

## RESEARCH ARTICLE

# A linear discriminant analysis model of imbalanced associative learning in the mushroom body compartment

David Lipshutz<sup>1\*</sup>, Aneesh Kashalikar<sup>1</sup>, Shiva Farashahi<sup>1</sup>, Dmitri B. Chklovskii<sup>1,2</sup>

**1** Center for Computational Neuroscience, Flatiron Institute, New York, New York, United States of America, **2** Neuroscience Institute, New York University School of Medicine, New York, New York, United States of America

\* [dlipshutz@flatironinstitute.org](mailto:dlipshutz@flatironinstitute.org)

## Abstract

To adapt to their environments, animals learn associations between sensory stimuli and unconditioned stimuli. In invertebrates, olfactory associative learning primarily occurs in the mushroom body, which is segregated into separate compartments. Within each compartment, Kenyon cells (KCs) encoding sparse odor representations project onto mushroom body output neurons (MBONs) whose outputs guide behavior. Associated with each compartment is a dopamine neuron (DAN) that modulates plasticity of the KC-MBON synapses within the compartment. Interestingly, DAN-induced plasticity of the KC-MBON synapse is imbalanced in the sense that it only weakens the synapse and is temporally sparse. We propose a normative mechanistic model of the MBON as a linear discriminant analysis (LDA) classifier that predicts the presence of an unconditioned stimulus (class identity) given a KC odor representation (feature vector). Starting from a principled LDA objective function and under the assumption of temporally sparse DAN activity, we derive an online algorithm which maps onto the mushroom body compartment. Our model accounts for the imbalanced learning at the KC-MBON synapse and makes testable predictions that provide clear contrasts with existing models.

## OPEN ACCESS

**Citation:** Lipshutz D, Kashalikar A, Farashahi S, Chklovskii DB (2023) A linear discriminant analysis model of imbalanced associative learning in the mushroom body compartment. *PLoS Comput Biol* 19(2): e1010864. <https://doi.org/10.1371/journal.pcbi.1010864>

**Editor:** Michele Migliore, National Research Council, ITALY

**Received:** September 26, 2022

**Accepted:** January 10, 2023

**Published:** February 6, 2023

**Copyright:** © 2023 Lipshutz et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

**Data Availability Statement:** Code for reproducing the numerical experiments on the synthetic datasets may be found on GitHub (<https://github.com/flatironinstitute/bio-lda>).

**Funding:** The authors received no specific funding for this work.

**Competing interests:** The authors have declared that no competing interests exist.

## Author summary

To adapt to their environments, animals learn associations between sensory stimuli (e.g., odors) and unconditioned stimuli (e.g., sugar or heat). In flies and other insects, olfactory associative learning primarily occurs in a brain region called the mushroom body, which is partitioned into multiple compartments. Within a compartment, neurons that represent odors synapse onto neurons that guide behavior. The strength of these synapses is modulated by a dopamine neuron that responds to one type of unconditioned stimuli (e.g., sugar), which implicates these synapses as a biological substrate for associative learning in insects. Modifications of these synapses is imbalanced in the sense that dopamine-induced modifications only weaken the synapses and are temporally sparse. In this work, we propose a simple mechanistic model of learning in the mushroom body that accounts for this


imbalanced learning. Our model is interpretable as implementing an algorithm for linear discriminant analysis, a classical statistical method for linearly separating feature vectors that belong to different classes. Our model makes testable predictions that provide clear contrasts with existing models.

## Introduction

Behavioral responses of animals are shaped in part by learned associations between sensory stimuli (e.g., odors) and unconditioned stimuli (e.g., sugar, heat or electric shock). A challenge in neuroscience is to understand the neural mechanisms that underlie associative learning. In invertebrates, the mushroom body is a well-studied brain region that plays a central role in olfactory associative learning [1–4]. The goal of this work is to propose a normative, mechanistic model of associative learning in the mushroom body that accounts for experimental observations and provides clear contrasts with existing models.

The mushroom body is segregated into functionally independent compartments [5], Fig 1. Within each compartment, Kenyon cells (KCs), which encode sparse odor representations [6], form synapses with the dendrites of mushroom body output neurons (MBONs), whose outputs guide learned behavior [7]. Associated with each compartment is a single Dopamine neuron (DAN) that responds to an unconditioned stimulus [8, 9], and projects its axon into the mushroom body compartment where it innervates the KC-MBON synapses to modulate plasticity, implicating the KC-MBON synapse as the synaptic substrate for associative learning in invertebrates.

Experimental evidence suggests that learning at the KC-MBON synapse is imbalanced in the sense that DAN-induced plasticity is one-sided and temporally sparse. In particular, co-activation of a KC and the DAN *weakens* the KC-MBON synapse (see Fig 1, right) and DAN-



**Fig 1. A simplified schematic of a mushroom body compartment.** In both the left and right panels, the mushroom body compartment, indicated by the shaded box, is innervated by the axons of multiple KCs, the dendrites from one MBON and the axon terminals of one DAN. The bi-colored circles at the intersections of the KC axons and the MBON dendrites denote the KC-MBON synapses. Faintly shaded cell bodies indicate inactive neurons and boldly shaded cell bodies indicate active neurons. In the left panel, the DAN is inactive. In the right panel, the DAN is active and co-activation of the KCs and the DAN weakens the associated KC-MBON synapses (as illustrated by the smaller synapses).

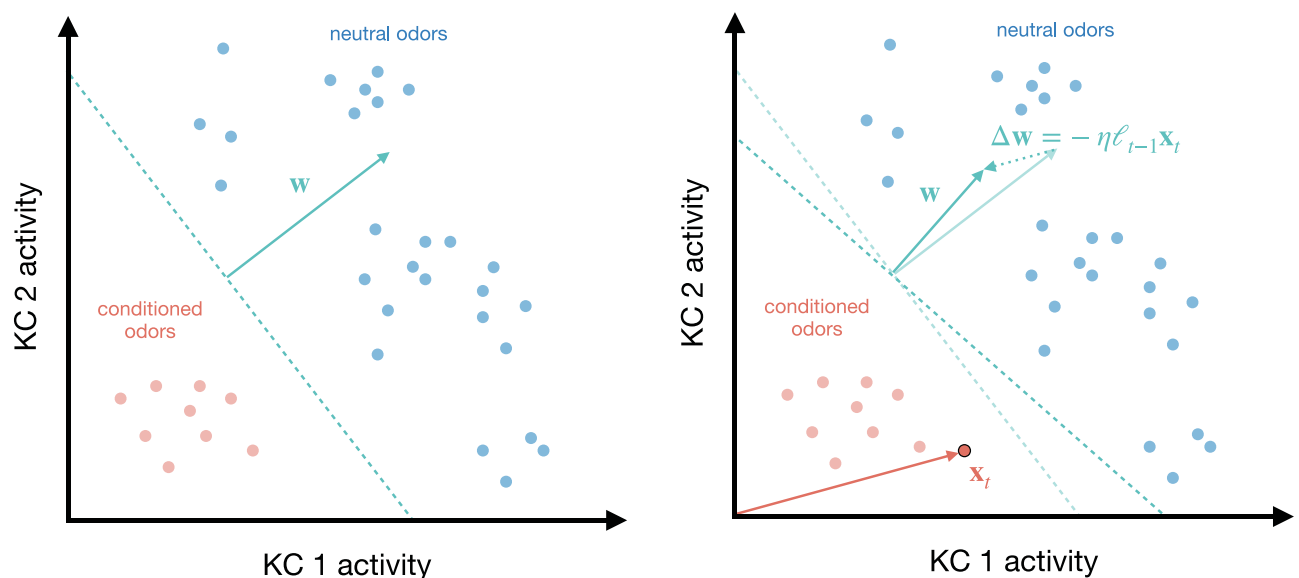
<https://doi.org/10.1371/journal.pcbi.1010864.g001>

induced plasticity is independent of the MBON activity [5]. This suggests that DAN-induced plasticity is one-sided and another mechanism such as homeostatic plasticity is responsible for strengthening the KC-MBON synapse. Furthermore, since each DAN responds to one type of unconditioned stimulus [2], which only constitutes a small fraction of all stimuli, the DAN activity is temporally sparse.

In this work, we propose a normative, mechanistic model of associative learning in the mushroom body that accounts for the imbalanced learning. We model each MBON as a linear discriminant analysis (LDA) classifier, which predicts if an associated unconditioned stimulus is present (the class label) given a KC odor representation (the feature vector). Under this interpretation, the KC-MBON synapses and an MBON bias term define a hyperplane in the high-dimensional space of KC odor representations that separates odor representations associated with the unconditioned stimulus from all other odor representations, Fig 2.

Here, ‘normative’ refers to the fact that our mechanistic model is interpretable as an algorithm for optimizing an LDA objective. The normative approach is top-down in the sense that first the circuit objective is proposed and then an optimization algorithm is derived and compared with known physiology. There are several advantages to this approach. First, it directly relates the circuit objective to its mechanism; for example, neural activities and synaptic weight updates are interpretable as steps in an algorithm for solving a relevant circuit objective. Second, the approach distills down what aspects of the physiological are essential for optimizing the circuit objective and what aspects are not captured by the objective. Third, normative models are often analytically tractable, which allows them to be analyzed for any input statistics without resorting to exhaustive numerical simulation.

To derive our algorithm, we start with a convex objective for LDA (in terms of the KC-MBON synaptic weights). The objective can be optimized in the offline setting by taking gradient descent steps with respect to the KC-MBON synaptic weights. To obtain an online



**Fig 2. Linear separation of odors in the space of KC activities.** Left: Illustration of the hyperplane  $\mathcal{H} = \{x : w \cdot x = b\}$  (dashed teal line) in the space of KC activities that separates conditioned odor responses from neutral odor responses. Each light red (resp. blue) dot denotes the KC response to a conditioned (resp. neutral) odor. The teal arrow denotes the vector of KC-MBON synaptic weights  $w$ , which is translated to show that it is orthogonal to the hyperplane  $\mathcal{H}$ . Right: Co-activation of the KCs and the DAN weakens the synaptic weights  $w$ . The KC activities  $x_t$  are denoted by the dark red dot with black border. The change of the synaptic weights  $\Delta w$  is in the direction  $-x_t$ . The hyperplane  $\mathcal{H}$  rotates to remain orthogonal to  $w$ . The change in bias  $\Delta b$ , which translates the hyperplane, is not depicted.

<https://doi.org/10.1371/journal.pcbi.1010864.g002>

algorithm that accounts for the imbalanced learning, we take advantage of the fact that DAN activity is temporally sparse to obtain online approximations of the input statistics. Finally, we show numerically that our algorithm performs well even when DAN activity is not temporally sparse.

Our model makes testable predictions that are a direct result of the learning imbalance. First, our model predicts that DAN-induced plasticity at the KC-MBON synapse is sensitive to the time elapsed since the DAN was last active. Second, our model predicts that if the DAN is never active, then the KC-MBON synapses adapt to align with the mean KC activity (normalized by the covariance of the KC activity).

## Results

### LDA model of the mushroom body compartment

We consider a simplified mushroom body compartment that consists of  $n$  KC axons, the axon terminals from one DAN and the dendrites of one MBON, Fig 1. At each time  $t = 1, 2, \dots$ , the vector  $\mathbf{x}_t \in \mathbb{R}^n$  encodes the KC activities and the scalar  $y_t \in \{0, 1\}$  indicates whether the DAN is active ( $y_t = 1$ ) or inactive ( $y_t = 0$ ). If the DAN is active, we refer to  $\mathbf{x}_t$  as a *conditioned odor response*, whereas if the DAN is inactive, we refer to  $\mathbf{x}_t$  as a *neutral odor response*. We assume the DAN activity is temporally sparse, which can be expressed mathematically as  $\pi_1 \ll 1$ , where  $\pi_1 := \langle y_t \rangle_t$  is the fraction of time that the DAN is active.

In our model, the MBON is a linear classifier that predicts the DAN activity  $y_t$  (class label) given the KC activities  $\mathbf{x}_t$  (feature vector). Let  $\mathbf{w} \in \mathbb{R}^n$  be a synaptic weight vector whose  $i^{\text{th}}$  component represents the strength of the synapse between the  $i^{\text{th}}$  KC and the MBON. At each time  $t$ , the KC activities  $\mathbf{x}_t$  are multiplied by the synaptic weight vector  $\mathbf{w}$  to generate the total input to the MBON, denoted  $c_t := \mathbf{w} \cdot \mathbf{x}_t$ . The output (firing rate) of the MBON is given by

$$z_t := \max(c_t - b, 0),$$

where  $b$  represents the ‘bias’ of the MBON; that is, the threshold below which the MBON does not fire. Under this interpretation, the KC-MBON synapses  $\mathbf{w}$  and MBON bias  $b$  define a hyperplane  $\mathcal{H} := \{\mathbf{x} \in \mathbb{R}^n : \mathbf{w} \cdot \mathbf{x} = b\}$  in the  $n$ -dimensional space of KC activities that separates conditioned odor responses from neutral odor responses, Fig 2. In this case,  $z_t > 0$  (resp.  $z_t = 0$ ) corresponds to the prediction  $y_t = 0$  (resp.  $y_t = 1$ ). In other words, the MBON is a linear classifier that is active when predicting there is no unconditioned stimulus and inactive when predicting there is an unconditioned stimulus, which is consistent with experimental observations [2].

We derive learning rules for the KC-MBON synaptic weights  $\mathbf{w}$  (and bias  $b$ ) that solve an LDA objective and are consistent with experimental observations [2, 5]. LDA is popular linear classification method that is optimal under the assumption that the neutral odor responses and conditioned odor responses are Gaussian with common covariance matrix, but works well in practice even when these assumptions do not hold [10].

Our starting point is the convex LDA objective

$$\min_{\mathbf{w}} L(\mathbf{w}), \quad L(\mathbf{w}) := -\mathbf{w} \cdot (\boldsymbol{\mu}_0 - \boldsymbol{\mu}_1) + \frac{1}{2} \mathbf{w}^\top \boldsymbol{\Sigma} \mathbf{w}, \quad (1)$$

where  $\boldsymbol{\mu}_0$  and  $\boldsymbol{\mu}_1$  denote the means of the neutral odor responses and conditioned odor responses, respectively, and  $\boldsymbol{\Sigma}$  denotes the covariance of the neutral odor response. In the off-line setting, we can minimize  $L(\mathbf{w})$  by taking gradient steps with respect to  $\mathbf{w}$ :

$$\Delta \mathbf{w} = \eta (\boldsymbol{\mu}_0 - \boldsymbol{\mu}_1 - \boldsymbol{\Sigma} \mathbf{w}), \quad (2)$$

where  $\eta > 0$  is the step size. However, computing the means  $\mu_0, \mu_1$  and covariance  $\Sigma$  requires the MBON to have access to the entire sequence of inputs, which is an unrealistic assumption.

To derive our online algorithm, we replace the averages  $\mu_0, \mu_1$  and  $\Sigma$  in Eq 2 with online estimates. When the DAN is inactive ( $y_t = 0$ ), we update the KC-MBON weights  $\mathbf{w}$  according to the homeostatic plasticity rule

$$y_t = 0 : \quad \Delta \mathbf{w} = \eta(\mu_{0,t} - (c_t - \zeta_t)(\mathbf{x}_t - \mu_{0,t})), \tag{3}$$

where  $\mu_{0,t}$  denotes the running estimate of the mean neutral odor response and  $\zeta_t$  denotes the running estimate of the mean total MBON input  $c_t$  conditioned on the DAN being inactive. Here,  $\mu_{0,t}$  and  $(c_t - \zeta_t)(\mathbf{x}_t - \mu_{0,t})$  are online estimates of  $\mu_0$  and  $\Sigma$ , respectively (see Methods section). The running means  $\mu_{0,t}$  and  $\zeta_t$  can be represented by biophysical quantities such as calcium concentrations at the pre- and postsynaptic terminals of the KC-MBON synapses.

When the DAN is active, we update the KC-MBON weights  $\mathbf{w}$  according to the following DAN-induced plasticity rule

$$y_t = 1 : \quad \Delta \mathbf{w} = -\eta \ell_{t-1} \mathbf{x}_t, \tag{4}$$

where  $\ell_{t-1}$  denotes the time elapsed since the last time the DAN was active; see Fig 2 (right) for a geometric interpretation of the plasticity rule. The update in Eq 4 is in line with experimental evidence showing that DAN-induced plasticity is independent of the MBON activity  $z_t$  and co-activation of the KCs and the DAN reduces the strength of the synapses between the KCs and the MBON [5]. Biologically, the scalar  $\ell_{t-1}$  can be represented as the sensitivity of the KC-MBON synapses to DAN-induced plasticity. Assuming the conditioned odor response is independent of the time elapsed between DAN activations, then on average, the update in Eq 4 is approximately equal to  $-\eta \mu_1$ , see Methods section. Therefore, the updates in Eqs 3–4 together account for all three terms in the offline update in Eq 2. The full model, including the bias updates (see Methods section), is summarized in Algorithm 1.

**Algorithm 1** LDA in the mushroom body compartment

```

input:  $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_T, y_T)$ 
initialize:  $\mathbf{w} = (w_1, \dots, w_n), b = 0, \ell_0 = 1, \eta > 0$ 
for  $t = 1, 2, \dots, T$  do
   $c_t \leftarrow \mathbf{w} \cdot \mathbf{x}_t$ 
   $z_t \leftarrow \max(c_t - b, 0)$ 
  if  $y_t = 0$  then
     $\mu_{0,t} \leftarrow \mu_{0,t-1} + \frac{1}{t}(\mathbf{x}_t - \mu_{0,t-1})$ 
     $\zeta_t \leftarrow \zeta_{t-1} + \frac{1}{t}(c_t - \zeta_{t-1})$ 
     $b \leftarrow b + \frac{1}{t}(\frac{1}{2}c_t - b)$ 
     $\mathbf{w} \leftarrow \mathbf{w} + \eta(\mu_{0,t} - (c_t - \zeta_t)(\mathbf{x}_t - \mu_{0,t}))$ 
     $\ell_t \leftarrow \ell_{t-1} + 1$ 
  else if  $y_t = 1$  then
     $b \leftarrow b + \frac{1}{t}(\frac{1}{2}\ell_{t-1}c_t - \log \ell_{t-1} - b)$ 
     $\mathbf{w} \leftarrow \mathbf{w} - \eta \ell_{t-1} \mathbf{x}_t$ 
     $\ell_t \leftarrow 1$ 
  end if
end for

```

Algorithm 1 only has one hyper-parameter—the learning rate  $\eta > 0$ —which corresponds to timescale for learning in the mushroom body compartment. Hige et al. [5] showed that mushroom body compartments have distinct timescales for learning, which can be modeled by choosing different learning rates  $\eta > 0$ .

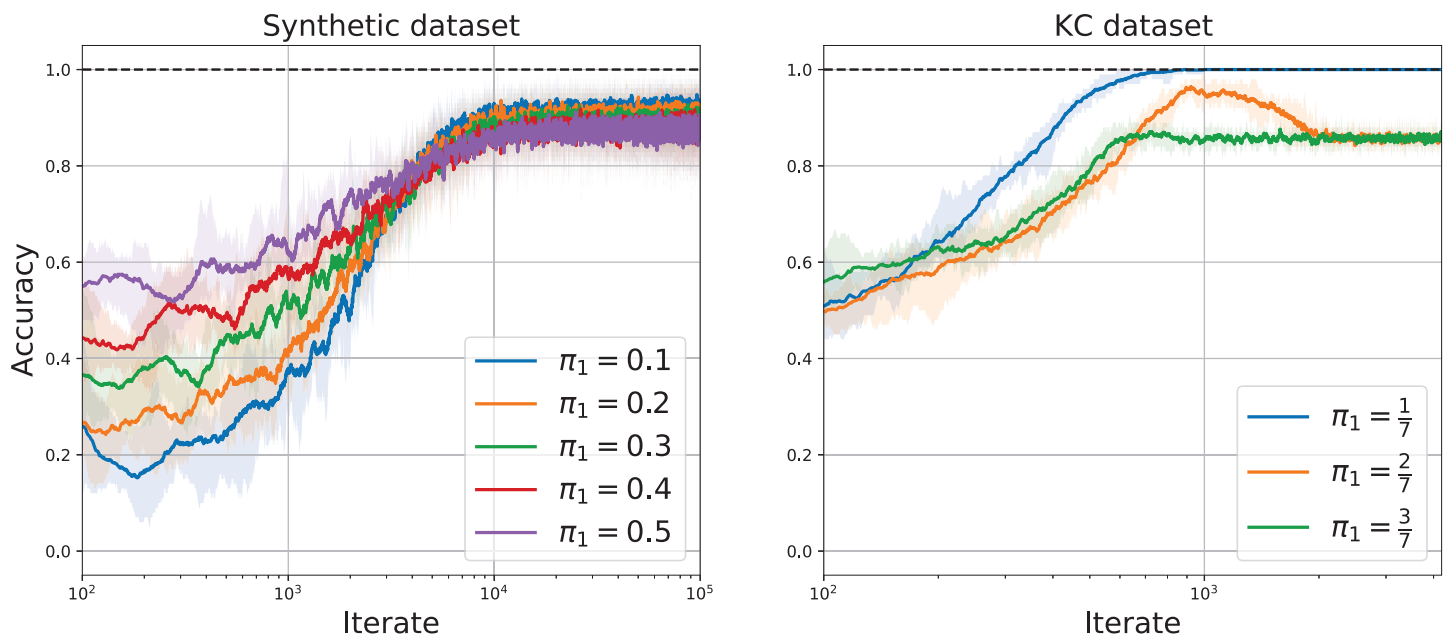
## Numerical experiments

Next, we test Algorithm 1 on synthetic and real datasets. We test our algorithm on inputs when our assumption  $\pi_1 \ll 1$  holds, but also on inputs when  $\pi_1 \approx 0.5$ . To evaluate our algorithm, we measure the running accuracy of the projections  $z_t$  over the previous  $\min(100, t)$  iterations, where the algorithm is accurate at the  $t^{\text{th}}$  iterate if  $z_t = 0$  and  $y_t = 1$  or if  $z_t > 0$  and  $y_t = 0$ .

**Synthetic dataset.** We begin by evaluating Algorithm 1 on a synthetic dataset generated by a mixture of 2 overlapping Gaussian distributions, so that the optimal accuracy is less than 1. The data points of the 2 classes are each drawn from a 2-dimensional mean with common covariance. We simulate datasets of  $10^5$  data points using the same mean and covariance in both classes but vary the frequency of class 1 samples encountered. We consider the cases  $\pi_1 = 0.1, 0.2, 0.3, 0.4, 0.5$ . In Fig 3 (left) we plot the error and the accuracy of our model for varying  $\pi_1$ . Remarkably, while the derivation of Algorithm 1 relied on the fact that  $\pi_1 \ll 1$ , the algorithm still performs well even when  $\pi_1 = 0.5$ .

**KC activities dataset.** We test our model on KC activities reported in [11]. Campbell et al. recorded odor-evoked KC responses in the fly mushroom body. The dataset we tested on contains the responses of 124 KCs in a single fly to the presentation of 7 odors, see [11, Figure 1]. To ensure the KC responses are well conditioned, we add Gaussian noise with covariance  $\epsilon \mathbf{I}_{124}$ , where  $\epsilon = 0.01$ . We apply Algorithm 1 to the KC dataset. We first consider the case that odor 1 denotes the class 1 odor and odors 2–7 denote the class 0 odors, so  $\pi_1 = \frac{1}{7}$ . We then consider the cases that odors 1–2 (resp. 1–3) odors denote the class 1 odors and the remaining odors denote the class 0 odors, so  $\pi_1 = \frac{2}{7}$  (resp.  $\pi_1 = \frac{3}{7}$ ). In Fig 3 (right) we plot the error and accuracy of our model for varying  $\pi_1$ . Impressively, the modified algorithm performs well (approximately 85% accuracy) even when the assumption  $\pi_1 \ll 1$  is violated.

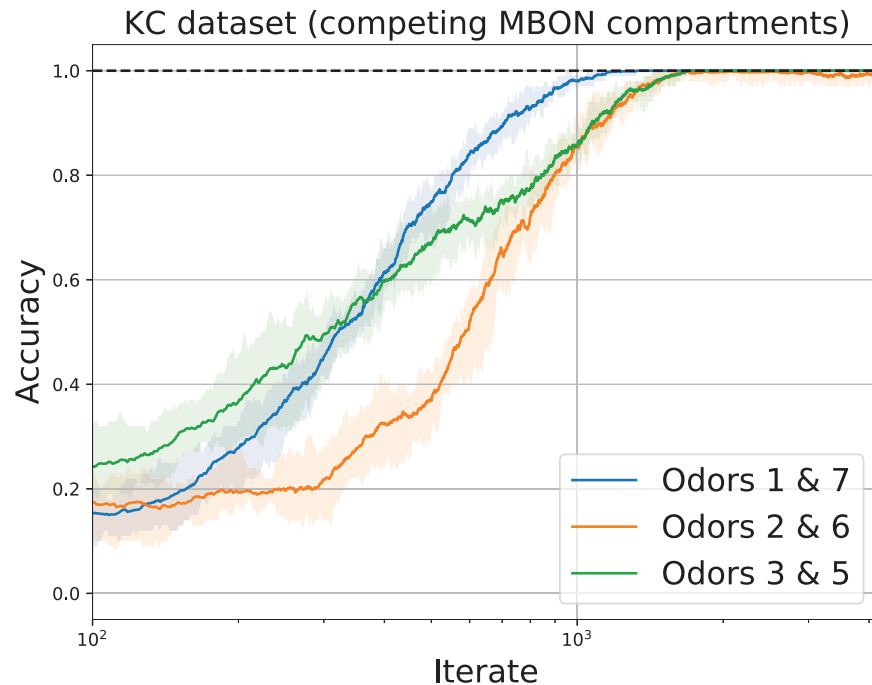
**Competing MBONs.** Using the KC activities dataset, we model 2 MBONs with competing valences by running 2 instances of Algorithm 1 in parallel with different class assignments for



**Fig 3. Performance of Algorithm 1.** Accuracy of Algorithm 1 on the synthetic datasets (left) and the KC dataset (right). Each line denotes the mean accuracy over 10 runs. Each shaded region indicates the area between the minimum and maximum accuracy over 10 runs.

<https://doi.org/10.1371/journal.pcbi.1010864.g003>





**Fig 4. Performance of competing MBONs.** Accuracy of 2 parallel runs of Algorithm 1 on the KC dataset to classify odors as aversive, attractive or neutral. Each line denotes the mean accuracy over 10 runs. Each shaded region indicates the area between the minimum and maximum accuracy over 10 runs.

<https://doi.org/10.1371/journal.pcbi.1010864.g004>

the odors. We consider the case that odor 1 is aversive, odor 7 is attractive and the remaining odors are neutral. For MBON 1 (resp. MBON 2), we assume that odor 1 (resp. odor 7) denotes the class 1 odor and odors 2–7 (resp. odors 1–6) denote the class 0 odors, so that MBON 1 (resp. MBON 2) activity promotes approach (resp. avoid) behavior. Let  $z_{i,t}$  denote the output of MBON  $i \in \{1, 2\}$ . At each iterate  $t$ , if odor 1 (resp. odor 7, odors 2–6) is presented, then the model is accurate if  $z_{1,t} > 0$  and  $z_{2,t} = 0$  (resp.  $z_{1,t} = 0$  and  $z_{2,t} > 0$ , resp.  $z_{1,t} > 0$  and  $z_{2,t} > 0$ ), and inaccurate otherwise. We then repeat the experiment two more times, but with odor 2 (resp. odor 3) labeled as aversive and odor 6 (resp. odor 5) as attractive. In Fig 4, we plot the performance of the competing MBONs.

## Discussion

### Summary

In this work, we proposed a normative model of the mushroom body compartment that accounts for imbalanced learning at the KC-MBON synapse. Testing our model on synthetic and real datasets shows that it performs well under a variety of conditions. In our model, DAN-induced plasticity at the KC-MBON synapse does not depend on the MBON activity, but rather on the time elapsed since the last time the DAN was active. This aspect of our model suggests testable predictions that provide clear contrasts with existing models of associative learning in the mushroom body.

### Model predictions

*Prediction 1*—In the absence of DAN activity, the KC-MBON synapses will align with the mean KC activity normalized by the covariance of their activities. When presented with

neutral odors, the synapses adapt according to the homeostatic update in Eq 3. Since this update is equal to  $\eta(\mu_0 - \Sigma \mathbf{w})$  on average (see [Methods](#) section), the KC-MBON synaptic weights equilibrate at  $\mathbf{w} = \Sigma^{-1} \mu_0$ . Experimentally, this prediction could be tested by first presenting a fly with neutral odors and simultaneously recording from multiple KCs and an MBON. The weights can be estimated from the neural activities (using, e.g., [12]) and compared with our prediction  $\mathbf{w} = \Sigma^{-1} \mu_0$ .

*Prediction 2*—DAN-induced plasticity is proportional to the time elapsed since the DAN was last active. According to update in Eq 4, DAN-induced plasticity is proportional to the time elapsed since the DAN was last active. Experimentally, this prediction could be tested by presenting a fly with conditioned odors with different time intervals between presentations and estimating the resulting change in the synaptic weights.

### Relation to existing models

There are a number of existing computational models of associative learning in the mushroom body [13–17], many of which are faithful to biophysical details and successfully capture important computational principles underlying associative learning in the mushroom body (see, e.g., [15]). Through extensive numerical simulations, these computational models can explain a number of phenomena. For example, Heurta and Nowotny [15] show that the organization of the mushroom body supports fast and robust associative learning, Bazhenov et al. [16] show that interactions between unsupervised and supervised forms of learning can explain how the timescale of associative learning depends on experimental conditions, and Peng and Chittka [17] show how complex forms of learning (e.g., peak shift) depend on different mechanistic aspects of learning in the mushroom body. In this work, we propose a top-down normative model of learning at the KC-MBON synapse, which contrasts with the bottom-up approach in these works that build models closely tied to physiological evidence. In this way, our model is interpretable as an algorithm for optimizing a circuit objective and the output can be predicted analytically for any environmental condition without needing to resort to numerical simulation. In addition, our normative model makes testable predictions that are in clear contrast with these models, providing a method for validating or invalidating our model.

In addition to these models, Bennett et al. [18] propose a reinforcement learning model of the KC-MBON synapses as minimizing reinforcement prediction errors. They first consider a model in which the reinforcement signal is computed as the difference between DAN activities, so their plasticity rule requires 2 DANs to innervate a single mushroom body compartment, which is in contrast to experiment evidence showing that most compartments only receive inputs from a single DAN [9]. To account for this experimental observation, they propose a heuristic modification that adds a constant source of synaptic potentiation, which can be viewed as a form of homeostatic plasticity and is in line with experimental evidence. However, the modification is not normative and can fail to minimize prediction errors.

A significant difference between our model and these existing models is that DAN-induced plasticity depends on  $\ell_{t-1}$ , the time elapsed since the DAN was last active. In our model, the variable  $\ell_{t-1}$  is critical for balancing homeostatic plasticity and DAN-induced plasticity. In [S1 Appendix](#), we consider a modification of our algorithm in which  $\ell_{t-1}$  is replaced by a fixed constant  $\ell^*$ .

### Comparison of LDA to other linear classification methods

LDA is a linear classifier that is optimal under strict assumptions on the inputs, so it is worth considering other linear classification methods such as logistic regression and support vector machines (SVMs). Logistic regression is classical method for estimating the probability of one



class versus another class. In terms of performance, there is evidence that there is not a substantial difference in the performance of logistic regression and LDA even when the assumptions for LDA are not met [19]. As a model of the insect mushroom body, we are unaware of an online algorithm for logistic regression that maps onto the mushroom body compartment and matches the experimental observations in [2, 5].

SVMs are flexible linear classifiers that do not make assumptions about the underlying data distribution. Huerta et al. [13, 15] proposed models of the mushroom body that are closely related to SVMs [20, 21]; however, the DAN-induced synaptic update rules depend on the MBON activity, which is in contrast to recent experimental evidence [5].

## Limitations

Our model is a dramatic simplification of the mushroom body focused on providing a normative account of learning at the KC-MBON synapse that can account for how balance between DAN-induced plasticity and homeostatic plasticity is optimally maintained. Consequently, our model does not account for a number of the physiological details. For example, in order to implement an LDA algorithm, we do not sign-constrain the synaptic weight vector  $\mathbf{w}$ , which violates Dale's law. In addition, we assume that the DAN activity is binary. In reality the DAN may fire at different rates depending on the strength of the unconditioned stimulus and the firing rate may affect the DAN-induced plasticity. We can modify our model to allow  $y_t$  to be any nonnegative scalar and replace the update in Eq 4 with  $\Delta \mathbf{w} = -\eta \ell_{t-1} y_t \mathbf{x}_t$ . However, in this case the algorithm is not derived from an objective function for LDA and so it is more challenging to understand the output. In addition to such simplifications, there are other features such as feedback connections in the mushroom body that have been recently discovered and are relevant for associative learning [7, 22], which are also not captured by our model.

## Methods

### Linear discriminant analysis

LDA is a statistical method for linear classification [23, section 4.3], which makes the following simplifying assumption: the conditional probability distributions  $p(\mathbf{x}|y=0)$  and  $p(\mathbf{x}|y=1)$  are both Gaussian with common full-rank  $n \times n$  covariance matrix  $\Sigma$ ; that is

$$p(\mathbf{x}|y=0) \sim \mathcal{N}(\boldsymbol{\mu}_0, \Sigma), \quad p(\mathbf{x}|y=1) \sim \mathcal{N}(\boldsymbol{\mu}_1, \Sigma),$$

where  $\boldsymbol{\mu}_0$  and  $\boldsymbol{\mu}_1$  denote the means of the class 0 and class 1 feature vectors. In this case, the optimal decision criteria for assigning class 0 (resp. class 1) to feature vector  $\mathbf{x}$  is  $\mathbf{w} \cdot \mathbf{x} > b$  (resp.  $\mathbf{w} \cdot \mathbf{x} < b$ ), where

$$\mathbf{w}_{\text{opt}} := \Sigma^{-1}(\boldsymbol{\mu}_0 - \boldsymbol{\mu}_1), \quad b_{\text{opt}} := \frac{1}{2} \mathbf{w} \cdot (\boldsymbol{\mu}_0 + \boldsymbol{\mu}_1) + \log \frac{\pi_1}{\pi_0}, \quad (5)$$

and  $\pi_i$  denotes the probability that a samples belongs to class  $i$ , for  $i = 0, 1$ . In particular, the hyperplane  $\mathcal{H} = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{w}_{\text{opt}} \cdot \mathbf{x} = b_{\text{opt}}\}$  defines the optimal separation boundary for predicting whether a feature vector belongs to class 0 or class 1. While LDA assumes a specific generative model, it performs well in practice even when the assumptions do not hold [10].

The optimal weights  $\mathbf{w}_{\text{opt}}$  can be expressed as the solution of the convex minimization problem in Eq 1, which we can solve for by taking the gradient descent steps (Eq 2). Formally, taking the step size  $\eta$  to zero in Eq 2 yields the linear gradient flow

$$\dot{\mathbf{w}}(t) = \boldsymbol{\mu}_0 - \boldsymbol{\mu}_1 - \Sigma \mathbf{w}(t),$$

whose solution is given by

$$\mathbf{w}(t) = e^{-\Sigma t} \mathbf{w}(0) + (\mathbf{I}_n - e^{-\Sigma t}) \Sigma^{-1} (\boldsymbol{\mu}_0 - \boldsymbol{\mu}_1).$$

In particular, we see that the solution  $\mathbf{w}(t)$  converges exponentially to the optimal solution  $\mathbf{w}_{\text{opt}}$ .

### An online algorithm for imbalanced learning

In the online setting, the class means  $\boldsymbol{\mu}_0, \boldsymbol{\mu}_1$  and the covariance  $\Sigma$  are not available. Instead, at each time  $t$  the algorithm has access to the feature vector  $\mathbf{x}_t$  and class label  $y_t$ . To derive our online algorithm, we make online approximations of the offline quantities  $\boldsymbol{\mu}_0, \boldsymbol{\mu}_1$  and  $\Sigma$  that are based on the fact that the unconditioned stimuli are sparse in time, i.e.,  $\pi_1 \ll 1$ , where we recall that  $\pi_1$  denotes the proportion of conditioned odors. First, we note that we can rewrite the sample class means

$$\boldsymbol{\mu}_0 = \frac{1}{\pi_0} \langle (1 - y_t) \mathbf{x}_t \rangle_t, \quad \boldsymbol{\mu}_1 = \frac{1}{\pi_1} \langle y_t \mathbf{x}_t \rangle_t,$$

where  $\pi_0 := \langle 1 - y_t \rangle_t \approx 1$  is the fraction of odors that are neutral, and the sample covariance

$$\Sigma = \langle (1 - y_t) (\mathbf{x}_t - \boldsymbol{\mu}_0) (\mathbf{x}_t - \boldsymbol{\mu}_0)^\top \rangle_t + \langle y_t (\mathbf{x}_t - \boldsymbol{\mu}_1) (\mathbf{x}_t - \boldsymbol{\mu}_1)^\top \rangle_t.$$

**Estimating the mean response to a neutral odor and the covariance.** Since  $\pi_0 \approx 1$ , we approximate

$$\boldsymbol{\mu}_0 \approx \langle (1 - y_t) \mathbf{x}_t \rangle_t, \quad \Sigma \approx \langle (1 - y_t) (\mathbf{x}_t - \boldsymbol{\mu}_0) (\mathbf{x}_t - \boldsymbol{\mu}_0)^\top \rangle_t. \tag{6}$$

Therefore, in the online setting, we can keep a running estimate of  $\boldsymbol{\mu}_0$  and  $\zeta := \mathbf{w} \cdot \boldsymbol{\mu}_0 \approx \langle (1 - y_t) c_t \rangle_t$ , where we recall that  $c_t = \mathbf{w} \cdot \mathbf{x}_t$ , by performing the updates

$$\boldsymbol{\mu}_{0,t} \leftarrow \boldsymbol{\mu}_{0,t-1} + \frac{1}{t} (1 - y_t) (\mathbf{x}_t - \boldsymbol{\mu}_{0,t-1}), \quad \zeta_t \leftarrow \zeta_{t-1} + \frac{1}{t} (1 - y_t) (c_t - \zeta_{t-1}). \tag{7}$$

In view of Eq 6 and the definitions of  $c_t$  and  $\zeta$ , we can replace  $\Sigma \mathbf{w}$  with the online approximation

$$(1 - y_t) (\mathbf{x}_t - \boldsymbol{\mu}_0) (\mathbf{x}_t - \boldsymbol{\mu}_0)^\top \mathbf{w} = (1 - y_t) (c_t - \zeta_t) (\mathbf{x}_t - \boldsymbol{\mu}_{0,t}). \tag{8}$$

We replace the first and third terms in the offline update in Eq 2,  $\eta(\boldsymbol{\mu}_0 - \Sigma \mathbf{w})$ , with the online estimate  $\eta(1 - y_t) (\boldsymbol{\mu}_{0,t} - (c_t - \zeta_t) (\mathbf{x}_t - \boldsymbol{\mu}_{0,t}))$ .

**Estimating the mean response to a conditioned odor.** To obtain an online approximation of  $\boldsymbol{\mu}_1$ , we first note that  $\frac{1}{\pi_1}$  is approximately equal to the average time elapsed between class 1 samples. To see this, let  $t_1, t_2, \dots$  denote the subset of times such that  $y_t = 1$ . Then, letting  $t_0 = 0$ , we have

$$\frac{1}{\pi_1} = \lim_{r \rightarrow \infty} \frac{t_r}{r} = \lim_{r \rightarrow \infty} \frac{1}{r} \sum_{j=1}^r (t_j - t_{j-1}).$$

Thus, in the online setting, when the  $j^{\text{th}}$  class 1 sample is presented (i.e.,  $y_t = 1$ ), we can use the

time elapsed since the last class 1 sample,  $t_j - t_{j-1}$ , as an online estimate of  $\frac{1}{\pi_1}$ . Setting  $\ell_0 = 1$  and

$$\ell_t = \begin{cases} \ell_{t-1} + 1 & \text{if } y_t = 0 \\ 1 & \text{if } y_t = 1, \end{cases}$$

we see that at time  $t$  such that  $y_t = 1$ ,  $\ell_{t-1}$  denotes the time elapsed since the last class 1 sample, so  $\langle \ell_{t-1} | y_t = 1 \rangle_t = \frac{1}{\pi_1}$ . Assuming that the variables  $\ell_{t-1}$  and  $\mathbf{x}_t$  are independent given  $y_t = 1$ —i.e., the KC representation  $\mathbf{x}_t$  of a conditioned odor is independent of the time elapsed since the last conditioned odor  $\ell_{t-1}$ —we see that  $\boldsymbol{\mu}_1 = \langle \ell_{t-1} | y_t = 1 \rangle_t \langle \mathbf{x}_t | y_t = 1 \rangle_t = \langle \ell_{t-1} \mathbf{x}_t | y_t = 1 \rangle_t$ . We replace the second term in the offline update in Eq 2,  $-\eta \boldsymbol{\mu}_1$ , with the online approximation  $-\eta y_t \ell_{t-1} \mathbf{x}_t$ .

**Estimating the bias.** To estimate the bias  $b$ , we note that because  $\pi_0 \approx 1$  and  $\frac{1}{\pi_1} \approx \langle \ell_{t-1} | y_t = 1 \rangle_t$

$$\log \frac{\pi_1}{\pi_0} \approx \log \pi_1 \approx -\log \langle \ell_{t-1} | y_t = 1 \rangle_t \leq \langle \log \ell_{t-1} | y_t = 1 \rangle_t,$$

where the final inequality follows from the fact that log is concave and Jensen’s inequality (with equality holding when the variance of  $\ell_{t-1}$  given  $y_t = 1$  is zero). Thus, assuming the variance of the time elapsed between conditioned odors is small, we can estimate the bias  $b$  in the online setting with the updates:

$$b \leftarrow b + \frac{1}{t} \left[ \frac{1}{2} (1 - y_t + \ell_{t-1} y_t) \mathbf{w}^\top \mathbf{x}_t - y_t \log \ell_{t-1} - b \right].$$

Substituting these approximations into the offline update rules in Eq 2 yields our online algorithm (Algorithm 1).

In view of Jensen’s inequality, if the variance of the time elapsed between conditioned odors is large, then the bias  $b$  will be overestimated, meaning that the MBON will be less active than optimal. In other words, irregular intervals between DAN activity biases the MBON to be less active (i.e., predict that the unconditioned stimulus is present more often).

### Details of numerical experiments

The experiments were performed on an Apple iMac with a 3.2 GHz 8-Core Intel Xeon W processor. For each experiment, we used a learning rate of the form

$$\eta_t = \frac{\eta_0}{1 + \gamma t}.$$

We chose the parameters  $\eta_0$  and  $\gamma$  by performing a grid search over  $\eta_0 \in \{1, 10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}\}$  and  $\gamma \in \{10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}, 10^{-6}\}$ . The optimal parameters for the synthetic dataset (resp. KC dataset) are  $\eta_0 = 10^{-1}$  and  $\gamma = 10^{-3}$  (resp.  $\eta_0 = 10^{-1}$  and  $\gamma = 10^{-4}$ ).

### Supporting information

**S1 Appendix. Comparison with a modified algorithm.** We consider a modification of Algorithm 1 in which the DAN-induced plasticity of the KC-MBON synapses does not depend on the time elapsed since the last time the DAN was active.

(PDF)

## Acknowledgments

We are grateful to Lucy Reading-Ikkanda for creating Fig 1. We thank Yanis Bahroun, Siavash Golkar, Jason Moore and Tiberiu Teşileanu for helpful feedback on an earlier draft of this work.

## Author Contributions

**Conceptualization:** David Lipshutz, Aneesh Kashalika, Dmitri B. Chklovskii.

**Data curation:** David Lipshutz, Shiva Farashahi.

**Formal analysis:** David Lipshutz, Aneesh Kashalika, Shiva Farashahi.

**Investigation:** David Lipshutz, Aneesh Kashalika, Shiva Farashahi.

**Methodology:** David Lipshutz, Dmitri B. Chklovskii.

**Project administration:** David Lipshutz.

**Software:** David Lipshutz, Aneesh Kashalika, Shiva Farashahi.

**Supervision:** David Lipshutz, Dmitri B. Chklovskii.

**Validation:** David Lipshutz.

**Visualization:** David Lipshutz.

**Writing – original draft:** David Lipshutz.

**Writing – review & editing:** David Lipshutz, Shiva Farashahi, Dmitri B. Chklovskii.

## References

1. Heisenberg M. Mushroom body memoir: from maps to models. *Nature Reviews Neuroscience*. 2003; 4(4):266–275. <https://doi.org/10.1038/nrn1074> PMID: 12671643
2. Oswald D, Waddell S. Olfactory learning skews mushroom body output pathways to steer behavioral choice in *Drosophila*. *Current Opinion in Neurobiology*. 2015; 35:178–184. <https://doi.org/10.1016/j.conb.2015.10.002> PMID: 26496148
3. Eichler K, Li F, Litwin-Kumar A, Park Y, Andrade I, Schneider-Mizell CM, et al. The complete connectome of a learning and memory centre in an insect brain. *Nature*. 2017; 548(7666):175. <https://doi.org/10.1038/nature23455> PMID: 28796202
4. Modi MN, Shuai Y, Turner GC. The *Drosophila* mushroom body: from architecture to algorithm in a learning circuit. *Annual Review of Neuroscience*. 2020; 43:465–484. <https://doi.org/10.1146/annurev-neuro-080317-0621333> PMID: 32283995
5. Hige T, Aso Y, Modi MN, Rubin GM, Turner GC. Heterosynaptic plasticity underlies aversive olfactory learning in *Drosophila*. *Neuron*. 2015; 88(5):985–998. <https://doi.org/10.1016/j.neuron.2015.11.003> PMID: 26637800
6. Honegger KS, Campbell RA, Turner GC. Cellular-resolution population imaging reveals robust sparse coding in the *Drosophila* mushroom body. *Journal of Neuroscience*. 2011; 31(33):11772–11785. <https://doi.org/10.1523/JNEUROSCI.1099-11.2011> PMID: 21849538
7. Eschbach C, Fushiki A, Winding M, Schneider-Mizell CM, Shao M, Arruda R, et al. Multilevel feedback architecture for adaptive regulation of learning in the insect brain. *bioRxiv*. 2019; p. 649731.
8. Waddell S. Reinforcement signalling in *Drosophila*; dopamine does it all after all. *Current Opinion in Neurobiology*. 2013; 23(3):324–329. <https://doi.org/10.1016/j.conb.2013.01.005> PMID: 23391527
9. Aso Y, Hattori D, Yu Y, Johnston RM, Iyer NA, Ngo TT, et al. The neuronal architecture of the mushroom body provides a logic for associative learning. *Elife*. 2014; 3:e04577. <https://doi.org/10.7554/eLife.04577> PMID: 25535793
10. Michie D, Spiegelhalter DJ, Taylor CC, editors. *Machine Learning, Neural and Statistical Classification*. Ellis Horwood; 1994.

11. Campbell RA, Honegger KS, Qin H, Li W, Demir E, Turner GC. Imaging a population code for odor identity in the *Drosophila* mushroom body. *Journal of Neuroscience*. 2013; 33(25):10568–10581. <https://doi.org/10.1523/JNEUROSCI.0682-12.2013> PMID: 23785169
12. Linderman S, Stock CH, Adams RP. A framework for studying synaptic plasticity with neural spike train data. *Advances in Neural Information Processing Systems*. 2014; 27.
13. Huerta R, Nowotny T, García-Sánchez M, Abarbanel HDI, Rabinovich MI. Learning classification in the olfactory system of insects. *Neural Computation*. 2004; 16(8):1601–1640. <https://doi.org/10.1162/089976604774201613> PMID: 15228747
14. Smith D, Wessnitzer J, Webb B. A model of associative learning in the mushroom body. *Biological Cybernetics*. 2008; 99(2):89–103. <https://doi.org/10.1007/s00422-008-0241-1> PMID: 18607623
15. Huerta R, Nowotny T. Fast and robust learning by reinforcement signals: explorations in the insect brain. *Neural Computation*. 2009; 21(8):2123–2151. <https://doi.org/10.1162/neco.2009.03-08-733> PMID: 19538091
16. Bazhenov M, Huerta R, Smith BH. A computational framework for understanding decision making through integration of basic learning rules. *Journal of Neuroscience*. 2013; 33(13):5686–5697. <https://doi.org/10.1523/JNEUROSCI.4145-12.2013> PMID: 23536082
17. Peng F, Chittka L. A simple computational model of the bee mushroom body can explain seemingly complex forms of olfactory learning and memory. *Current Biology*. 2017; 27(2):224–230. <https://doi.org/10.1016/j.cub.2016.10.054> PMID: 28017607
18. Bennett JE, Philippides A, Nowotny T. Learning with reinforcement prediction errors in a model of the *Drosophila* mushroom body. *Nature Communications*. 2021; 12(1):1–14. <https://doi.org/10.1038/s41467-021-22592-4>
19. Lei PW, Koehly LM. Linear discriminant analysis versus logistic regression: A comparison of classification errors in the two-group case. *The Journal of Experimental Education*. 2003; 72(1):25–49. <https://doi.org/10.1080/00220970309600878>
20. Huerta R, Vembu S, Amigó JM, Nowotny T, Elkan C. Inhibition in multiclass classification. *Neural Computation*. 2012; 24(9):2473–2507. [https://doi.org/10.1162/NECO\\_a\\_00321](https://doi.org/10.1162/NECO_a_00321) PMID: 22594829
21. Huerta R. Learning pattern recognition and decision making in the insect brain. In: *AIP Conference Proceedings*. vol. 1510. American Institute of Physics; 2013. p. 101–119.
22. Li F, Lindsey JW, Marin EC, Otto N, Dreher M, Dempsey G, et al. The connectome of the adult *Drosophila* mushroom body provides insights into function. *Elife*. 2020; 9:e62576. <https://doi.org/10.7554/eLife.62576> PMID: 33315010
23. Hastie T, Tibshirani R, Friedman JH. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. vol. 2. Springer; 2009.