analysis revealed that the contrast between representations of different odor classes increased via the same mechanism described in the biophysical model: units representing components unique to an odor were enhanced, while units representing components common to many odors were suppressed. Together, these results suggest that contrast enhancement is an efficient strategy for performing a categorization task.

In summary, our study has uncovered novel circuit-level mechanisms of olfactory learning in the honeybee AL that enhance odor discrimination. This learning paradigm relies on inhibitory network plasticity, which allows the coding space to be reshaped to align with the current task and environment. These findings suggest an effective computational strategy for the perceptual learning of complex natural odors, achieved through changes in the inhibitory network during early sensory processing.

4 Methods and Materials

Animals and calcium imaging experiments

Pollen forager honey bees (*Apis mellifera*) were collected from regular hives and restrained in individual harnesses suited for olfactory conditioning and calcium imaging recordings. Appetitive olfactory conditioning followed standard protocols (Smith and Burden, 2014). Two groups of bees received different conditioning protocols using odor blends that mimicked the natural odor variation within and between two varieties of Snapdragon flowers, Potomac Pink (PP) and Pale Hybrid (PH) (Wright et al., 2005; Locatelli et al., 2016). One group of bees was subject to differential conditioning (PH+/PP-) by rewarding with 1M sucrose solution odor blends from the PH variety and not rewarding blends of the PP variety. The second group of bees (PH+/MO-) was trained rewarding the same PH blends as for the first group, but the odor solvent mineral oil (MO) was used instead PP blends for unrewarded trials. Thus, the difference between groups was in the unrewarded trials. Eight hours after conditioning, uniglomerular projections neurons of the antennal lobe were stained by backfilling with the calcium sensor dye fura-dextran (potassium salt, 10.000 MW, ThermoFisher Scientific) (Sachse and Galizia, 2002). Calcium imaging was performed on the next day. The head-capsule was opened and rinsed with Ringer solution (130mM NaCl, 6mM KCl, 4mM

MgCl2, 5mM CaCl2, 160mM sucrose, 25mM glucose, 10mM HEPES, pH 6.7, 500mOsmol; all from Sigma-Aldrich). Glands and trachea covering the antennal lobes were removed. Only one antennal lobe per bee was used for calcium imaging. Calcium imaging was done using a CCD camera (SensiCamQE, T.I.L.L. Photonics) mounted on an upright fluorescence microscope (Olympus BX-50WI, Japan) equipped with a 20x objective, NA 0.95 (Olympus), 505 DRLPXR dichroic mirror, and 515nm LP filter (TILL Photonics). Monochromatic excitation light provided by a PolichromeV (TILL Photonics) alternated between 340nm and 380nm. Images were recorded at a sampling rate of 8Hz. Spatial resolution was 172 x 130 pixels, after a binning of 8 x 8 on a chip of 1376 x 1040 pixels, resulting in a spatial sampling of 2.6µm per pixel side. Each data set consisted in a double sequence of 80 images, obtained at 340nm and 380nm excitation light respectively (Fi340, Fi380, where i is the number of images from 1 to 80). Calcium signals were calculated using a ratiometric method: for each pair of images Fi we calculated Ri = (Fi340/Fi380) x 100 and subtracted the background Rb that was obtained by averaging the Ri values 1s before odor onset. The analysis was based on the calcium signals from the same 24 glomeruli identified in all bees (figure 1C) (Galizia et al., 1999). The glomerular activation was considered as representing the activity of uniglomerular PNs, as only these neurons in the AL were stained. The odor delivery device consisted of a main charcoal filtered air stream (500ml/min) and fifteen identical odor channels, each composed of one valve and one odor cartridge. The odor cartridges consisted of a 1ml syringe containing a paper strip loaded with 10ul odor solution. A three-way valve (LFAA1200118H; The LEE Company) controlled the airflow through the odor cartridge. Opening of the valves was synchronized with the optical recordings by the acquisition software TILLVisION (TILL Photonics). During stimulation with odor, 4ml of air loaded with the odor was added to the main air stream during 4 seconds. Calcium imaging analysis was designed to determine how both trainings protocols affect the neural representation of PH and PP blends. The calcium activity measured between 375ms and 625ms after odor onset was averaged, thus the activation pattern elicited by each odor was reduced to a single vector of 24 elements (glomeruli) (see figure 1A for examples). This time interval includes the time points in which odor representations reach the minimal correlation after odor onset. The similarity between 2 glomerular odor patterns was calculated as the correlation between the respective 24-dimensional vectors. As we measured glomerular activity patterns elicited by 6 PH blends and 6 PP blends, we could calculate for each animal the average distance between 36 pairs of PH-PP blends. This was done for 10 absolute conditioned bees and 7 differentially conditioned bees. This data has been previously published in (Locatelli et al., 2016). In this manuscript, we re-analyze this calcium imaging data.

Computational Modeling

We constructed a biophysical model of the honeybee antennal lobe (AL), containing 100 excitatory projection neurons (PNs) and 280 inhibitory local interneurons (LNs). Each PN and LN was modeled by a single compartment that included voltage- and calcium-dependent currents described by Hodgkin-Huxley kinetics (Hodgkin and Huxley, 1952). The model here was adapted from our previous studies (Haney et al., 2018), (Chen et al., 2015) and is available for download from the Bazhenov lab website. Model equations were solved using a fourth order Runge-Kutta method with an integration time step of 0.04ms.

Membrane Potentials

PN and LN membrane potential equations (Hodgkin and Huxley, 1952) are given by

$$C_m \frac{dV_{PN}}{dt} = -g_L(V_{PN}-E_L)-I_{Na}-I_K-I_A-I_T-I_h-g_{KL}(V_{PN}-E_{KL})-I_{GABA_A}-I_{\text{slow-inh}}-I_{nACh}-I_{stim}$$

$$C_m \frac{dV_{LN}}{dt} = -g_L(V_{LN}-E_L)-I_{Na}-I_K-I_T-g_{KL}(V_{LN}-E_{KL})-I_{GABA_A}-I_{nACh}-I_{stim}$$

The passive parameters are as follows: For PNs, $C_m = 2.9 * 10^{-4}\mu F$, $g_L = 0.01mS/cm^2$, $g_{KL} = 0.012mS/cm^2, E_L = -70mV, E_{KL} = -95mV$.

For LNs, $C_m = 1.43 * 10^{-4}\mu F$, $g_L = 0.05mS/cm^2$, $g_{KL} = 0.018mS/cm^2, E_L = -70mV, E_{KL} = -95mV$. An external DC input was introduced to each neuron through I_{stim}.

The intrinsic currents, including a fast sodium current, I_{Na}, a fast potassium current, I_K, a transient potassium A-current, I_A, a low threshold transient Ca^{2+} current, I_T, a hyperpolarization activated cation current I_h are given by equation

$$I = gm^M h^N (V - E)$$

where the maximal conductances for PNs are $g_{Na} = 90mS/cm^2, g_K = 10mS/cm2, g_A = 10mS/cm^2, g_T = 2mS/cm^2, g_h = 0.02mS/cm^2$. The maximal conductances for LNs are $g_{Na} = 100mS/cm2, g_K = 10mS/cm^2$, and $g_T = 1.75mS/cm^2$.

For all cells, $E_{Na} = 50mV, E_K = 95mV$, and $E_{Ca} = 140mV$. The gating variables $0 \leq m(t), h(t) \leq 1$ satisfy the following:

$$\frac{dm}{dt} = \frac{m_\infty(V) - m}{\tau_m(V)}$$
$$\frac{dh}{dt} = \frac{h_\infty(V) - h}{\tau_h(V)}$$

where the steady state (e.g., $m_\infty(V)$) and relaxation times (e.g., $\tau_h(V)$) are derived from experimental recordings of the specific ionic currents. These distinct voltage-dependent functions are given in (Chen et al., 2015). For all cells, intracellular Ca^{2+} dynamics were described by a simple first-order model as follows:

$$\frac{d[Ca^{2+}]}{dt} = -A \cdot I_T - \frac{[Ca^{2+}] - [Ca^{2+}]_\infty}{\tau}$$

where $[Ca^{2+}]_\infty = 2.4 \cdot 10^{-4} mM$ is the equilibrium of intracellular Ca^{2+} concentration, $A = 5.2 \cdot 10^{-5} mM \cdot cm^2/(ms \cdot \mu A)$ and $\tau = 5ms$.

Synaptic Currents

Fast GABA and cholinergic synaptic currents to LNs and PNs were modeled by first-order activation schemes (Destexhe et al., 1994). Fast GABA and cholinergic synaptic currents are given by:

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

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

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

. For the transmitter concentration of the cholinergic synapses, $[T]$, is given by

$$[T] = A \cdot H(t_0 + t_{max} - t) \cdot H(t - t_0)$$

and for the GABAergic synapses

$$[T] = \frac{1}{1 + \exp(-(V - V_0)/\sigma)},$$

where H is the Heaviside step function, t_0 is the time of receptor activation, $A = 0.5, t_{max} = 0.3ms, V_0 = -20mV$ and $\sigma = 1.5$. The rate constants were given $\alpha = 10ms^{-1}$ and $\beta = 0.2ms^{-1}$ for fast

$$GABA_A$$

synapses, and $\alpha = 1ms^{-1}$ and $\beta = 0.2ms^{-1}$ for cholinergic synapses.

The slow inhibitory synaptic current is given by the following equations:

$$I_{\text{slow-inh}} = g_{\text{slow-inh}} \frac{[G]^4}{[G]^4 + K}(V - E_k)$$

,

$$\frac{d[R]}{dt} = r_1(1-[R])[T] - r_2[R]$$

,

$$\frac{d[G]}{dt} = r_3[R] - r_4[G]$$

where $[R]$ is the fraction of activated receptors, $[G]$ is the concentration of G proteins, $E_K = -95\text{mV}$ is the potassium reversal potential. The rate constants were $r_1 = 0.5\text{mM}^{-1}\text{ms}^{-1}$, $r_2 = 0.0013\text{ms}^{-1}$, $r_3 = 0.1\ \text{ms}^{-1}$, $r_4 = 0.033\text{ms}^{-1}$ and $K = 100\ \mu\text{M}^4$. The maximum synaptic conductances were set at $g_{GABA_A} = 0.02\mu S$ between LN, $g_{GABA_A} = 0.015\mu S$ from LN to PN, $g_{ACh} = 0.3\mu S$ from PN to LN and $g$slow-inh $= 0.02\ \mu$ from LN to PN.

For simulating the larger network with the gas sensor array drift dataset as input, the maximum synaptic conductances were as follows: $g_{GABA_A} = 0.024\mu S$ between LNs, $g_{GABA_A} = 0.019\mu S$ from LNs to PNs and $g_{ACh} = 0.075\mu S$ from PNs to LNs.

Plasticity

A simple phenomenological model of synaptic facilitation was used to model inhibitory plasticity (from LN to LN and LN to PN). Specifically, the inhibitory network underwent presynaptic facilitation when presented with a rewarded odor and postsynaptic facilitation when presented with a habituated odor during learning (Chen et al., 2015). Presynaptic (associated with reward) facilitation was assumed to depend on the octopamine receptors AmOA1 expressed in inhibitory LNs in the honeybee AL (Sinakevitch et al., 2013), (Sinakevitch et al., 2011) and it was based on the spiking events of the presynaptic neurons. In contrast, postsynaptic facilitation (associated with odor habituation) was based solely on the spiking event of the postsynaptic neuron (Sudhakaran et al., 2012), (Das et al., 2011).

To model facilitation, the maximum synaptic conductance is multiplied by a facilitation variable, F. F is updated each time there is a spike:

$$F_{t+1} = F_t + dF_{pre} + dF_{post}$$

where dF is facilitation rate ($dF_{pre} = 0.15$ and $dF_{post} = 0$ for presynaptic facilitation and $dF_{pre} = 0$ and $dF_{post} = 0.15$ for postsynaptic facilitation). F_t is the value of the facilitation variable just before the spike. In the absence of any spiking events, F had an exponential decay according to the following equation

$$F_t = 1 + (F_{t-1} - 1) * \exp\frac{-((t-1)-t_i)}{\tau}$$

In the above equation, $\tau = 30\text{s}$ is the time constant of the decay; t_i is the time of the ith spiking event; F_t corresponds to the value of F at time t with the initial value $F_0 = 1$. The time constant of decay is the same for pre and postsynaptic plasticity. The updated synaptic weights at the end of the training period (30 presentations, 60,000 ms) were frozen and used as the synaptic weights in the testing phase.

Network Geometry

The AL network included 100 PNs and 280 LNs. All these PNs and LNs were organised into 20 glomeruli, with each glomerulus containing 12 LNs (unipolar LNs) and 5 PNs. In addition to these, there are 40 LNs that contribute to global inhibition (multipolar LNs). The connections between each group of neurons were generated randomly based on the probabilities described in Table 2

In order to accommodate the ML-dataset inputs to the model, in some simulations we increased the number of neurons to 400 PNs and 1120 LNs. The larger network has random connectivity between LNs, LNs to PNs and PNs to LNs with a probability of 0.125. There were no connections between the PNs.

Odor Stimulation

To model odor stimulation, a fraction of LNs and PNs was activated by a current pulse with a rise time constant of 66.7ms and a decay time constant of 200ms for a period of 500ms as can be seen in Figure 2F. In addition, small amplitude current in the form of Gaussian Noise was added to each cell, to ensure random and independent membrane potential fluctuations. Each odor was modelled as a mixture of 'pure' chemicals. Each pure chemical is defined here as a percept. Each percept activates a group of 'virtual

Table 2 Connection probabilities between different neuron groups in the AL.

Source Unit	Target Unit	Connection Probability
All PNs	All PNs	0
	Unipolar LNs	0.4
	Multipolar LNs	0.4
Unipolar LNs	PNs from same glomerulus	0
	PNs from different glomerulus	0.5
	Unipolar LNs from same glomerulus	0
	Unipolar LNs from different glomerulus	0.4
	Multipolar LNs	0.3
Multipolar LNs	All PNs	0.3
	Mutlipolar LNs	0.1
	Unipolar LNs	0.4

ORNs', which in turn activate a set of PNs. The odors we created activate 3 out of 7 percepts, which means each odor is a mixture of 3 'pure' chemicals out of a total of 7 'pure' chemicals.

For example, for simulations shown in (Figure 3), the odor classes (A and B) were defined based on which 3 out of the 7 percepts were active for a given class. One odor class (A) was associated with reward and another one (B) with no reinforcement, with an overlap of 2 percepts. Individual odors within a class were defined by changing the level of LN/PN activation triggered by percepts.

In some simulations we modeled "PH-PP" odors which were created from the chemical compositions of two varieties of snapdragon odor blends (flowers) – PH (Pale Hybrid) and PP (Potomac Pink) as used in (Locatelli et al., 2016). Since there are 6 chemical components present in each odor, each chemical component was assigned a percept and the level of activation of each percept (width of the gaussian centered around the center of the percept in Fig5A) was calculated based on the proportion of that chemical component present in the natural odor blend. The 7th percept remained inactive for all odors. For example, in the odor PH1, Oci chemical is 37.2% of the mixture. Hence, the standard deviation, which determines the width for the Gaussian, for percept 1 is set to 0.372 (Figure 5A).

Finally, we also used odor classes derived from the UCI Gas Sensor Array Drift dataset (Vergara et al., 2012). This dataset consists of data for 6 odors, obtained from 16 metal oxide sensors. Each odor consisted of 16 percepts, which activated a certain group of PNs and LNs. The first 200 out of 400 PNs and first 560 LNs out of 1120 LNs were provided with inputs, whereas the second half of PNs and LNs received no odor related inputs. This was done in order to maintain the E/I balance in the larger network. The first half of PNs and LNs were provided with odor pulses, created from 3 features of the data from each metal oxide sensor - ΔR, $ema_max_{\alpha=0.001}$ and $ema_min_{\alpha=0.001}$ (Vergara et al., 2012). From these 3 features, we created an odor pulse with a rise time constant proportional to $1/ema_max_{\alpha=0.001}$, decay time constant proportional to $1/ema_min_{\alpha=0.001}$ and the maximum height of the pulse proportional to ΔR. Of the 6 odors, 3 were associated with reward and 3 were associated with no reinforcement.

Machine Learning Model - GCN model

In the last section of this paper, we applied a Graph Convolutional Network (GCN), a type of Graph Neural Network (GNN), to simulate AL training using machine learning approach. GCN is an approach for semi-supervised learning on graph structured data. It is convolutional in nature, because filter parameters are typically shared over all locations in the graph. GCN is a generalized version of CNNs which operate directly on arbitrarily structured graphs (Duvenaud et al., 2015). In (Kipf and Welling, 2016), the choice of convolutional architecture is motivated via a localized first-order approximation of spectral graph convolutions.

For the GCN model, the goal is to learn a function of signals/features on a graph which takes as input:

- A feature description x_i for every node i; summarized in a N×D feature matrix X (N: number of nodes, D: number of input features)
- A representative description of the graph structure in matrix form; typically in the form of an adjacency matrix A (or some function thereof) and produces a node-level output Z (an N×F feature matrix, where F is the number of output features per node). Graph-level outputs can be modeled by introducing some form of pooling operation. (Duvenaud et al., 2015)

Every neural network layer can then be written as a non-linear function $H(l+1) = f(H(l), A)$ with $H(0) = X$ and $H(L) = Z$, L being the number of layers. The specific models then differ only in how f is chosen and parameterized.

The UCI Gas Sensor Array Drift Dataset was taken as the input to the GCN, which consists of 2 parts, i.e, the graph convolutional layers and the fully connected layers. This dataset contains data for 6 different odors from 16 metal oxide sensors, with 8 features extracted from each sensor response, as described above. Thus, each odor is made of 128 features from 16 sensors. The data was normalized by subtracting the mean and dividing by standard deviation for each sensor value over the training and test sets. The GCN model consists of 16 (sensor) + 1 (reward/habituation) nodes and 2 layers of graph convolution. The output of the GCN is converted to a vector and sent to 2 fully connected layers for the final classification. The 6 odors from the dataset are split into 2 groups, one group associated with reward and one group associated with habituation. The task being performed by the model is to determine which group the given odor belongs to. The equation of this GCN model is, then, $H(l+1) = h(A \cdot H(l) \cdot W(l))$, where l is the current layer and A is the graph adjacency matrix with all connections between nodes belonging to sensors set to 1. The reward/habituation nodes' connections to the main 16 node graph was set during training as follows: if a rewarded odor was presented to the model during training, the connections between the reward and sensor nodes was set to 1. Conversely, if a habituated odor was presented to the model during training, the connection between reward node and sensor nodes was set to -1.

The model was trained via backpropagation, with the matrix W and the fully connected head of the model being updated during training. The cross entropy loss was optimized with the Adam Optimizer with a learning rate of 0.01 for 20 epochs. The graph convolution part of the model takes in 8 features from the 16 sensors and transform them into 4 abstract ones at the second layer. These abstract features encode the relationships between the different sensor features during training, similar to the AL which learn the relationships between the different odor percepts during training. Hence, we can say that the GCN model performs input preprocessing by learning the relationships between input features similar to the AL in honeybees. The first layer of the GCN can be thought of as analogous to the untrained AL network and the second layer can be thought of as the trained AL network.

Analysis

Correlation Analysis

Correlation between odor representations was calculated based on the normalized binned spike counts across all the PNs for a single trial during the period of stimulation, with a bin size of 100ms. The correlation was calculated between arrays of PN activity for each time bin and each trial and averaged over both time bins and trials.

PCA trajectory

The response trajectory of the entire PN population to a specific input was constructed by binning the spike trains in 40ms bins. For visualization, the 100 dimensional space of PN responses was reduced to 2D/3D space using Principal Component Analysis (PCA).

Uniqueness index calculation

The uniqueness index for each glomerulus/percept (unit) was calculated as follows:

$$\text{UI} = \frac{(act_unit^{\text{rew}}_{\text{naive}} - act_unit^{\text{hab}}_{\text{naive}})}{act_unit^{\text{hab}}_{\text{naive}}}$$

and the change in activity of the neurons:

$$\text{Change in activity} = \frac{act_unit^{\text{rew}}_{\text{diff}} - act_unit^{\text{rew}}_{\text{naive}}}{act_unit^{\text{rew}}_{\text{naive}}}$$

where $act_unit^{\text{rew}}_{\text{naive}}$ denotes the activation of the glomerulus by the rewarded odor and $act_unit^{\text{hab}}_{\text{naive}}$ denotes the activation by the habituated odor before training (in the naive condition) and $act_unit^{\text{rew}}_{\text{diff}}$ denotes the activation of the glomerulus by the rewarded odor after differential conditioning. A high positive Uniqueness index (UI) indicates that the given glomerulus is uniquely activated by the rewarded odors and a UI close to 0 would indicate the glomerulus is activated to a common extent by both rewarded and habituated odors. This method of calculating the UI was used for the biophysical models.

The uniqueness index for the GCN model was calculated based on the average activation (average of 8 feature values) for each sensor in GCN layer 1 for rewarded and habituated odors. The UI for the GCN is:

$$\mathrm{UI} = \frac{act_unit_{\mathrm{L1}}^{\mathrm{rew}} - act_unit_{\mathrm{L1}}^{\mathrm{hab}}}{act_unit_{\mathrm{L1}}^{\mathrm{hab}}}$$

where $act_unit_{\mathrm{L1}}^{\mathrm{rew}}$ corresponds to the average activation of the GCN unit for the rewarded odors in layer 1 and $act_unit_{\mathrm{L1}}^{\mathrm{hab}}$ corresponds to the average activation of the GCN unit for the habituated odors in layer 1. The change in activity was calculated based on the difference in average activation of each sensor node between layer 1 and layer 2 of the GCN:

$$\text{Change in activity} = act_unit_{\mathrm{L2}}^{\mathrm{rew}} - act_unit_{\mathrm{L1}}^{\mathrm{rew}}$$

Statistics and Reproducibility

Statistical analysis was performed using unpaired two-tailed Student's t-tests to compare the correlation between the different conditions for the computational model. The Wilcoxon signed rank test was used for the comparison of correlation between representations for the calcium imaging data in Figure 1. All statistical tests were performed in MATLAB (MathWorks). Exact p-values are reported in the figure legends. For computational modeling experiments, n represents the number of independent simulation runs with different random seeds, which we call trials. For the calcium imaging experiments, n represents the number of different bees used for each condition. The computational model code used in this study is available at Dr. Bazhenov's lab website. All simulations are deterministic given the same random seed.

5 Acknowledgments

This work was supported by grants - NSF (2223839, 2323241), NIH (R01DC020892) and NSF/CIHR/DFG/FRQ/UKRI-MRC Next Generation Networks for Neuroscience Program (2014217).

Declarations

- Funding - This work was supported by grants - NSF (2223839, 2323241), NIH (R01DC020892) and NSF/CIHR/DFG/FRQ/UKRI-MRC Next Generation Networks for Neuroscience Program (2014217).
- Competing interests - The authors declare no competing interests.
- Ethics approval and consent to participate - Not applicable.
- Consent for publication - All authors consent to the publication of this manuscript.
- Data availability - The data will be made available after publication.
- Materials availability - Not applicable.
- Code availability - The code is available on Dr Bazhenov's lab website.
- Author contributions - S.J., S.H., Z.W., M.B. and Y.C. conceived and designed the research. S.J., S.H., F.L. and Z.W. performed the experiments. S.J. and S.H. analyzed the data. S.J. and S.H. drafted the manuscript. S.J. prepared the figures. S.J., S.H., Z.W., F.L., B.S., Y.C., M.B. edited and revised the manuscript.

References

Antonov, I., Antonova, I., Kandel, E.R., Hawkins, R.D.: Activity-dependent presynaptic facilitation and hebbian LTP are both required and interact during classical conditioning in Aplysia. Neuron **37**(1), 135–147 (2003) https://doi.org/10.1016/s0896-6273(02)01129-7

Abel, R., Rybak, J., Menzel, R.: Structure and response patterns of olfactory interneurons in the honeybee, Apis mellifera. Journal of Comparative Neurology **437**(3), 363–383 (2001) https://doi.org/10.1002/cne.1289 . _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/cne.1289. Accessed 2023-08-07

Aso, Y., Siwanowicz, I., Bräcker, L., Ito, K., Kitamoto, T., Tanimoto, H.: Specific dopaminergic neurons for the formation of labile aversive memory. Current biology: CB **20**(16), 1445–1451 (2010) https://doi.org/10.1016/j.cub.2010.06.048

Assisi, C., Stopfer, M., Bazhenov, M.: Using the structure of inhibitory networks to unravel mechanisms of spatiotemporal patterning. Neuron **69**(2), 373–386 (2011) https://doi.org/10.1016/j.neuron.2010.12.019 . Accessed 2023-08-16

Assisi, C., Stopfer, M., Bazhenov, M.: Optimality of sparse olfactory representations is not affected by network plasticity. PLOS Computational Biology **16**(2), 1007461 (2020) https://doi.org/10.1371/journal.pcbi.1007461 . Publisher: Public Library of Science. Accessed 2023-08-16

Bazhenov, M., Huerta, R., Smith, B.H.: A Computational Framework for Understanding Decision Making through Integration of Basic Learning Rules. The Journal of Neuroscience **33**(13), 5686–5697 (2013) https://doi.org/10.1523/JNEUROSCI.4145-12.2013 . Accessed 2023-08-07

Bicker, G., Kreissl, S.: Calcium imaging reveals nicotinic acetylcholine receptors on cultured mushroom body neurons. Journal of Neurophysiology **71**(2), 808–810 (1994) https://doi.org/10.1152/jn.1994.71.2.808 . Publisher: American Physiological Society. Accessed 2023-08-07

Bazhenov, M., Stopfer, M., Rabinovich, M., Abarbanel, H.D., Sejnowski, T.J., Laurent, G.: Model of cellular and network mechanisms for odor-evoked temporal patterning in the locust antennal lobe. Neuron **30**(2), 569–581 (2001) https://doi.org/10.1016/s0896-6273(01)00286-0

Bazhenov, M., Stopfer, M., Rabinovich, M.I., Huerta, R., Abarbanel, H.D.I., Sejnowski, T.J., Laurent, G.: Model of Transient Oscillatory Synchronization in the Locust Antennal Lobe. Neuron **30**(2), 553–567 (2001) https://doi.org/10.1016/s0896-6273(01)00284-7 . MAG ID: 2148452798

Bazhenov, M., Stopfer, M., Sejnowski, T.J., Laurent, G.: Fast Odor Learning Improves Reliability of Odor Responses in the Locust Antennal Lobe. Neuron **46**(3), 483–492 (2005) https://doi.org/10.1016/j.neuron.2005.03.022 . Accessed 2024-12-06

Capogna, M., Gähwiler, B.H., Thompson, S.M.: Presynaptic enhancement of inhibitory synaptic transmission by protein kinases A and C in the rat hippocampus in vitro. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience **15**(2), 1249–1260 (1995) https://doi.org/10.1523/JNEUROSCI.15-02-01249.1995

Chen, J.-Y., Marachlian, E., Assisi, C., Huerta, R., Smith, B.H., Locatelli, F., Bazhenov, M.: Learning Modifies Odor Mixture Processing to Improve Detection of Relevant Components. Journal of Neuroscience **35**(1), 179–197 (2015) https://doi.org/10.1523/JNEUROSCI.2345-14.2015 . Accessed 2022-04-26

Chandra, S.B.C., Wright, G.A., Smith, B.H.: Latent inhibition in the honey bee, Apis mellifera: Is it a unitary phenomenon? Animal Cognition **13**(6), 805–815 (2010) https://doi.org/10.1007/s10071-010-0329-6

Duvenaud, D., Maclaurin, D., Aguilera-Iparraguirre, J., Gómez-Bombarelli, R., Hirzel, T., Aspuru-Guzik, A., Adams, R.P.: Convolutional networks on graphs for learning molecular fingerprints. In: Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2. NIPS'15, pp. 2224–2232. MIT Press, Cambridge, MA, USA (2015)

Destexhe, A., Mainen, Z.F., Sejnowski, T.J.: Synthesis of models for excitable membranes, synaptic transmission and neuromodulation using a common kinetic formalism. Journal of Computational Neuroscience **1**(3), 195–230 (1994) https://doi.org/10.1007/BF00961734 . Accessed 2022-04-27

Das, S., Sadanandappa, M.K., Dervan, A., Larkin, A., Lee, J.A., Sudhakaran, I.P., Priya, R., Heidari, R., Holohan, E.E., Pimentel, A., Gandhi, A., Ito, K., Sanyal, S., Wang, J.W., Rodrigues, V., Ramaswami, M.: Plasticity of local GABAergic interneurons drives olfactory habituation. Proceedings of the National Academy of Sciences **108**(36), 646–654 (2011) https://doi.org/10.1073/pnas.1106411108 . Publisher: Proceedings of the National Academy of Sciences. Accessed 2023-09-15

Daly, K.C., Wright, G.A., Smith, B.H.: Molecular Features of Odorants Systematically Influence Slow Temporal Responses Across Clusters of Coordinated Antennal Lobe Units in the Moth *Manduca sexta*.

Journal of Neurophysiology **92**(1), 236–254 (2004) https://doi.org/10.1152/jn.01132.2003 . Accessed 2024-12-06

Finelli, L.A., Haney, S., Bazhenov, M., Stopfer, M., Sejnowski, T.J.: Synaptic Learning Rules and Sparse Coding in a Model Sensory System. PLOS Computational Biology **4**(4), 1000062 (2008) https://doi.org/10.1371/journal.pcbi.1000062 . Publisher: Public Library of Science. Accessed 2023-08-16

Faber, T., Joerges, J., Menzel, R.: Associative learning modifies neural representations of odors in the insect brain. Nature Neuroscience **2**(1), 74–78 (1999) https://doi.org/10.1038/4576 . Number: 1 Publisher: Nature Publishing Group. Accessed 2023-08-07

Fernandez, P.C., Locatelli, F.F., Person-Rennell, N., Deleo, G., Smith, B.H.: Associative conditioning tunes transient dynamics of early olfactory processing. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience **29**(33), 10191–10202 (2009) https://doi.org/10.1523/JNEUROSCI.1874-09.2009

Farooqui, T., Robinson, K., Vaessin, H., Smith, B.H.: Modulation of Early Olfactory Processing by an Octopaminergic Reinforcement Pathway in the Honeybee. The Journal of Neuroscience **23**(12), 5370–5380 (2003) https://doi.org/10.1523/JNEUROSCI.23-12-05370.2003 . Accessed 2023-08-08

Fulton, K.A., Zimmerman, D., Samuel, A., Vogt, K., Datta, S.R.: Common principles for odour coding across vertebrates and invertebrates. Nature Reviews Neuroscience **25**(7), 453–472 (2024) https://doi.org/10.1038/s41583-024-00822-0 . Accessed 2024-12-05

Galizia, C.G.: Olfactory coding in the insect brain: data and conjectures. The European Journal of Neuroscience **39**(11), 1784–1795 (2014) https://doi.org/10.1111/ejn.12558 . Accessed 2023-09-12

Grohmann, L., Blenau, W., Erber, J., Ebert, P.R., Strünker, T., Baumann, A.: Molecular and functional characterization of an octopamine receptor from honeybee (*Apis mellifera*) brain. Journal of Neurochemistry **86**(3), 725–735 (2003) https://doi.org/10.1046/j.1471-4159.2003.01876.x . Accessed 2024-12-05

Galizia, C.G., McIlwrath, S.L., Menzel, R.: A digital three-dimensional atlas of the honeybee antennal lobe based on optical sections acquired by confocal microscopy. Cell and Tissue Research **295**(3), 383–394 (1999) https://doi.org/10.1007/s004410051245

Gronenberg, W.: Physiological and anatomical properties of optical input-fibres to the mushroom body in the bee brain. Journal of Insect Physiology **32**(8), 695–704 (1986) https://doi.org/10.1016/0022-1910(86)90111-3 . Accessed 2023-09-13

Hammer, M.: An identified neuron mediates the unconditioned stimulus in associative olfactory learning in honeybees. Nature **366**(6450), 59–63 (1993) https://doi.org/10.1038/366059a0 . Number: 6450 Publisher: Nature Publishing Group. Accessed 2023-08-07

Hodgkin, A.L., Huxley, A.F.: A quantitative description of membrane current and its application to conduction and excitation in nerve. The Journal of Physiology **117**(4), 500–544 (1952). Accessed 2023-08-25

Hayes, T.L., Krishnan, G.P., Bazhenov, M., Siegelmann, H.T., Sejnowski, T.J., Kanan, C.: Replay in Deep Learning: Current Approaches and Missing Biological Elements. Neural Computation **33**(11), 2908–2950 (2021) https://doi.org/10.1162/neco_a_01433 https://direct.mit.edu/neco/article-pdf/33/11/2908/1966599/neco_a_01433.pdf

Hassabis, D., Kumaran, D., Summerfield, C., Botvinick, M.: Neuroscience-Inspired Artificial Intelligence. Neuron **95**(2), 245–258 (2017) https://doi.org/10.1016/j.neuron.2017.06.011 . Accessed 2023-02-24

Hammer, M., Menzel, R.: Learning and memory in the honeybee. Journal of Neuroscience **15**(3), 1617–1630 (1995) https://doi.org/10.1523/JNEUROSCI.15-03-01617.1995 . Publisher: Society for Neuroscience Section: Articles. Accessed 2023-08-07

Hammer, M., Menzel, R.: Multiple Sites of Associative Odor Learning as Revealed by Local Brain Microinjections of Octopamine in Honeybees. Learning & Memory **5**(1), 146–156 (1998). Accessed 2023-08-07

Han, K.-A., Millar, N.S., Davis, R.L.: A Novel Octopamine Receptor with Preferential Expression inDrosophila Mushroom Bodies. The Journal of Neuroscience **18**(10), 3650–3658 (1998) https://doi.org/10.1523/JNEUROSCI.18-10-03650.1998 . Accessed 2023-09-17

Hildebrand, J.G., Shepherd, G.M.: Mechanisms of olfactory discrimination: converging evidence for common principles across phyla. Annual Review of Neuroscience **20**, 595–631 (1997) https://doi.org/10.1146/annurev.neuro.20.1.595

Haney, S., Saha, D., Raman, B., Bazhenov, M.: Differential effects of adaptation on odor discrimination. Journal of Neurophysiology **120**(1), 171–185 (2018) https://doi.org/10.1152/jn.00389.2017 . Accessed 2022-04-26

Ito, I., Bazhenov, M., Ong, R.C.-y., Raman, B., Stopfer, M.: Frequency transitions in odor-evoked neural oscillations. Neuron **64**(5), 692–706 (2009) https://doi.org/10.1016/j.neuron.2009.10.004 . Accessed 2023-08-16

Kudithipudi, D., Aguilar-Simon, M., Babb, J., Bazhenov, M., Blackiston, D., Bongard, J., Brna, A.P., Chakravarthi Raja, S., Cheney, N., Clune, J., *et al.*: Biological underpinnings for lifelong learning machines. Nature Machine Intelligence **4**(3), 196–210 (2022)

Komischke, B., Giurfa, M., Lachnit, H., Malun, D.: Successive Olfactory Reversal Learning in Honeybees. Learning & Memory **9**(3), 122–129 (2002) https://doi.org/10.1101/lm.44602 . Accessed 2024-03-24

Kipf, T.N., Welling, M.: Semi-Supervised Classification with Graph Convolutional Networks (2016). https://arxiv.org/abs/1609.02907v4 Accessed 2023-08-16

Locatelli, F.F., Fernandez, P.C., Smith, B.H.: Learning about natural variation of odor mixtures enhances categorization in early olfactory processing. The Journal of Experimental Biology **219**(17), 2752–2762 (2016) https://doi.org/10.1242/jeb.141465 . Accessed 2021-07-12

Locatelli, F.F., Fernandez, P.C., Villareal, F., Muezzinoglu, K., Huerta, R., Galizia, C.G., Smith, B.H.: Nonassociative plasticity alters competitive interactions among mixture components in early olfactory processing. The European Journal of Neuroscience **37**(1), 63–79 (2013) https://doi.org/10.1111/ejn.12021

Larkin, A., Karak, S., Priya, R., Das, A., Ayyub, C., Ito, K., Rodrigues, V., Ramaswami, M.: Central synaptic mechanisms underlie short-term olfactory habituation in Drosophila larvae. Learning & Memory **17**(12), 645–653 (2010) https://doi.org/10.1101/lm.1839010 . Company: Cold Spring Harbor Laboratory Press Distributor: Cold Spring Harbor Laboratory Press Institution: Cold Spring Harbor Laboratory Press Label: Cold Spring Harbor Laboratory Press Publisher: Cold Spring Harbor Lab. Accessed 2023-08-07

Latshaw, J.S., Mazade, R.E., Petersen, M., Mustard, J.A., Sinakevitch, I., Wissler, L., Guo, X., Cook, C., Lei, H., Gadau, J., Smith, B.: Tyramine and its *Amtyr1* receptor modulate attention in honey bees (*Apis mellifera*). eLife **12**, 83348 (2023) https://doi.org/10.7554/eLife.83348

Laurent, G., MacLeod, K., Stopfer, M., Wehr, M.: Spatiotemporal Structure of Olfactory Inputs to the Mushroom Bodies. Learning & Memory **5**(1), 124–132 (1998). Accessed 2023-08-07

Linster, C., Menon, A.V., Singh, C.Y., Wilson, D.A.: Odor-specific habituation arises from interaction of afferent synaptic adaptation and intrinsic synaptic potentiation in olfactory cortex. Learning & Memory **16**(7), 452–459 (2009) https://doi.org/10.1101/lm.1403509 . Accessed 2023-08-07

Linster, C., Smith, B.H.: A computational model of the response of honey bee antennal lobe circuitry to odor mixtures: overshadowing, blocking and unblocking can arise from lateral inhibition. Behavioural

Brain Research **87**(1), 1–14 (1997) https://doi.org/10.1016/s0166-4328(96)02271-1

Lee, H.-G., Seong, C.-S., Kim, Y.-C., Davis, R.L., Han, K.-A.: Octopamine receptor OAMB is required for ovulation in Drosophila melanogaster. Developmental Biology **264**(1), 179–190 (2003) https://doi.org/10.1016/j.ydbio.2003.07.018 . Accessed 2024-12-05

Menzel, R.: Learning, memory, and "cognition" in honey bees. In: Neurobiology of Comparative Cognition. Comparative cognition and neuroscience, pp. 237–292. Lawrence Erlbaum Associates, Inc, Hillsdale, NJ, US (1990)

Marachlian, E., Klappenbach, M., Locatelli, F.: Learning-dependent plasticity in the antennal lobe improves discrimination and recognition of odors in the honeybee. Cell and Tissue Research **383**(1), 165–175 (2021) https://doi.org/10.1007/s00441-020-03396-2 . Accessed 2024-12-06

MaBouDi, H., Shimazaki, H., Giurfa, M., Chittka, L.: Olfactory learning without the mushroom bodies: Spiking neural network models of the honeybee lateral antennal lobe tract reveal its capacities in odour memory tasks of varied complexities. PLOS Computational Biology **13**(6), 1005551 (2017) https://doi.org/10.1371/journal.pcbi.1005551 . Accessed 2024-12-05

Mobbs, P.G., Young, J.Z.: The brain of the honeybee Apis mellifera. I. The connections and spatial organization of the mushroom bodies. Philosophical Transactions of the Royal Society of London. B, Biological Sciences **298**(1091), 309–354 (1997) https://doi.org/10.1098/rstb.1982.0086 . Publisher: Royal Society. Accessed 2023-09-13

Müller, U.: Learning in honeybees: from molecules to behaviour. Zoology (Jena, Germany) **105**(4), 313–320 (2002) https://doi.org/10.1078/0944-2006-00075

Negre, F., Kish, C.M., Boatright, J., Underwood, B., Shibuya, K., Wagner, C., Clark, D.G., Dudareva, N.: Regulation of Methylbenzoate Emission after Pollination in Snapdragon and Petunia Flowers. The Plant Cell **15**(12), 2992–3006 (2003) https://doi.org/10.1105/tpc.016766 . Accessed 2023-08-08

Pavlov, I.P.: Lectures on Conditioned Reflexes...: Twenty-five Years of Objective Study of the Higher Nervous Activity (behaviour) of Animals... Translated from the Russian by WH Gantt... with the Collaboration of G. Volborth... and an Introduction by WB Cannon vol. 1. International publishers, ??? (1928)

Pichersky, E., Gershenzon, J.: The formation and function of plant volatiles: perfumes for pollinator attraction and defense. Current Opinion in Plant Biology **5**(3), 237–243 (2002) https://doi.org/10.1016/S1369-5266(02)00251-0 . Accessed 2023-08-07

Paoli, M., Galizia, G.C.: Olfactory coding in honeybees. Cell and Tissue Research **383**(1), 35–58 (2021) https://doi.org/10.1007/s00441-020-03385-5 . Accessed 2023-09-12

Perez-Orive, J., Mazor, O., Turner, G.C., Cassenaer, S., Wilson, R.I., Laurent, G.: Oscillations and sparsening of odor representations in the mushroom body. Science (New York, N.Y.) **297**(5580), 359–365 (2002) https://doi.org/10.1126/science.1070502

Rajagopalan, A., Assisi, C.: Effect of Circuit Structure on Odor Representation in the Insect Olfactory System. eNeuro **7**(3) (2020) https://doi.org/10.1523/ENEURO.0130-19.2020 . Publisher: Society for Neuroscience Section: Research Article: New Research. Accessed 2023-09-13

Rein, J., Mustard, J.A., Strauch, M., Smith, B.H., Galizia, C.G.: Octopamine modulates activity of neural networks in the honey bee antennal lobe. Journal of Comparative Physiology. A, Neuroethology, Sensory, Neural, and Behavioral Physiology **199**(11), 947–962 (2013) https://doi.org/10.1007/s00359-013-0805-y . Accessed 2023-08-07

Reiser, P., Neubert, M., Eberhard, A., Torresi, L., Zhou, C., Shao, C., Metni, H., Hoesel, C., Schopmans, H., Sommer, T., Friederich, P.: Graph neural networks for materials science and chemistry. Communications Materials **3**(1), 1–18 (2022) https://doi.org/10.1038/s43246-022-00315-6 . Number: 1 Publisher:

Nature Publishing Group. Accessed 2023-08-16

Sandoz, J.-C.: [olfactory perception and learning in the honey bee (apis mellifera): calcium imaging in the antenna lobe]. Journal de la Societe de biologie **197**(3), 277–282 (2003)

Smith, B.H., Burden, C.M.: A proboscis extension response protocol for investigating behavioral plasticity in insects: application to basic, biomedical, and agricultural research. Journal of Visualized Experiments: JoVE (91), 51057 (2014) https://doi.org/10.3791/51057

Stopfer, M., Bhagavan, S., Smith, B.H., Laurent, G.: Impaired odour discrimination on desynchronization of odour-encoding neural assemblies. Nature **390**(6655), 70–74 (1997) https://doi.org/10.1038/36335 . Accessed 2024-06-12

Sachse, S., Galizia, C.G.: Role of Inhibition for Temporal and Spatial Odor Representation in Olfactory Output Neurons: A Calcium Imaging Study. Journal of Neurophysiology **87**(2), 1106–1117 (2002) https://doi.org/10.1152/jn.00325.2001 . Publisher: American Physiological Society. Accessed 2023-08-07

Szyszka, P., Galkin, A., Menzel, R.: Associative and non-associative plasticity in Kenyon cells of the honeybee mushroom body. Frontiers in Systems Neuroscience **2** (2008). Accessed 2023-09-13

Sudhakaran, I.P., Holohan, E.E., Osman, S., Rodrigues, V., VijayRaghavan, K., Ramaswami, M.: Plasticity of Recurrent Inhibition in the Drosophila Antennal Lobe. Journal of Neuroscience **32**(21), 7225–7231 (2012) https://doi.org/10.1523/JNEUROSCI.1099-12.2012 . Publisher: Society for Neuroscience Section: Articles. Accessed 2023-03-30

Stopfer, M., Jayaraman, V., Laurent, G.: Intensity versus identity coding in an olfactory system. Neuron **39**(6), 991–1004 (2003) https://doi.org/10.1016/j.neuron.2003.08.011

Sanda, P., Kee, T., Gupta, N., Stopfer, M., Bazhenov, M.: Classification of odorants across layers in locust olfactory pathway. Journal of Neurophysiology **115**(5), 2303–2316 (2016) https://doi.org/10.1152/jn.00921.2015 . Accessed 2023-08-16

Stopfer, M., Laurent, G.: Short-term memory in olfactory network dynamics. Nature **402**(6762), 664–668 (1999) https://doi.org/10.1038/45244

Sanchez-Lengeling, B., Wei, J.N., Lee, B.K., Gerkin, R.C., Aspuru-Guzik, A., Wiltschko, A.B.: Machine Learning for Scent: Learning Generalizable Perceptual Representations of Small Molecules. arXiv. arXiv:1910.10685 [physics, stat] (2019). http://arxiv.org/abs/1910.10685 Accessed 2023-07-25

Sinakevitch, I., Mustard, J.A., Smith, B.H.: Distribution of the octopamine receptor AmOA1 in the honey bee brain. PloS One **6**(1), 14536 (2011) https://doi.org/10.1371/journal.pone.0014536

Sinakevitch, I., Niwa, M., Strausfeld, N.J.: Octopamine-like immunoreactivity in the honey bee and cockroach: Comparable organization in the brain and subesophageal ganglion. Journal of Comparative Neurology **488**(3), 233–254 (2005) https://doi.org/10.1002/cne.20572 . _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/cne.20572. Accessed 2023-08-07

Sinakevitch, I., Smith, A., Locatelli, F., Huerta, R., Bazhenov, M., Smith, B.: Apis mellifera octopamine receptor 1 (AmOA1) expression in antennal lobe networks of the honey bee (Apis mellifera) and fruit fly (Drosophila melanogaster). Frontiers in Systems Neuroscience **7** (2013). Accessed 2023-08-16

Smith, B.H., Wright, G.A., Daly, K.C.: Learning-Based Recognition and Discrimination of Floral Odors, pp. 263–295 (2006). https://doi.org/10.1201/9781420004007.ch12

Thompson, R.F., Spencer, W.A.: Habituation: A model phenomenon for the study of neuronal substrates of behavior. Psychological Review **73**(1), 16–43 (1966) https://doi.org/10.1037/h0022681 . Place: US Publisher: American Psychological Association

Vergara, A., Vembu, S., Ayhan, T., Ryan, M.A., Homer, M.L., Huerta, R.: Chemical gas sensor drift

compensation using classifier ensembles. Sensors and Actuators B: Chemical **166-167**, 320–329 (2012) https://doi.org/10.1016/j.snb.2012.01.074 . Accessed 2023-07-18

Wilson, D.A., Linster, C.: Neurobiology of a simple memory. Journal of Neurophysiology **100**(1), 2–7 (2008) https://doi.org/10.1152/jn.90479.2008

Wright, G.A., Lutmerding, A., Dudareva, N., Smith, B.H.: Intensity and the ratios of compounds in the scent of snapdragon flowers affect scent discrimination by honeybees (Apis mellifera). Journal of Comparative Physiology A **191**(2), 105–114 (2005) https://doi.org/10.1007/s00359-004-0576-6 . Accessed 2022-04-26

Wilson, R.I., Turner, G.C., Laurent, G.: Transformation of Olfactory Representations in the Drosophila Antennal Lobe. Science **303**(5656), 366–370 (2004) https://doi.org/10.1126/science.1090782 . Publisher: American Association for the Advancement of Science. Accessed 2023-08-07

Supplementary figures

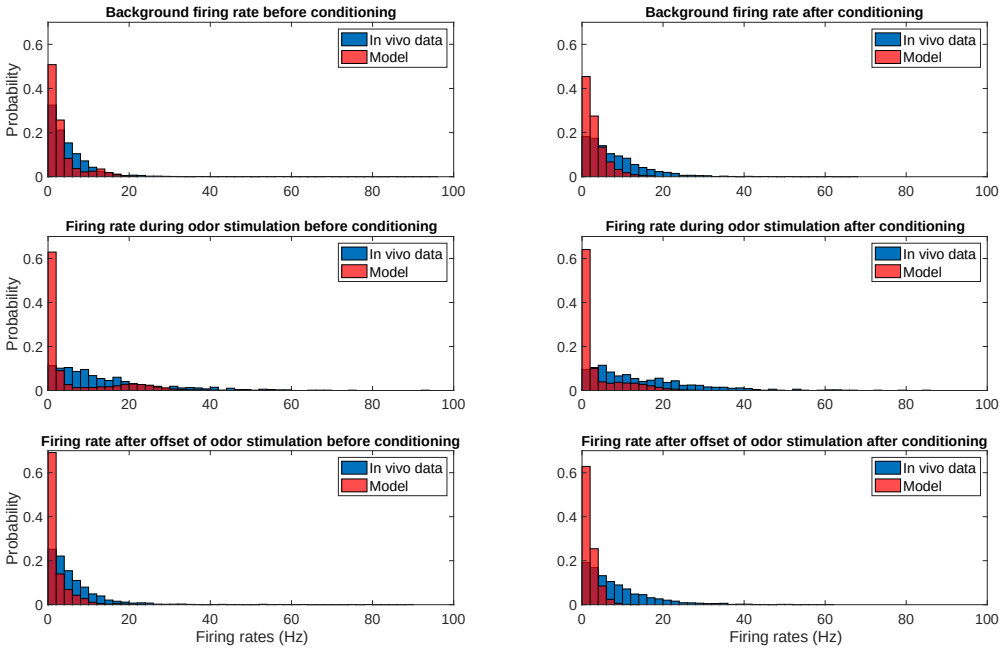

Fig. 1 Histograms of firing rates for in vivo and model PN responses for before and after conditioning. The three rows show firing rates calculated during different time periods. Row 1 shows firing rates for background activity. Row 2 shows firing rates calculated during 500ms of odor stimulation. Row 3 shows firing rates calculated 1s after the offset of the odor stimulation.

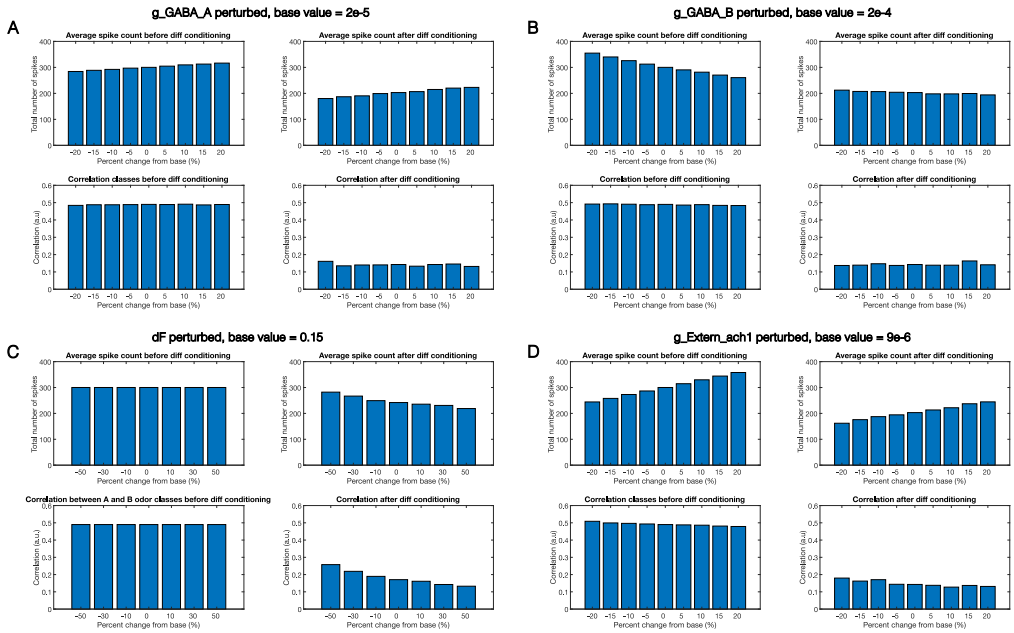

Fig. 2 Parameter sensitivity analysis demonstrates robustness of model behavior. (A-D) Each panel shows the effect of perturbing a different parameter: (A) fast GABA inhibition (g_GABA_A, base value = 2.0e-5), (B) slow GABA inhibition (g_GABA_B, base value = 2.0e-4), (C) facilitation rate (dF, base value = 1.5e-1), and (D) input synaptic conductance (g_Extern_ach1, base value = 9.0e-6). For each parameter, the top row shows the average spike count across all PNs before (left) and after (right) differential conditioning, while the bottom row shows the correlation between odor classes A and B before (left) and after (right) differential conditioning. The x-axis represents the percent change from the baseline parameter value. Despite variations in absolute spike counts, the key finding of reduced correlation between odor representations after differential conditioning remains robust across parameter perturbations, indicating that the model's core behavior does not depend on precise parameter tuning.

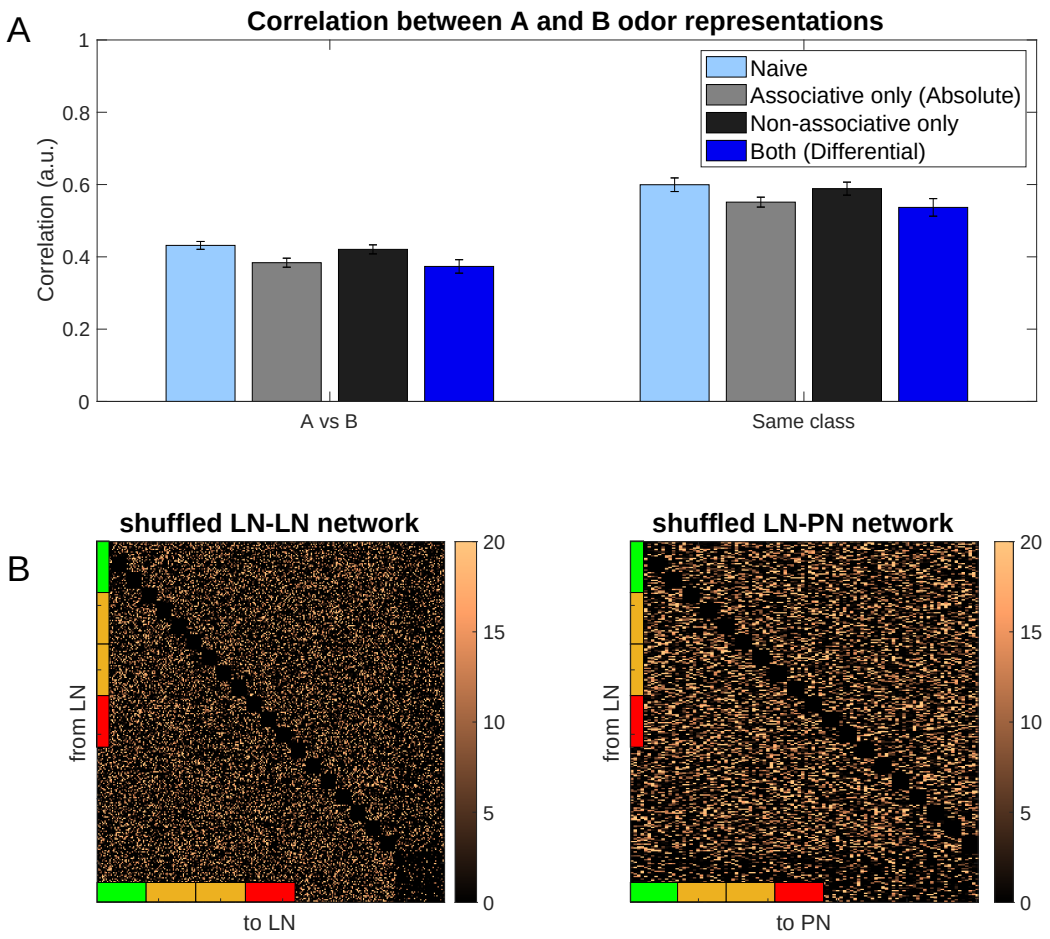

Fig. 3 Analysis of network connectivity structure and odor discrimination. A) Correlation between odor representations comparing naive, associative-only, non-associative-only, and combined (differential) conditioning for odors between different classes (left) and for odors within the same class (right). B) Connectivity matrices showing shuffled LN-LN and LN-PN networks after differential conditioning. When the structured connectivity patterns are disrupted through shuffling while maintaining the same total inhibitory strength, the enhanced odor discrimination is lost, demonstrating that the specific spatial organization of inhibitory connections, not just increased inhibition, is essential for improved odor discrimination. Side bars indicate percept identities (green: unique to rewarded odor, orange: shared, red: unique to habituated odor). Color bars show relative synaptic weights.