Report

Current Biology

A Simple Computational Model of the Bee Mushroom Body Can Explain Seemingly Complex Forms of Olfactory Learning and Memory

Highlights

- We model the olfactory neural circuitry of the honeybee brain
- Model outputs reproduce peak shift, positive and negative patterning discrimination
- We identify synapses that can control the generalization-discrimination trade-off
- Configural coding capacity is dependent on Kenyon cell activation sparseness level

Authors

Fei Peng, Lars Chittka

Correspondence

l.chittka@qmul.ac.uk

In Brief

Peak shift, positive and negative patterning are perceived as relatively complex forms of learning, involving the integration of multiple stimuli with opposite meanings. Peng and Chittka show that a neuronal model of the bee mushroom body circuitry with simple learning rules can account for these multiple forms of learning in the olfactory domain.

Peng & Chittka, 2017, Current Biology *27*, 224–230
January 23, 2017
http://dx.doi.org/10.1016/j.cub.2016.10.054

CellPress

CellPress

Current Biology
Report

A Simple Computational Model of the Bee Mushroom Body Can Explain Seemingly Complex Forms of Olfactory Learning and Memory

Fei Peng[1] and Lars Chittka[1,2,*]
[1]Biological and Experimental Psychology, School of Biological and Chemical Sciences, Queen Mary University of London, London E1 4NS, UK
[2]Lead Contact
*Correspondence: l.chittka@qmul.ac.uk
http://dx.doi.org/10.1016/j.cub.2016.10.054

SUMMARY

Honeybees are models for studying how animals with relatively small brains accomplish complex cognition, displaying seemingly advanced (or "non-elemental") learning phenomena involving multiple conditioned stimuli. These include "peak shift" [1–4]—where animals not only respond to entrained stimuli, but respond even more strongly to similar ones that are farther away from non-rewarding stimuli. Bees also display negative and positive patterning discrimination [5], responding in opposite ways to mixtures of two odors than to individual odors. Since Pavlov, it has often been assumed that such phenomena are more complex than simple associate learning. We present a model of connections between olfactory sensory input and bees' mushroom bodies [6], incorporating empirically determined properties of mushroom body circuitry (random connectivity [7], sparse coding [8], and synaptic plasticity [9, 10]). We chose not to optimize the model's parameters to replicate specific behavioral phenomena, because we were interested in the emergent cognitive capacities that would pop out of a network constructed solely based on empirical neuroscientific information and plausible assumptions for unknown parameters. We demonstrate that the circuitry mediating "simple" associative learning can also replicate the various non-elemental forms of learning mentioned above and can effectively multi-task by replicating a range of different learning feats. We found that PN-KC synaptic plasticity is crucial in controlling the generalization-discrimination trade-off—it facilitates peak shift and hinders patterning discrimination—and that PN-to-KC connection number can affect this trade-off. These findings question the notion that forms of learning that have been regarded as "higher order" are computationally more complex than "simple" associative learning.

RESULTS AND DISCUSSION

Based on structural and functional characteristics of the insect olfactory pathway, we built a three-layer neuronal network model (Figure 1A) to test whether such simple circuits can reproduce empirical behavioral results. The bee olfactory pathway (Figure S1) recruits a divergence-convergence structure where ~800–900 projection neurons (PNs) expand onto ~170,000 Kenyon cells (KCs), and these are then read out by ~400 mushroom body extrinsic neurons (ENs). The connections between projection neurons and Kenyon cells are relatively sparse: each Kenyon cell is thought to be innervated by approximately ten projection neurons [8], and connections are random in *Drosophila* [7], a property that we adopt here for bees. With a biologically realistic PN-KC neuronal number ratio (1:40, ~500 lateral antennal lobe tract projection neurons [11] onto ~20,000 clawed Kenyon cells [8]), we generate a random connectivity matrix; each Kenyon cell receives input from five to 15 projection neurons. This connectivity can transform olfactory inputs into sparse representations—experimental estimates for *Drosophila* [12] indicate that ~5% of Kenyon cells are activated for any given stimulus. Moreover, a putative GABAergic feedback inhibitory pathway onto mushroom body calyces [13] may control the sparseness level of the Kenyon cell population, as suggested for locusts [14] and *Drosophila* [15]. We simulate feedback inhibition by selecting 5% of Kenyon cells that receive the largest summed inputs and we label these Kenyon cells as activated for each stimulus (Supplemental Experimental Procedures). To verify the efficiency of this implementation, we compare the sets of activated Kenyon cells so generated with the ones by a spiking network with a global feedback inhibition, showing that the two are almost identical (Figure S2A), indicating that this implementation can adequately substitute a computationally more complex spiking model. We apply a fixed level of feedback inhibition in the model; it is, however, possible that this feedback pathway is also subject to learning-related plasticity [16, 17].

The convergence between the putative reward pathway of the VUMmx1 neuron [18] and the PN-KC synapses suggests the possibility of learning-induced synaptic changes [9]. Mushroom body extrinsic neurons encode stimulus valence [10, 19], and KC-EN synapses display learning-related plasticity [20, 21]. Thus, we expand previous models that focus on KC-EN plastic synapses [22–25] by implementing plasticity among PN-KC

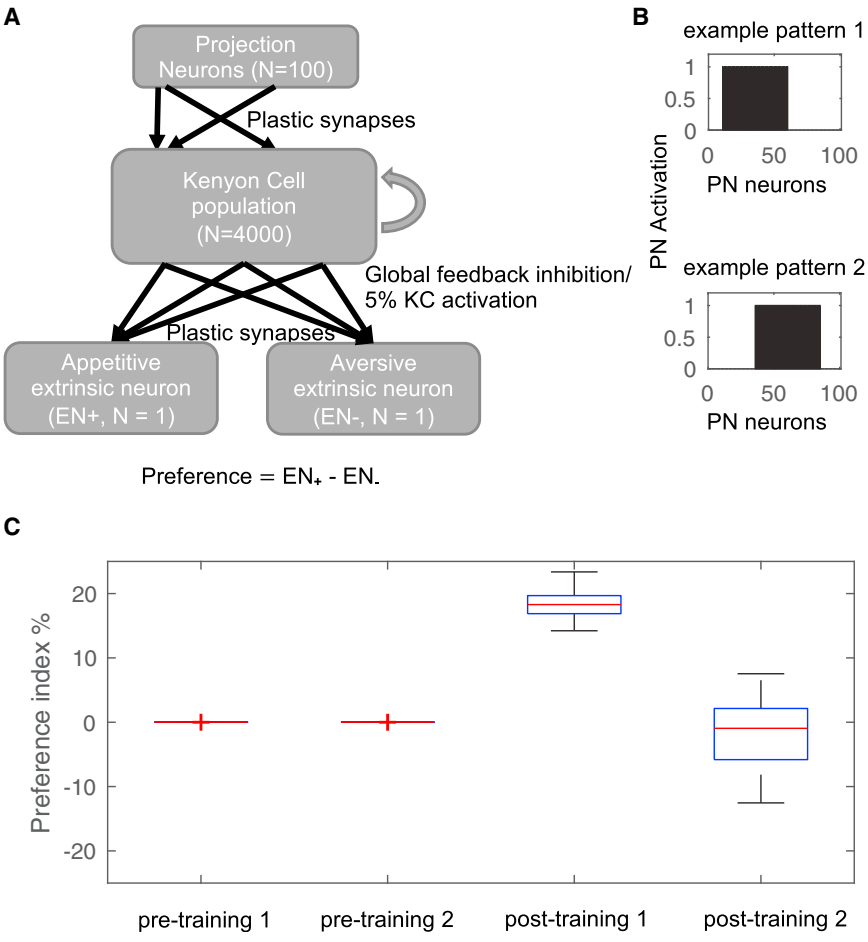

Figure 1. The Simple Model of the Bee Olfactory Pathway and Elemental Associative Learning

(A) The model has a three-layer feedforward architecture, with feedback inhibition onto Kenyon cells. PN-KC and KC-EN synapses both undergo associative learning. Feedback inhibition controls Kenyon cells populational sparseness to generate 5% activated Kenyon cells toward any given stimulus. The two model extrinsic neurons, each has an all-to-one connection from Kenyon cells, are encoding appetitive (EN_+) and aversive valance (EN_-), respectively. A preference is generated by a simple subtraction of the responses of the two extrinsic neurons [19].

(B) Example projection neuron firing patterns as artificial inputs to the model. The x axis shows the projection neuron population space, i.e., all of the projection neurons, and the y axis gives the profile for each projection neuron—a combination of binary activated projection neurons (intensity = 1) and inactivated projection neurons (intensity = 0).

(C) The preference index is generated by simple addition of the two extrinsic neuron responses with opposite directions of valence. Overall a minus sign is added to the preference index, such that, for example, a decreased response to appetitive stimuli now would have a positive preference—the same direction toward behavioral bias (see also the Supplemental Experimental Procedures). Before training (pre-training 1, pre-training 2), the model's naive responses toward example patterns 1 and 2 are tested and shown as 0%. After differential training with the two example patterns (post-training 1, post-training 2), the model preference toward both example pattern 1 (CS+) and example pattern 2 (CS−) is significantly different from baseline. Data for all four cases are represented as mean responses over 100 network realizations (or 100 virtual bees) ± SD.

See also Figures S1 and S2.

and KC-EN synapses. We use a reward-based synaptic weight modification rule, such that for PN-KC synapses, if a stimulus is rewarded (CS+), the corresponding synapses between activated neurons will be strengthened; for a stimulus paired with punishment (CS−), activated synapses are weakened (Supplemental Experimental Procedures). For KC-EN synapses, the opposite is applied, as decreased extrinsic neuron response for rewarded stimulus is consistently shown in bees [20], flies [26], and locusts [21]; it is, however, worth noting that in bees, the PE1 neuron could respond decreasingly [20], whereas non-PE1 extrinsic neurons could respond increasingly [10], to CS+. Aversive learning in *Drosophila* induces synaptic depression [27], but among distinct KC-EN synapses other than the ones used for appetitive valence encoding, suggesting a separation of appetitive and aversive valence encoding among those extrinsic neurons [19]. Moreover, the insect's final behavioral decision is proposed to come from a simple integration of those different valence-encoding extrinsic neurons [19]. We thus introduce two extrinsic neurons (EN_+ and EN_-), one for appetitive and another for aversive value encoding. The simple learning rule for PN-KC synapses is modified here, such that rewarded and punished stimuli will lead to synaptic depression among KC-EN_+ and KC-EN_-, respectively. Following [19], a preference index is introduced to measure the learned response of the model.

With equal weights initially assigned for PN-KC and KC-EN synapses and a fixed Kenyon cell activation level, without learning the model scores 0% in the preference index, as the valence for positive and negative is balanced out. Plasticity will then skew the model preference, and, assuming an inhibitory link between extrinsic neurons and the motor response [20], we view this change as a direct reflection of the final behavioral response. In other words, rewarding stimuli induce decreased extrinsic neuron responses, but finally higher preference, and vice versa for punished stimuli.

Model Input and Output

Olfactory projection neuron activation patterns forming inputs to the model are generated based on two empirical characteristics. First, honeybee olfactory inputs to antennal lobes are encoded as combinatorial glomerular activation patterns [28]. Similarly, olfactory projection neurons permeating the mushroom bodies also show a combinatorial spatial-temporal response profile [29]. Second, ~50% of projection neurons are responsive to each odor [30, 31]. In terms of encoding odor identity, Kenyon

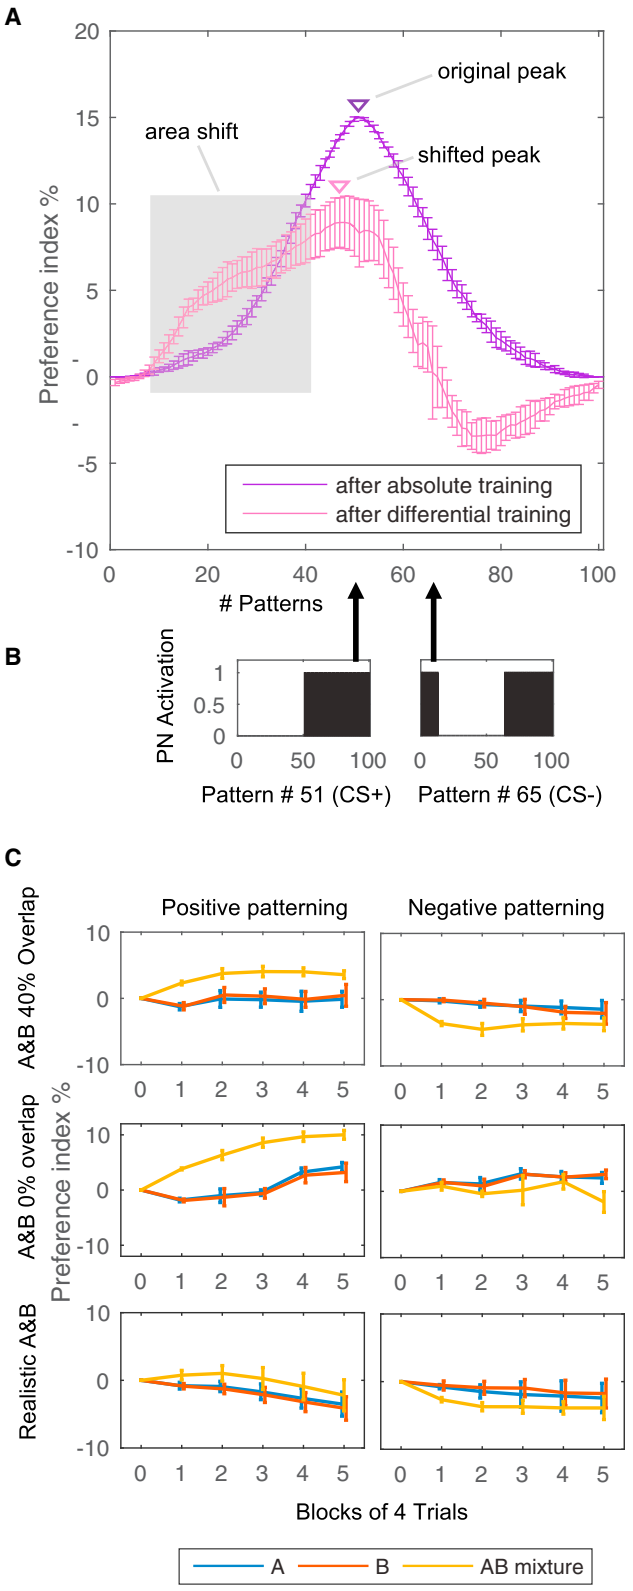

Figure 2. Model Exhibits Peak Shift and Positive and Negative Patterning

(A) Peak shift phenomenon generated by the model. Two groups, each containing 100 network realizations (or 100 virtual bees; mean ± SD), have been recorded. The magenta line represents the preference index of the absolute training group with stimulus 51 (CS+), and the pink line represents the preference index of differential training group on stimuli 51 (CS+) and 65 (CS−). The response curve (shown as mean ± SD) peaked after differential training with a novel position that is slightly farther away from CS− than CS+, similar to the curve in [1]. The gray area points out the so-called area shift [1].

(B) Two input patterns, CS+ (pattern 51) and CS− (pattern 65), were used in the training and test.

(C) The model learns to solve both positive patterning and negative patterning tasks (shown as mean of 100 network realizations ± SD). Top: model response for artificial inputs A and B, where these inputs have a moderate level (40%) of overlap. Middle: model response for artificial input A and B, where input A and B have no overlap (0%) at all. Bottom: model response for realistic projection neuron inputs [41] with normalization. Each panel contains five blocks of training, with each block containing four learning trials as follows: one exposure to A, one to B, and two to the AB mixture in pseudo-random order. Yellow lines and error bars indicate model preference toward the AB mixture, while red and blue represent the preference toward A and B, respectively. Note that red and blue curves are practically indistinguishable in several of the panels. After five blocks of trials, the extrinsic neuron response for A/B is significantly different from the AB mixture and from baseline in both positive and negative patterning, in all six cases (paired-sample t test, $p < 0.001$). For training and testing procedure of peak shift and patterning tasks, see the Supplemental Experimental Procedures.

See also Figures S2 and S3.

cells might only use the spatial, rather than the temporal, component of projection neuron activity since Kenyon cells respond in a temporally sharpened manner [8]. It is thus plausible to model projection neuron patterns simply as a binary combinatorial code, and this simplification allows us to have precise knowledge about the similarity between different inputs. In reality, projection neuron firing patterns can be more specific and selective, and thus we recruit both binary and more complex projection neuron patterns for model testing. Following [32], we model different inputs as partially overlapping projection neuron binary activation patterns (Figure 1B). The resulting similarity between stimuli is therefore proportional to the number of shared projection neurons that are activated by each stimulus; this is in accordance with evidence that similarity of antennal lobe activity patterns correlates with generalization in odor learning [33]. As a result of training, the model preference to CS+ and to CS− is changed significantly (Figure 1C; paired-sample t test, $t_{99,\ CS+} = 95.4028$, $t_{99,\ CS-} = -4.0262$, $p < 0.001$).

Model Replicates Olfactory Peak Shift and Patterning Tasks

Honeybees, like other animals [34], display peak shift in olfactory learning on a similarity continuum by mixtures of two odorants with different ratios [1, 3]. We assume the similarity in odors is encoded as partially overlapping projection neuron activation patterns. A set of artificial olfactory projection neuron patterns that constitutes a similarity continuum was generated (Figure S2B). Trained with absolute and differential training procedures (Supplemental Experimental Procedures) for CS+ with reward and CS− with punishment (Figure 2B), the model is tested with all of the stimuli patterns in the input continuum. The model response closely resembles peak shift (Figure 2A) in honeybees [1]: after differential training, the model responds maximally to novel stimuli farther away from CS− than CS+, while after absolute training the peak response occurs at CS+

(also see Figure S2C). There is also broadened generalization toward the opposite side of CS−, the so-called area shift [1, 35]. Peak shift has been interpreted as a form of relationship learning [4] or uncertainty coding [2], but a standard artificial neural network model with back-propagation can reproduce this phenomenon, indicating an emergent property of a simple network [36]. We confirm this with a model based on neurobiological evidence.

We then tested the model with positive and negative patterning (equivalent to the exclusive-or/XOR problem [37]: one or the other, but not both), following the same experimental procedures as [5]. According to [38], the olfactory compound input pattern to the honeybee antennal lobe is close to the arithmetic sum of the independent input patterns by the two odors. Moreover, odor mixture responses among projection neurons show only a slight deviation from linearity [30, 39]. We here assume the projection neuron activation pattern for an AB mixture is a linear summation of A and B odor, such that the overlapping projection neurons for an AB mixture respond additively and that the task cannot be solved by classical conditioning, e.g., the Rescorla-Wagner rule [40]. Hence, the differentiation for input A or B from its binary mixture will be a reflection of the model's capacity in solving this "non-elemental" problem. We employed two sets of input patterns that had different similarity and a final set of realistic projection neuron firing patterns [41], which were used independently for training and testing of the patterning discrimination (Supplemental Experimental Procedures). In each block of training, a pseudo-random sequence of A, B (once in each block), and the AB mixture (twice in each block) was applied (Figure 2C). In all cases, the model reaches similar levels of preferences for inputs A and B and successfully discriminates A and B from their binary mixture ($p < 0.001$ in paired-sample t test for A versus AB and B versus AB). Our model exhibits positive and negative patterning discrimination, which confirms previous theoretical work [23] and shows that it can robustly deal with inputs that have various similarity levels, as well as with realistic projection neuron firing patterns. Note that the model was not optimized to replicate the above learning phenomena. Instead, peak shift and patterning discrimination emerge as the direct consequence of the model structure, which was derived from known anatomy and physiology, and from requiring parameters, e.g., learning rates, to be in a neurobiologically plausible range.

Having shown that the model can account for these behaviors, we investigated which particular features of the neuronal circuitry could explain these phenomena. The insect mushroom body circuitry resembles three-layer perceptron [37, 42] and has been proposed to work similarly as the support vector machine [43]. Thus, the transformation of dense codes among projection neurons with linear mixture summation into a much higher dimensional representation by the Kenyon cells is likely to be the key for configural coding and non-linear problem solving. We here test whether different sparseness levels among Kenyon cells for given stimuli affects the model's ability to solve patterning tasks. We show that a high sparseness level among Kenyon cells is indeed necessary to separate overlapping input, generating an inbuilt ability to solve the seemingly complex "non-linear" patterning tasks through simple associative learning (Figure S3). Interestingly, a recent study found that the insect mushroom body GABAergic feedback pathway (the putative source of inhibitory inputs to Kenyon cells [15]) is crucial for solving so-called "non-elemental" forms of learning [5]. It is thus possible that the necessity of GABAergic feedback lies in modulating Kenyon cell population sparseness and so affecting configural coding capacity.

Plasticity in PN-KC Synapses Affects Peak Shift and Patterning Task in Opposite Directions

To explore how learning among PN-KC and KC-EN synapses contributes to performance in peak shift and patterning discrimination, we trained the model with different sets of learning rates (Supplemental Experimental Procedures). Positive rate here refers to the learning rate for CS+ and negative rate refers to learning for CS−, for modifying the activated synapses. Beginning with a fixed positive rate (0.006) for PN-KC and KC-EN synapses and various negative rates, heatmaps for peak shift are generated by detecting the proportion of groups (ten groups in total, each trained with ten virtual bees) that have peak responses to novel stimuli that are farther away from CS− than CS+ and are significantly higher ($p < 0.05$) than the responses to CS+ (Figure 3A). Similarly, heatmaps for patterning are made by detecting the proportion of groups (ten group in total, each trained with ten virtual bees on inputs of various similarities) that can significantly differentiate ($p < 0.05$) single odors from their binary mixture. Thus, performance under each set of learning rates is shown as a probability of successful behavioral reproduction. We found that a large negative rate (relative to the positive rate) among KC-EN synapses is critical to induce peak shift, consistent with experimental findings [1]. Moreover, when we fix the PN-KC synapses but implement plasticity in KC-EN, the model can no longer reproduce peak shift (Figure 3B); in contrast, for positive and negative patterning, almost all of the sets can induce good patterning discrimination, but simulations with plastic PN-KC synapses actually perform slightly worse than with fixed PN-KC synapses. This suggests that plasticity in PN-KC synapses facilitates peak shift but may hinder patterning discrimination.

Effects of Classes of Kenyon Cells on Peak Shift and Patterning Tasks

Honeybee mushroom bodies contain two classes of intrinsic cells: class I, or spiny Kenyon cells with wide-field dendritic arborizations, and class II, or clawed Kenyon cells with small-field arborizations [44]. Above we have modeled class II. Since class I is less researched in bees, we substitute information from locusts, where class I Kenyon cells [44, 45] may receive input from ~50% of the projection neuron population [46]. We build a variety of models, each having only one specific number of PN-to-KC connections, from 1% to 70% of all 100 projection neurons. Using the Kenyon cell activation for pattern 51 (in the middle of the input continuum) as the reference, we calculate the Kenyon cell activation similarity between the reference pattern and the entire input continuum (Figure S4A) for each of the different models. Before learning, class I Kenyon cell models (PN-to-KC = 45–55) show good specificity, in agreement with previous suggestions [46] (Figure S4A, the "naïve" case); absolute training on class I models leads to a broad generalization to novel input patterns, and differential training

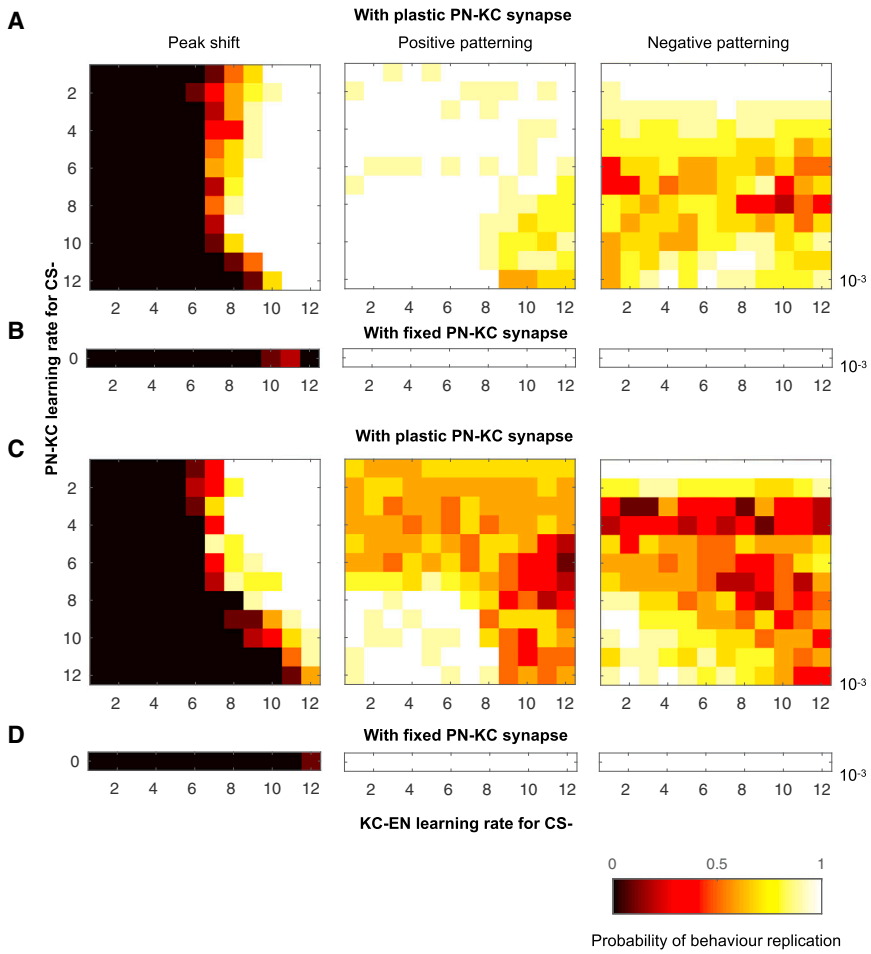

Figure 3. Model Performance on Peak Shift and Patterning with a Different Set of Learning Rates

(A) With plastic PN-KC synapses, three heatmaps show the probability of the model in reproducing peak shift and in significantly solving positive and negative patterning under a different set of learning rates, respectively. Note that the positive learning rate for both PN-KC and KC-EN synapses is fixed at 0.006, and the negative learning rate for the two levels of synapses varies from 0.001 to 0.012 with an increment of 0.001. The probability for peak shift is calculated as the frequency of groups of virtual bees (of ten groups) that display significant peak shift as follows: peak responses to a novel response farther away from CS− than CS+, and peak responses that are significantly higher than the responses to CS+ (paired-sample t test with a significance level of 0.05). In positive and negative patterning, each square indicates the frequency of groups of virtual bees (of ten groups) that can significantly (paired-sample t test with a significance level of 0.05) discriminate single odor patterns from the binary mixture.

(B) The same as (A), but with fixed PN-KC synapses, i.e., no learning among PN-KC is allowed.

(C) The probability of peak shift and patterning reproduction in class I Kenyon cell model with plastic PN-KC synapses, where each Kenyon cell receives input from ~50% of all projection neurons (45–55 out of 100 projection neurons).

(D) The same as (C), but with fixed PN-KC synapses.

See also Figure S4.

results in moderate generalization (Figure S4A, "abs" and "dif"). In contrast, class II models (PN-to-KC = 5–15) show consistently narrow generalization to novel patterns both before and after learning. The area under the response similarity curve over the input continuum for each model (Figure S4B) confirms this, suggesting that class I Kenyon cells may facilitate generalization, while class II Kenyon cells may be more reliable for discrimination.

We thus predict that a model with a connectivity scheme of ~50% projection neuron per Kenyon cell is likely to do well in peak shift and not so well in patterning discrimination, as the two tasks require generalization and fine discrimination capacity, respectively. Indeed, the class I model perform similarly in peak shift and clearly worse in patterning tasks (Figure 3C) than the original class II model (Figure 3A). The comparison between the class I model performance with and without PN-KC synaptic plasticity (Figures 3C and 3D) also confirms that PN-KC plasticity contributes to peak shift and patterning discrimination differently. We also test performance of models with different PN-to-KC connection numbers for both peak shift and patterning discrimination (Figure S4C), finding that with increasing PN-to-KC connection numbers, the probability of peak shift increases while the probability of both positive and negative patterning decreases.

PN-KC Plasticity and PN-to-KC Connection Number Can Control the Generalization-Discrimination Trade-off

Sparse coding in mushroom body neurons in insects [8, 15] reduces the overlap between stimulus representations and thus may support the encoding of specific addressable memory and discrimination [15]. However, it might have a disadvantage for generalization [47], e.g., very sparse representation may separate overlapping stimuli into non-overlapped configural codes. Optimal discrimination will be feasible in this case, but generalization is not possible, and vice versa, generating a generalization-discrimination trade-off [47]. We hypothesized that PN-KC plasticity and PN-to-KC connection number can affect this trade-off. We thus compare the generalization and discrimination capacity of class I model (PN-to-KC = 45–55) and class II model (PN-to-KC = 5–15) with and without plastic PN-KC synapses. For generalization, various pairs of two patterns (both as CS+) of decreased similarity are used to train the class I and class II models. The generalization score is the difference between trained model responses to a novel pattern (between the two training patterns) and to the two training patterns (Figure 4A). Similarly, for discrimination (Figure 4A), models are trained with various pairs of two patterns of decreasing similarity—one as CS+ and another CS−. The discrimination score is the response difference of trained model to CS+ and CS−.

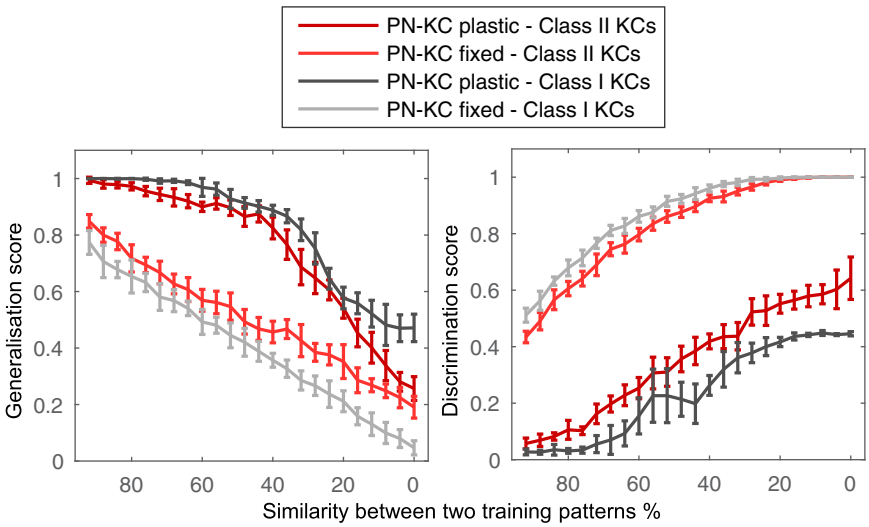

Figure 4. Kenyon Cell Classes and PN-KC Plasticity in Generalization-Discrimination

A generalization score (mean of ten network realizations ± SD) is given as summed differences—after being trained on two patterns (both as CS+) with various levels of similarity—between model learned preferences to a novel pattern that lies in the middle (relative to their positions in the input continuum) of the two training patterns and the two patterns, normalized to [0, 1]. A high score here indicates a good generalization capacity and vice versa. A discrimination score (mean of ten network realizations ± SD) is measured as the difference between learned preference of the model to CS+ and CS−, with various pairs of two input patterns that have different levels of similarity, also normalized into the range [0, 1]. A high score indicates a good discrimination capacity and vice versa. Class I Kenyon cell model used here refers to a model with PN-to-KC connection number of 45–55, while class II model has a PN-to-KC connection number of 5–15. Both class I and class II models with fixed PN-KC synapses show a lower generalization score and a higher discrimination score than the models with plastic PN-KC synapses, suggesting that KC sparse representation with random PN-KC connection is naturally efficient at fine discrimination while limited at generalization. Plasticity in PN-KC (left, generalization: PN-KC plastic versus PN-KC fixed) can alter this and turn the model to more readily generalize. Comparing the class I model with the class II model that implements learning, the former is better at generalization (left: PN-KC plastic class I versus PN-KC plastic class II) but worse at discrimination (right: PN-KC plastic class I versus PN-KC plastic class II) than the latter. See also Figure S4.

Both class I and class II models with fixed PN-KC synapses show lower generalization scores but higher discrimination scores than models with plastic PN-KC synapses. This suggests that with fixed PN-KC synapses, both models show naturally optimal discrimination but limited generalization capacity. Thus, the learning capacity in PN-KC synapses, or the microcircuits in mushroom body calyces [48], might be able to control the generalization-discrimination trade-off. The class I model shows a better generalization score but poorer discrimination than the class II model, suggesting that PN-to-KC connection number might also contribute to this trade-off, with class I Kenyon cells being better at generalization and class II Kenyon cells being better at discrimination. Taken together, PN-KC plasticity and different classes of Kenyon cells may serve as a compensation in the natural inadequacy of Kenyon cell sparse coding in generalization. Mixtures of different classes of Kenyon cells thus may grant flexibility in adapting to generalization and discrimination contexts; this needs further experimental study.

SUPPLEMENTAL INFORMATION

Supplemental Information includes Supplemental Experimental Procedures and four figures and can be found with this article online at http://dx.doi.org/10.1016/j.cub.2016.10.054.

AUTHOR CONTRIBUTIONS

Conceptualization, F.P. and L.C.; Methodology, Software, and Investigation, F.P.; Writing – Original Draft, F.P. and L.C.; Writing – Review & Editing, F.P. and L.C.

ACKNOWLEDGMENTS

We thank Mark Roper, HaDi MaBouDi, Vera Vasas, Barbara Webb, Wolfgang Rössler, and Thomas Nowotny for comments. Work was funded by HFSP program grant RGP0022/2014 and PhD grant 201208440235 from the China Scholarship Council.

Received: May 31, 2016
Revised: September 6, 2016
Accepted: October 26, 2016
Published: December 22, 2016

REFERENCES

1. Wright, G.A., Choudhary, A.F., and Bentley, M.A. (2009). Reward quality influences the development of learned olfactory biases in honeybees. Proc. Biol. Sci. *276*, 2597–2604.
2. Leonard, A.S., Dornhaus, A., and Papaj, D.R. (2011). Flowers help bees cope with uncertainty: signal detection and the function of floral complexity. J. Exp. Biol. *214*, 113–121.
3. Andrew, S.C., Perry, C.J., Barron, A.B., Berthon, K., Peralta, V., and Cheng, K. (2014). Peak shift in honey bee olfactory learning. Anim. Cogn. *17*, 1177–1186.
4. Martínez-Harms, J., Márquez, N., Menzel, R., and Vorobyev, M. (2014). Visual generalization in honeybees: evidence of peak shift in color discrimination. J. Comp. Physiol. A Neuroethol. Sens. Neural Behav. Physiol. *200*, 317–325.
5. Devaud, J.M.M., Papouin, T., Carcaud, J., Sandoz, J.C.C., Grünewald, B., and Giurfa, M. (2015). Neural substrate for higher-order learning in an insect: mushroom bodies are necessary for configural discriminations. Proc. Natl. Acad. Sci. USA *112*, E5854–E5862.
6. Menzel, R. (2012). The honeybee as a model for understanding the basis of cognition. Nat. Rev. Neurosci. *13*, 758–768.
7. Caron, S.J.C., Ruta, V., Abbott, L.F., and Axel, R. (2013). Random convergence of olfactory inputs in the Drosophila mushroom body. Nature *497*, 113–117.
8. Szyszka, P., Ditzen, M., Galkin, A., Galizia, C.G., and Menzel, R. (2005). Sparsening and temporal sharpening of olfactory representations in the honeybee mushroom bodies. J. Neurophysiol. *94*, 3303–3313.

9. Szyszka, P., Galkin, A., and Menzel, R. (2008). Associative and non-associative plasticity in kenyon cells of the honeybee mushroom body. Front. Syst. Neurosci. *2*, 3.
10. Strube-Bloss, M.F., Nawrot, M.P., and Menzel, R. (2011). Mushroom body output neurons encode odor-reward associations. J. Neurosci. *31*, 3129–3140.
11. Rybak, J. (2012). The digital honey bee brain atlas. In Honeybee Neurobiology and Behavior: A Tribute to Randolf Menzel, C.G. Galizia, D. Eisenhardt, and M. Giurfa, eds. (Springer), pp. 125–140.
12. Honegger, K.S., Campbell, R.A.A., and Turner, G.C. (2011). Cellular-resolution population imaging reveals robust sparse coding in the Drosophila mushroom body. J. Neurosci. *31*, 11772–11785.
13. Grünewald, B. (1999). Morphology of feedback neurons in the mushroom body of the honeybee, Apis mellifera. J. Comp. Neurol. *404*, 114–126.
14. Papadopoulou, M., Cassenaer, S., Nowotny, T., and Laurent, G. (2011). Normalization for sparse encoding of odors by a wide-field interneuron. Science *332*, 721–725.
15. Lin, A.C., Bygrave, A.M., de Calignon, A., Lee, T., and Miesenböck, G. (2014). Sparse, decorrelated odor coding in the mushroom body enhances learned odor discrimination. Nat. Neurosci. *17*, 559–568.
16. Haehnel, M., and Menzel, R. (2010). Sensory representation and learning-related plasticity in mushroom body extrinsic feedback neurons of the protocerebral tract. Front. Syst. Neurosci. *4*, 161.
17. Filla, I., and Menzel, R. (2015). Mushroom body extrinsic neurons in the honeybee (Apis mellifera) brain integrate context and cue values upon attentional stimulus selection. J. Neurophysiol. *114*, 2005–2014.
18. Hammer, M. (1993). An identified neuron mediates the unconditioned stimulus in associative olfactory learning in honeybees. Nature *366*, 59–63.
19. Aso, Y., Sitaraman, D., Ichinose, T., Kaun, K.R., Vogt, K., Belliart-Guérin, G., Plaçais, P.-Y., Robie, A.A., Yamagata, N., Schnaitmann, C., et al. (2014). Mushroom body output neurons encode valence and guide memory-based action selection in Drosophila. eLife *3*, e04580.
20. Okada, R., Rybak, J., Manz, G., and Menzel, R. (2007). Learning-related plasticity in PE1 and other mushroom body-extrinsic neurons in the honeybee brain. J. Neurosci. *27*, 11736–11747.
21. Cassenaer, S., and Laurent, G. (2012). Conditional modulation of spike-timing-dependent plasticity for olfactory learning. Nature *482*, 47–52.
22. Huerta, R., and Nowotny, T. (2009). Fast and robust learning by reinforcement signals: explorations in the insect brain. Neural Comput. *21*, 2123–2151.
23. Wessnitzer, J., Young, J.M., Armstrong, J.D., and Webb, B. (2012). A model of non-elemental olfactory learning in Drosophila. J. Comput. Neurosci. *32*, 197–212.
24. Bazhenov, M., Huerta, R., and Smith, B.H. (2013). A computational framework for understanding decision making through integration of basic learning rules. J. Neurosci. *33*, 5686–5697.
25. Ardin, P., Peng, F., Mangan, M., Lagogiannis, K., and Webb, B. (2016). Using an insect mushroom body circuit to encode route memory in complex natural environments. PLoS Comput. Biol. *12*, e1004683.
26. Owald, D., Felsenberg, J., Talbot, C.B., Das, G., Perisse, E., Huetteroth, W., and Waddell, S. (2015). Activity of defined mushroom body output neurons underlies learned olfactory behavior in Drosophila. Neuron *86*, 417–427.
27. Hige, T., Aso, Y., Modi, M.N., Rubin, G.M., and Turner, G.C. (2015). Heterosynaptic plasticity underlies aversive olfactory learning in Drosophila. Neuron *88*, 985–998.
28. Joerges, J., Küttner, A., Galizia, C.G., and Menzel, R. (1997). Representations of odours and odour mixtures visualized in the honeybee brain. Nature *387*, 285–288.
29. Yamagata, N., Schmuker, M., Szyszka, P., Mizunami, M., and Menzel, R. (2009). Differential odor processing in two olfactory pathways in the honeybee. Front. Syst. Neurosci. *3*, 16.
30. Krofczik, S., Menzel, R., and Nawrot, M.P. (2009). Rapid odor processing in the honeybee antennal lobe network. Front. Comput. Neurosci. *2*, 9.
31. Brill, M.F., Rosenbaum, T., Reus, I., Kleineidam, C.J., Nawrot, M.P., and Rössler, W. (2013). Parallel processing via a dual olfactory pathway in the honeybee. J. Neurosci. *33*, 2443–2456.
32. McLaren, I.P.L., and Mackintosh, N.J. (2002). Associative learning and elemental representation: II. Generalization and discrimination. Anim. Learn. Behav. *30*, 177–200.
33. Guerrieri, F., Schubert, M., Sandoz, J.C., and Giurfa, M. (2005). Perceptual and neural olfactory similarity in honeybees. PLoS Biol. *3*, e60.
34. Ghirlanda, S., and Enquist, M. (2003). A century of generalization. Anim. Behav. *66*, 15–36.
35. Lynn, S.K. (2010). Decision-making and learning: the peak shift behavioral response. In Encyclopedia of Animal Behavior, M.D. Breed, and J. Moore, eds. (Elsevier), pp. 470–475.
36. Ghirlanda, S., and Enquist, M. (1998). Artificial neural networks as models of stimulus control. Anim. Behav. *56*, 1383–1389.
37. Rumelhart, D.E., Hinton, G.E., and Williams, R.J. (1986). Learning representations by back-propagating errors. Nature *323*, 533–536.
38. Deisig, N., Giurfa, M., Lachnit, H., and Sandoz, J.C. (2006). Neural representation of olfactory mixtures in the honeybee antennal lobe. Eur. J. Neurosci. *24*, 1161–1174.
39. Deisig, N., Giurfa, M., and Sandoz, J.C. (2010). Antennal lobe processing increases separability of odor mixture representations in the honeybee. J. Neurophysiol. *103*, 2185–2194.
40. Rescorla, R.A., and Wagner, A.R. (1972). A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. In Classical Conditioning II: Current Research and Theory, A.H. Black, and W.F. Prokasy, eds. (Appleton-Century-Crofts), pp. 64–99.
41. Luo, S.X., Axel, R., and Abbott, L.F. (2010). Generating sparse and selective third-order responses in the olfactory system of the fly. Proc. Natl. Acad. Sci. USA *107*, 10713–10718.
42. Schmuker, M., Pfeil, T., and Nawrot, M.P. (2014). A neuromorphic network for generic multivariate data classification. Proc. Natl. Acad. Sci. USA *111*, 2081–2086.
43. Huerta, R., Nowotny, T., García-Sanchez, M., Abarbanel, H.D.I., and Rabinovich, M.I. (2004). Learning classification in the olfactory system of insects. Neural Comput. *16*, 1601–1640.
44. Strausfeld, N.J. (2002). Organization of the honey bee mushroom body: representation of the calyx within the vertical and gamma lobes. J. Comp. Neurol. *450*, 4–33.
45. Fahrbach, S.E. (2006). Structure of the mushroom bodies of the insect brain. Annu. Rev. Entomol. *51*, 209–232.
46. Jortner, R.A., Farivar, S.S., and Laurent, G. (2007). A simple connectivity scheme for sparse coding in an olfactory system. J. Neurosci. *27*, 1659–1669.
47. Spanne, A., and Jörntell, H. (2015). Questioning the role of sparse coding in the brain. Trends Neurosci. *38*, 417–427.
48. Hourcade, B., Muenz, T.S., Sandoz, J.C., Rössler, W., and Devaud, J.M. (2010). Long-term memory leads to synaptic reorganization in the mushroom bodies: a memory trace in the insect brain? J. Neurosci. *30*, 6461–6465.

Current Biology, Volume *27*

Supplemental Information

A Simple Computational Model of the Bee Mushroom Body Can Explain Seemingly Complex Forms of Olfactory Learning and Memory

Fei Peng and Lars Chittka

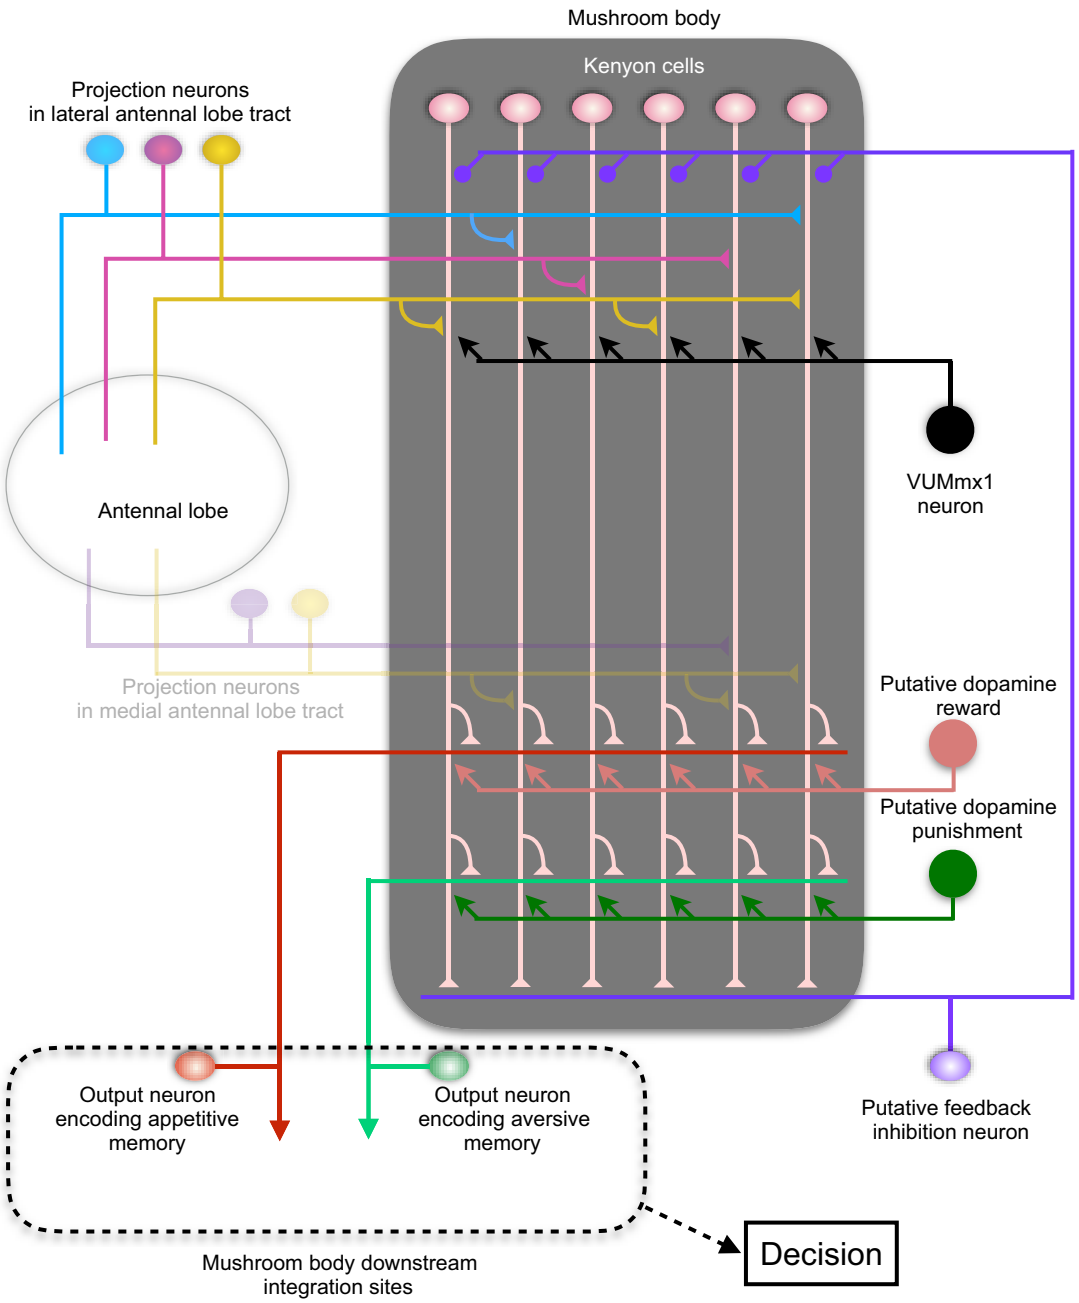

Figure S1 related to Figure 1. Diagram of the honeybee olfactory circuit. Odorant molecules are first sensed by ~60,000 olfactory receptor neurons (ORNs), which project to the antennal lobe (AL), the primary olfactory processing centre. Within the antennal lobe, the olfactory information is converged onto ~160 spherical structures called glomeruli (GL). The upper hemilobe of the antennal lobe is innervated by ~510 lateral antennal lobe tract (l-ALT) projection neurons (PNs), and the lower hemilobe is innervated by ~410 medial antennal lobe tract (m-ALT) projection neurons [S1]. Carrying largely redundant information [S2], both tracts project to the mushroom body (MBs) and the lateral horn (LH), but in different order. Projection neurons in the lateral antennal lobe tract are considered to encode odour identity, while the medial antennal lobe tract is thought to encode the odour concentration [S3]. Here we focus on the lateral antennal lobe tract projection neurons. In bees, there are about 20,000 clawed, or class II, Kenyon cells [S4], so the ratio between lateral antennal lobe tract projection neurons and

clawed Kenyon cells neuronal number is ~1:40. Kenyon cells synapse onto ~400 mushroom body extrinsic neurons (ENs). A subset of extrinsic neurons project back to Kenyon cell calyx (Kenyon cell input region), forming a putative inhibitory feedback pathway [S5]. Evidence from *Drosophila* [S6], locust [S7] and honeybee [S8] suggests that the KC-EN synapse is plastic, and appetitive learning might lead to synaptic depression among KC-EN synapses. In *Drosophila*, aversive learning also induces synaptic depression [S9], and appetitive and aversive valence are suggested to be encoded by distinct group of extrinsic neurons [S10]. One possibility is that the appetitive and aversive memories are stored in distinct KC-EN synapses. In a simplified version, we only show two extrinsic neurons here, with one for appetitive and another for aversive valence. Putative synapses for appetitive memory are shown in red, while the ones for aversive memory are shown in green. In addition, extrinsic neurons for appetitive and aversive valence have been shown to be simply additive in behavioural decision making [S10].

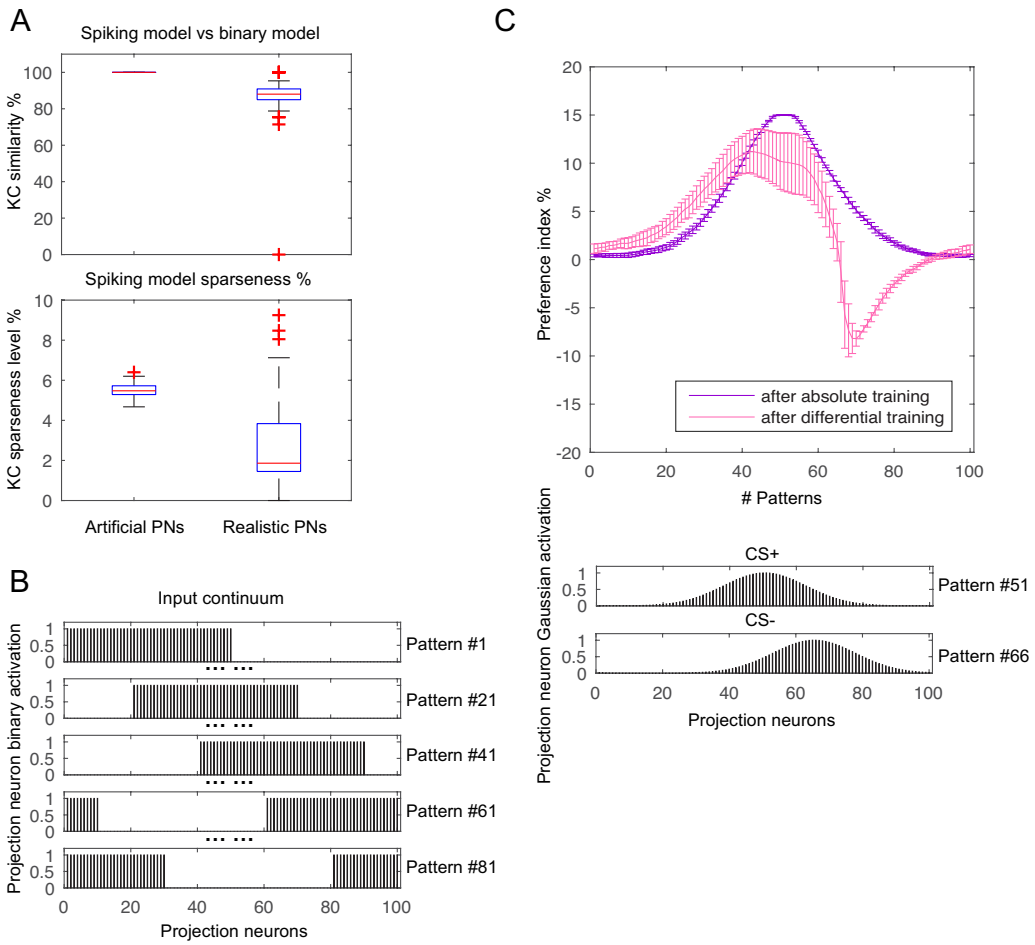

Figure S2 related to Figure 1 and Figure 2. Binary and spiking model comparison, artificial input continuum and peak shift. (A) For the upper panel, both artificially generated binary (100 different patterns in the input continuum, B), and more realistically modelled projection neuron firing patterns (110 different odour patterns in [S11]) are adopted (see also Supplemental Experimental Procedures). We apply each pattern to both the binary model and a standard spiking network model using Izhikevich neurons [S12, S13] and a global feedback inhibition neuron. Same PN-KC connectivity matrix is applied to both models, allowing a direct comparison of the input driven Kenyon cell activation. We show that in terms of activated Kenyon cells, binary model and spiking model generate highly similar results (close to 100%), indicating that the binary model is sufficient in substituting a spiking model for sparse Kenyon cell representation. In other words, the cut-off method (Supplemental Experimental Procedures) for Kenyon cell sparseness has a similar effect to a global feedback inhibition, without requiring any additional assumptions. For the lower panel, the spiking network model Kenyon cell sparseness toward artificial stimuli is shown to vary moderately (4-6%), while vary hugely (0-10%) toward realistic stimuli. Note that the comparison of each case is done by matching the actual sparseness level of Kenyon cells from the spiking model, with the binary model, as the latter one allows a precise control of the sparseness. It thus also demonstrates that for theoretical exploration, the binary model might be better since it would be difficult for the spiking model to maintain a certain level of sparseness among Kenyon cells, irrespective of the input characteristics. **(**B) A set of projection neuron input patterns that have systematically different levels of similarity is generated. As described in Figure 1B, each stimulus is modelled as a combination of activated projection neurons (50%, firing intensity =1) and inactivated projection neurons (50%, firing intensity =0). Here we show five different inputs that slightly differ from each other. The total number of artificial patterns are 100, with the smallest similarity among two patterns being 2% (see Supplemental Experimental Procedures for similarity measurement). (C) Using projection neuron patterns of various amplitudes other than the binary activation patterns, we show that the model also manifests peak shift phenomenon (mean of 100 network realisations ± standard deviation), consistent with Figure 2A. The projection neuron responses shown here (CS+ and CS-) are generated from a Gaussian probability density function.

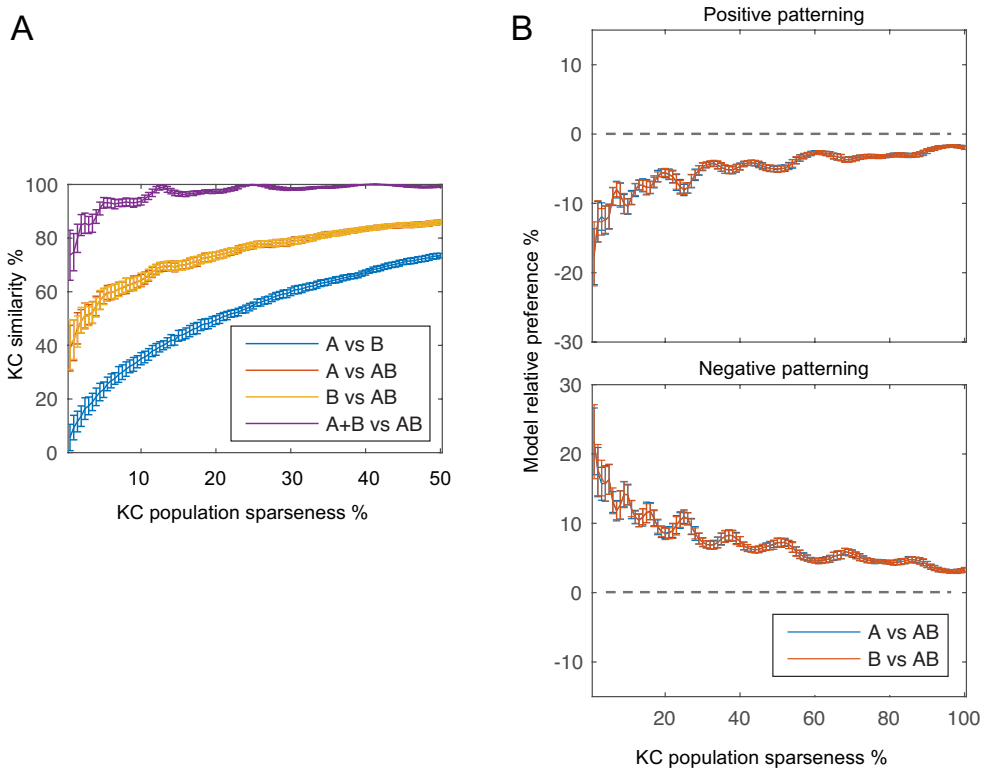

Figure S3 related to Figure 2. The effect of Kenyon Cell population sparseness on patterning discriminations. (A) By exhaustively varying the sparseness level from 0.5% to 50% and comparing the joint Kenyon cell representation of odours A & B, denoted as A+B, with the Kenyon cell representation of AB mixture and others, less sparse (or denser) Kenyon cell representation leads to a higher similarity between A+B and AB mixture, and between single odours and AB mixture, indicating a deterioration of configural code generation. (B) Along with the decrease of Kenyon cell sparseness level (or increase of Kenyon cell activation), the model's relative preference is measured by the difference of trained model response to A and AB, to B and AB respectively. This shows a similar trend of that shown in (A), i.e., decrease of Kenyon cell sparseness leads to naturally smaller difference between single odours and the binary mixture. Overall this suggests the necessity of sufficiently sparse Kenyon cell representation in mediating the configural coding capacity (mean of 100 network realisations ± standard deviation).

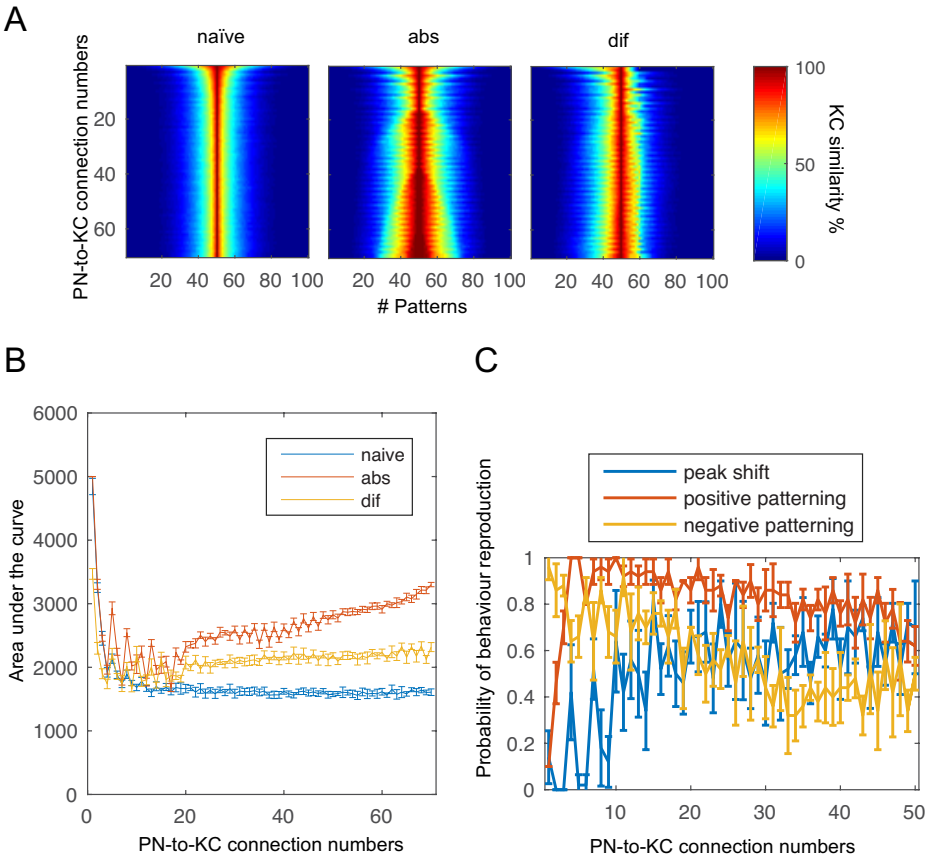

Figure S4 related to Figure 3 and Figure 4. Performance of models with different PN-to-KC connection numbers. (A) The three graphs are generated by taking the Kenyon cell activation pattern for the central pattern as the reference, and compare the model Kenyon cell activation similarity for each pattern in the input continuum with it. Different models are generated by allocating Kenyon cells with different projection neuron numbers. Before training (naïve case), the models have a rather narrow band of generalisation, which actually decreases further slightly along with the increase of PN-to-KC connection numbers. After absolute training on the central pattern, the generalisation curve from models with larger PN-to-KC connection numbers is much broader than the models that only have a few projection neuron connections per Kenyon cell. The trend is similar in differential training group, but in that case, generalisation is not as broad as after absolute training. This is even more visible in (B), where the area under each generalisation curve has been plotted and all the graphs in (A) have been included (mean of 100 network realisations for each PN-to-KC connection number ± standard deviation). This suggests that models with a large number of PN-to-KC connections (e.g., class I model with PN-to-KC = 45-55) may facilitate broad generalisation while a few PN-KC connections (e.g., class II model with PN-to-KC = 5-15) may actually be reliable in discrimination. (C) Model performance (probability of behavioural reproduction over 10 groups of 10 network realisations for each model, and shown as mean of 5 repetitions ± standard deviation) on peak shift and patterning discrimination is shown as a probability of behaviour reproduction, similarly as shown in Figure 3. With increased PN-to-KC connection number, the probability for peak shift is tend to increase, but is tend to decrease for both positive and negative patterning. This is in line with the result shown in Figure 3 and Figure 4, suggesting that PN-to-KC connection numbers, when coupled with plastic synapses, may contribute to the generalisation-discrimination trade-off.

Supplemental Experimental Procedures

The mushroom body circuit model

We constructed a biologically constrained neuronal network, using binary neurons with global feedback inhibition, and implemented simple reward-based synaptic modifications. We implemented learning in both mushroom body input region – calyces (PN-KC synapses), and output region – lobes (KC-EN synapses).

Projection neuron to Kenyon cell connectivity

The random connectivity between projection neurons and Kenyon cells is generated in two steps. First we draw for each Kenyon cell j a number of connections, C_{KC}^{j}, from a uniform distribution over a user defined interval (5-15 for class II and 45-55 for class I). Then, we randomly choose C_{KC}^{j} unrepeated projection neurons that Kenyon cell j connects to. The chosen connectivity is stored in the connectivity matrix C_{PN-KC} using 1 for existing connections and 0 otherwise. The synaptic weights among PN-KC synapses are stored in the matrix W_{PN-KC} and initialised by assigning a non-zero identical weight g_0 to all existing PN-KC connections (Supplemental Table).

Parameter	PN-KC	KC-EN
α^+	0.006	0.006
α^-	0.007	0.008
g_0	0.2	0.2
g_{max}	0.4	0.4
g_{min}	0	0
class II model	PN-to-KC = 5-15	all KCs to two ENs
class I model	PN-to-KC = 45-55	all KCs to two ENs

Supplemental Table, related to Supplemental Experimental Procedures. Model parameter used in producing elemental, peak shift and patterning tasks, and different classes of Kenyon cells. The first five parameters are related to the reward-based synaptic weight modification rule, where α^+ and α^- are positive and negative learning rates respectively. The initial weights are set as g_0 = 0.2 for both PN-KC and KC-EN synapses, and are bounded by the range [0, 0.4] by g_{min} and g_{max}. The class II Kenyon cell model has a PN-to-KC connection number uniform randomly drawn from 5-15 (out of 100 projection neurons), and class I Kenyon cell model has a PN-to-KC connection number uniform randomly drawn from 45-55 (out of 100 projection neurons). Note that a broad range of parameters other than the above can also reproduce the behavioural tasks described here.

Kenyon cell sparse coding by feedback inhibition

For each stimulus activated projection neuron (N_{PN}=100) firing pattern, Kenyon cell (N_{KC}=4000) representation is given by a function of the product of weight matrix PN-KC and projection neuron inputs:

$$R_{KC} = h(W_{PN-KC} * R_{PN}) \quad (1)$$

where h is a function that selects Kenyon cells that receive most intense inputs with a given Kenyon cell population sparseness level (e.g. 5%), and outputs a binary vector that only selected Kenyon cells are labelled as 1s (fired), and others are labelled as 0s (not fired). We assume that GABAergic inhibitory inputs onto the Kenyon cells' dendritic regions (calyces) serve as a feedback inhibition mechanism to regulate Kenyon cell population sparseness. This is implemented with Matlab's built-in '*sort*' function, i.e. a simplified version of feedback inhibition. We also use a spiking model with Izhikevich spiking neurons [S12, S13] and a one-to-all global feedback neuron to simulate the feedback inhibition on Kenyon cells, and we show that the '*sort*' function implemented on the present binary

network generates almost identical Kenyon cell population codes as with the spiking network (Figure S2A), indicating that this simplification is efficient in substituting the essential function of global feedback inhibition. The comparison of Kenyon cell activation is done by using the similarity calculation defined below.

A simple learning rule among PN-KC synapses

Despite the evidence of associative plasticity among PN-KC synapses in bees [S14], the detailed plasticity rule is yet to be discovered. We here recruit a simple supervised learning rule, consistent with [S15, S16], such that:

$$W_{PN-KC}^{ji}(t+1) = \begin{cases} W_{PN-KC}^{ji}(t) + \alpha^{+}, & if\ r = 1,\ R_{PN}^{i} > 0,\ R_{KC}^{j} = 1,\ C_{PN-KC}^{ji} = 1 \\ W_{PN-KC}^{ji}(t) - \alpha^{-}, & if\ r = -1,\ R_{PN}^{i} > 0,\ R_{KC}^{j} = 1,\ C_{PN-KC}^{ji} = 1 \\ W_{PN-KC}^{ji}(t), & otherwise \end{cases} \quad (2)$$

where $W_{PN-KC}^{ji}(t)$ denotes the synaptic weight between activated projection neuron i and Kenyon cell j that have existing connection ($C_{PN-KC}^{ji} = 1$), at trial number t. In terms of neuronal activation, projection neuron i is activated, if the firing strength R_{PN}^{i} is greater than 0, both in the case of binary activations (Figure S2B) and in the case of activations with various amplitudes (Figure S2C). Kenyon cell j is 'activated' if the corresponding firing strength R_{KC}^{j} is equal to 1 (binary neurons). r is the reward signal that tells the model whether a stimulus is paired with reward ($r = 1$) or punishment ($r = -1$), and α^{+} and α^{-} are the learning rates for CS+ and CS- respectively. Overall, the synaptic weights W_{PN-KC}^{ji} might increase or decrease linearly within a range of $[g_{min}, g_{max}]$. Model parameters are listed in the Supplemental Table. Note that the very same set of parameters were used to produce elemental learning, peak shift and patterning learning. We argue that a robust model should support a broad range of behaviours that do not require tuning of parameters for any specific task.

Appetitive and aversive extrinsic neurons and KC-EN plasticity

Evidence in *Drosophila* suggests a separation of appetitive and aversive valence encoding among distinct extrinsic neurons via different KC-EN synapses. To simplify the problem, but capture the essential structural segregation, two extrinsic neurons are recruited, with EN_{+} for appetitive valence encoding and EN_{-} for aversive valence encoding. Both of them receive all-to-one connections from all the Kenyon cells, forming parallel synapses. The Kenyon cell representation R_{KC} is then transformed into the extrinsic neuron responses separately, in a linear fashion:

$$\begin{aligned} R_{EN_{+}} &= W_{KC-EN_{+}} * R_{KC} \\ R_{EN_{-}} &= W_{KC-EN_{-}} * R_{KC} \end{aligned} \quad (3)$$

where $W_{KC-EN_{+}}$ and $W_{KC-EN_{-}}$ are the weight matrices from all the Kenyon cells to EN_{+} and EN_{-} respectively. Both of the two matrices are initialised as all-to-one connectivity and with an equal weight g_0, denoted as W_{init}. For KC-EN synapses, it has been shown in fly [S6], locust [S7] and honeybee [S8] that appetitive learning in general leads to synaptic depression, despite some exceptions [S17]. Moreover, aversive learning in *Drosophila* [S9] also results in synaptic depression. We thus modified the learning rule for PN-KC synapses (Equation 2), such that:

$$\begin{aligned} W_{KC-EN_{+}}^{j}(t+1) &= \begin{cases} W_{KC-EN_{+}}^{j}(t) - \alpha^{+}, & if\ r = 1,\ R_{KC}^{j} = 1 \\ W_{KC-EN_{+}}^{j}(t), & otherwise \end{cases} \\ W_{KC-EN_{-}}^{j}(t+1) &= \begin{cases} W_{KC-EN_{-}}^{j}(t) - \alpha^{-}, & if\ r = 1,\ R_{KC}^{j} = 1 \\ W_{KC-EN_{-}}^{j}(t), & otherwise \end{cases} \end{aligned} \quad (4)$$

where $W_{KC-EN_{+}}^{j}(t)$ and $W_{KC-EN_{-}}^{j}(t)$ denote the synaptic weight between activated Kenyon cell j ($R_{KC}^{j} = 1$) and EN_{+} and EN_{-} respectively, at trial number t. Both $W_{KC-EN_{+}}^{j}$ and $W_{KC-EN_{-}}^{j}$ are bounded by the range $[g_{min}, g_{max}]$. Appetitive learning thus will lead to a depression among only KC-EN$_{+}$ synapses but not among KC-EN$_{-}$ synapses, and *vice versa*.

Model performance metric

Recent evidence from *Drosophila* suggest that the behavioural decision of the insect might be underpinned by a simple summation of distinct valence coding output neurons. In other words, the extrinsic neuron group is balanced out among naïve flies, but biased toward learned odours among trained flies. Following this assumption of simple integration, the model response here is measured by the preference index (PI), which is shown as a percentage

increase or decrease of the model behavioural preference toward a stimulus (e.g., Figure 1C, Figure 2A, Figure 2C), calculated by:

$$PI = -\frac{R_{EN_+} - R_{EN_-}}{W_{init} * R_{KC}} * 100\% \quad (5)$$

where R_{EN_+} and R_{EN_-} are given by Equation 3, and W_{init} is the initial KC-EN weight matrix with an equal weighting. Initially the preference toward any stimulus is balanced to be 0% according to this definition, and any parallel changes among the two extrinsic neurons would bias the final model response toward appetitive or aversive stimuli. Note that the minus sign in the front here is to model the putative inhibition over sensory-motor connections [S8]. For example, a learned depression among EN_+ for appetitive stimuli will eventually transform into a behavioural preference or positive response, e.g., proboscis extension in bees.

Alternatively, recent evidence also indicates that some extrinsic neurons can encode positive and negative valence bi-directionally, which gave decreased responses toward appetitive stimuli but increased responses toward aversive stimuli [S6]. It can be easily inferred that this simple summation rule described above is equivalent to this bi-directional valence encoding among KC-EN synapses. Indeed, we also test a model with a single extrinsic neuron that have bi-directional KC-EN synapses, which produces identical result as the simple summation. However, further study is needed for a better understanding of the mushroom body output neuron circuit.

Similarity Calculation

To analyse the Kenyon cell activation similarity changes with different levels of sparseness, we introduce a similarity metric that is similar to Hamming distance, but only taking the activated Kenyon cells into account. Given two column vectors x and y, the Similarity S is defined as follows:

$$S = \frac{\left|\left\{j \mid R^j_{KC_x} = R^j_{KC_y} = 1\right\}\right|}{\left|\left\{j \mid R^j_{KC_x} = 1\right\}\right|} * 100\% \quad (6)$$

where $R^j_{KC_x}$ is the jth Kenyon cell response to one pattern, and $R^j_{KC_x}$ the jth Kenyon cell response to another. Note that the cardinality $|\{j \mid R^j_{KC_x} = 1\}|$ is exchangeable with the cardinality $\left|\left\{j \mid R^j_{KC_y} = 1\right\}\right|$, as the number of Kenyon cells activated for the two patterns is the same.

Absolute and differential training procedure

In experimental studies, absolute conditioning involves repeatedly pairing a conditioned stimulus (CS+) with reward, without exposing the animal to unrewarded stimuli [S18, S19]. In contrast, differential training refers to a sequence of training with alternating and equal amount of exposure to CS+ and CS-, which forces the animal to form both appetitive, and aversive (or no reward) memories. In bees [S18, S19] and flies [S20], it has been shown that the two training procedures produce different behavioural outcomes: after absolute training the animal tends to recognise similar (relative to CS+) but novel stimuli as rewarding, but not after differential training, which is shown to be necessary in fine discriminations.

Peak shift training and testing procedure

Honeybees display the 'peak shift' phenomenon in olfactory learning [S21, S22]: after being trained with a differential training procedure on two similar odours – one paired with reward (CS+) while the other is paired with punishment (CS-), the bee not only responds to CS+ more strongly than to CS-, but responds most strongly to a novel stimulus that is further away from CS- than CS+. In other words, a shifted response peak or bias is formed. We here follow the established training and testing procedure for peak shift in honeybee olfactory learning. Specifically, each network realisation (N=100) in the experiment group is trained with 10 trials of CS+ and 10 trials of CS- with a pseudorandom sequence (differential training), and then is tested with the entire input continuum (Figure S2B). Each network realisation (N=100) in the control group is trained with 5 trials of CS+, and then is tested again with the input continuum. Note that we apply fewer trials here for the absolute training than the differential training. The reason is that with the simplified linear learning rule, equal number of training trials with absolute training will drive the model response change to be much bigger than with differential training, which makes it visually difficult to compare the two groups.

Positive and negative patterning training and testing procedure

Bees also can discriminate two individual odours (A and B) from their binary mixture (AB) whether the individual odours are being rewarded (negative patterning discrimination) or their mixture (positive patterning discrimination) [S23, S24]. Such discriminations have been classified as configural or 'non-elemental' learning, requiring the compound to be treated differently from its elements [S25, S26], which is often viewed as a more complex form of learning that involves ambiguity and nonlinearity. We here follow the established procedure for patterning discrimination [S24], such that we train each network realisation with 5 blocks of trials in total, each containing a pseudorandom sequence of 4 trials. Specifically, in positive patterning, the 4 trials are A-, B-, and twice of AB+, and in negative patterning the 4 trials are A+, B+ and twice of AB-. After each block of training, the model responses to A, B, and AB mixture are tested.

Three different sets of input patterns used in patterning discrimination

To test if the model can solve the patterning task under various input conditions, three sets of inputs are used: set1 – artificial projection neuron patterns for input A and B that have moderate level (40%) of overlap; set2 – artificial projection neuron patterns for input A and B that have no overlap; set3 – realistic Projection neuron firing patterns used in [S11], which is generated from *Drosophila* olfactory receptor neuron responses to 110 different odours, and then modelled by [S11] to simulate the realistic olfactory projection neuron firing patterns. Since there are only 20 projection neuron responses for each odour [S11], we here replicate each projection neuron response for a given odour into 5 identical responses, such that each odour is represented by 100 projection neurons in total. As in real honeybees, each glomerulus is innervated by ~5 projection neurons that are likely to convey more or less the same information. Finally, we normalise those projection neuron firing rates to a range of [0, 1], with the low rates (<0.2) truncated to be 0. During patterning discrimination, two odours are randomly drawn from the 110 odours for each network realisation.

Testing peak shift performance in learning rate space

To evaluate the model's performance with different sets of learning rates – the amount of synaptic weight modification when the reward signal is indicated. A positive learning rate here refers to the amount of weight change in each trial for CS+, and negative learning rate to the amount of weight change in each trial for CS-. Each set of learning rates is generated (or each square in the heat maps) by fixing positive learning rate – that used to associate CS with reward – in both PN-KC synapses and KC-EN synapses, with a value of an arbitrary choice (0.006), and by varying the negative learning rate for both synapses from 0.001 to 0.012 with an incremental of 0.001, which displays as a 2D (12 x 12) learning rate space (heat maps in Figure 3).We define the probability of successful reproduction in peak shift as how many groups (N = 10) of virtual bees can reach statistical significant level ($P < 0.05$) in paired-sample t test on the differences between peak responses (if any) and the responses to CS+. Within each group, 10 network realisations (virtual bees) are tested for training induced peak responses, which refer to the maximum responses to a novel stimulus that is different from CS+, but also is further way from CS- than CS+. We show that a large number of learning rate sets can reproduce peak shift effect. Interestingly, a large negative learning rate on KC-EN synapse, in relative to the positive rate, seems to be crucial, which is in accordance with [S21]. In addition, when comparing with the heat map generated by only allow learning in KC-EN but not PN-KC, there is almost no peak shift generated, indicating that the PN-KC synapses is crucial in accounting for the peak shift effect, possibly by reshaping Kenyon cell representation for stimuli through learning.

Testing patterning discrimination performance in learning rate space

Similarly, we use the same setup for learning rate space to evaluate the performance of positive and negative patterning discrimination. The probability of behavioural reproduction here is defined as how many groups (each group N = 10) of virtual bees that can reach statistical significant level ($P < 0.05$) in paired-sample *t* test after training. In addition, to test the model on input patterns with various similarity levels, we use 10 different input pattern pairs that have similarity between 0% and 90% for each group, with an increment of 10%. Therefore, the maximum of 10 occurrences here means that the virtual bees have successfully learned to differentiate patterns and mixtures with the given learning rate sets, in all the input pattern pairs that have different similarity levels. It is clear that the model has good performance in most of the learning rate sets. Strikingly, when comparing again the both synapses case with KC-EN only, we found that the model performs even better with the latter, indicating that naïve Kenyon cell sparse representation is already suitable for discrimination; learning within PN-KC synapses here might even hinder downstream discrimination.

Testing model with a simple generalisation and discrimination task

To compare the model capacity with two different classes of Kenyon cells for generalisation and discrimination, a simple task for generalisation and discrimination is introduced. For generalisation, the class I Kenyon cell model (each Kenyon cell receives from 45-55 projection neurons) and the class II Kenyon cell model (each Kenyon cell receives from 5-15 projection neurons) are trained independently with various pairs of artificial patterns, and with plastic and fixed PN-KC synapses respectively. Specifically, each pair of training patterns has a particular level of similarity, depending on the relative positions of the two patterns in the input continuum. After absolute training on the two patterns (both are CS+), the models' performance for a novel pattern that lies in the middle of the positions of the two is measured, and is compared with the response to the two training patterns respectively. The generalisation score (GS) is given by the sum of the difference between the response to the middle novel pattern and the trained response to the two patterns separately ($GS = 2 * PI_{novel} - PI_{CS_+^1} - PI_{CS_+^2}$). Finally, the generalisation score is normalised to the range [0, 1] to allow an overall comparison of different conditions and models with different classes of Kenyon cells. Similarly, for discrimination score (DS), class I Kenyon cell model and class II Kenyon cell model are trained differentially with various pairs of artificial patterns, and the difference between the trained response to CS+ and to CS- is measured ($DS = PI_{CS_+} - PI_{CS\_}$). Finally, the discrimination score is also normalised to the range [0, 1]. In both cases, a high score indicates a good generalisation / discrimination capacity and *vice versa* (10 network realisations for each training pairs).

Supplemental References

S1. Rybak, J. (2012). The Digital Honey Bee Brain Atlas. In Honeybee Neurobiology and Behavior (Springer), pp. 125–140.

S2. Carcaud, J., Hill, T., Giurfa, M., and Sandoz, J. C. (2012). Differential coding by two olfactory subsystems in the honeybee brain. J. Neurophysiol. *108*, 1106–1121.

S3. Yamagata, N., Schmuker, M., Szyszka, P., Mizunami, M., and Menzel, R. (2009). Differential odor processing in two olfactory pathways in the honeybee. Front. Syst. Neurosci. *3*, 1–13.

S4. Szyszka, P., Ditzen, M., Galkin, A., Galizia, C. G., and Menzel, R. (2005). Sparsening and temporal sharpening of olfactory representations in the honeybee mushroom bodies. J. Neurophysiol. *94*, 3303–3313.

S5. Grünewald, B. (1999). Morphology of feedback neurons in the mushroom body of the honeybee, Apis mellifera. J. Comp. Neurol. *404*, 114–126.

S6. Owald, D., Felsenberg, J., Talbot, C. B., Das, G., Perisse, E., Huetteroth, W., and Waddell, S. (2015). Activity of defined mushroom body output neurons underlies learned olfactory behavior in Drosophila. Neuron *86*, 417–427.

S7. Cassenaer, S., and Laurent, G. (2012). Conditional modulation of spike-timing-dependent plasticity for olfactory learning. Nature *487*, 128–128.

S8. Okada, R., Rybak, J., Manz, G., and Menzel, R. (2007). Learning-related plasticity in PE1 and other mushroom body-extrinsic neurons in the honeybee brain. J. Neurosci. *27*, 11736–11747.

S9. Hige, T., Aso, Y., Modi, M. N., Rubin, G. M., and Turner, G. C. (2015). Heterosynaptic Plasticity Underlies Aversive Olfactory Learning in Drosophila. Neuron *88*, 985–998.

S10. Aso, Y., Sitaraman, D., Ichinose, T., Kaun, K. R., Vogt, K., Belliart-Guérin, G., Plaçais, P.-Y., Robie, A. A., Yamagata, N., Schnaitmann, C., et al. (2014). Mushroom body output neurons encode valence and guide memory-based action selection in Drosophila. Elife *3*, e04580.

S11. Luo, S. X., Axel, R., and Abbott, L. F. (2010). Generating sparse and selective third-order responses in the olfactory system of the fly. Proc. Natl. Acad. Sci. U. S. A. *107*, 10713–10718.

S12. Izhikevich, E. M. (2004). Which model to use for cortical spiking neurons? IEEE Trans. Neural Networks *15*, 1063–1070.

S13. Ardin, P., Peng, F., Mangan, M., Lagogiannis, K., and Webb, B. (2016). Using an Insect Mushroom Body Circuit to Encode Route Memory in Complex Natural Environments. PLoS Comput. Biol. *12*, e1004683.

S14. Szyszka, P., Galkin, A., and Menzel, R. (2008). Associative and non-associative plasticity in kenyon cells of the honeybee mushroom body. Front. Syst. Neurosci. *2*, 3.

S15. Bazhenov, M., Huerta, R., and Smith, B. H. (2013). A computational framework for understanding decision making through integration of basic learning rules. J. Neurosci. *33*, 5686–5697.

S16. Dehaene, S., and Changeux, J. P. (2000). Reward-dependent learning in neuronal networks for planning and decision making. Prog. Brain Res. *126*, 217–229.

S17. Strube-Bloss, M. F., Nawrot, M. P., and Menzel, R. (2011). Mushroom body output neurons encode odor-reward associations. J. Neurosci. *31*, 3129–3140.

S18. Giurfa, M. (2004). Conditioning procedure and color discrimination in the honeybee Apis mellifera. Naturwissenschaften *91*, 228–231.

S19. Dyer, A. G., and Chittka, L. (2004). Fine colour discrimination requires differential conditioning in bumblebees. Naturwissenschaften *91*, 224–227.

S20. Barth, J., Dipt, S., Pech, U., Hermann, M., Riemensperger, T., and Fiala, A. (2014). Differential Associative Training Enhances Olfactory Acuity in Drosophila melanogaster. J. Neurosci. *34*, 1819–37.

S21. Wright, G. A., Choudhary, A. F., and Bentley, M. A. (2009). Reward quality influences the development of learned olfactory biases in honeybees. Proc. R. Soc. B *276*, 2597–2604.

S22. Andrew, S. C., Perry, C. J., Barron, A. B., Berthon, K., Peralta, V., and Cheng, K. (2014). Peak shift in honey bee olfactory learning. Anim. Cogn. *17*, 1177–1186.

S23. Deisig, N., Lachnit, H., Giurfa, M., and Hellstern, F. (2001). Configural Olfactory Learning in Honeybees: Negative and Positive Patterning Discrimination. Learn. Mem. *8*, 70–78.

S24. Devaud, J. M. M., Papouin, T., Carcaud, J., Sandoz, J. C. C., Grünewald, B., and Giurfa, M. (2015). Neural substrate for higher-order learning in an insect: Mushroom bodies are necessary for configural discriminations. Proc. Natl. Acad. Sci. U. S. A. *112*, E5854–E5862.

S25. Pearce, J. M. (1994). Similarity and discrimination: a selective review and a connectionist model. Psychol. Rev. *101*, 587–607.

S26. Giurfa, M. (2003). Cognitive neuroethology: Dissecting non-elemental learning in a honeybee brain. Curr.

Opin. Neurobiol. *13*, 726–735.