

1

Biological Mechanisms for Learning: A Computational Model of Olfactory Learning in the *Manduca sexta* Moth, with Applications to Neural Nets

Charles B. Delahunt 1,* , Jeffrey A. Riffell 2 and J. Nathan Kutz 3

¹ Department of Electrical Engineering, University of Washington, Seattle, WA, USA; Computational Neuroscience Center, University of Washington, Seattle, WA, USA.

- ² Department of Biology, University of Washington, Seattle, WA, USA.
- ³ Department of Applied Mathematics, University of Washington, Seattle, WA, USA.

Correspondence*: Corresponding Author delahunt@uw.edu

2 ABSTRACT

The insect olfactory system, which includes the antennal lobe (AL), mushroom body (MB), and 3 ancillary structures, is a relatively simple neural system capable of learning. Its structural features, 4 which are widespread in biological neural systems, process olfactory stimuli through a cascade 5 of networks where large dimension shifts occur from stage to stage and where sparsity and 6 randomness play a critical role in coding. Learning is partly enabled by a neuromodulatory reward 7 mechanism of octopamine stimulation of the AL, whose increased activity induces synaptic weight 8 9 updates in the MB through Hebbian plasticity. Enforced sparsity in the MB focuses Hebbian growth on neurons that are the most important for the representation of the learned odor. Based 10 upon current biophysical knowledge, we have constructed an end-to-end computational firing-rate 11 model of the Manduca sexta moth olfactory system which includes the interaction of the AL and 12 13 MB under octopamine stimulation. Our model is able to robustly learn new odors, and neural firing rates in our simulations match the statistical features of *in vivo* firing rate data. From a biological 14 perspective, the model provides a valuable tool for examining the role of neuromodulators, like 15 octopamine, in learning, and gives insight into critical interactions between sparsity, Hebbian 16 growth, and stimulation during learning. Our simulations also inform predictions about structural 17 details of the olfactory system that are not currently well-characterized. From a machine learning 18 perspective, the model yields bio-inspired mechanisms that are potentially useful in constructing 19 neural nets for rapid learning from very few samples. These mechanisms include high-noise 20 layers, sparse layers as noise filters, and a biologically-plausible optimization method to train the 21 network based on octopamine stimulation, sparse layers, and Hebbian growth. 22

23 Keywords: antenna lobe, mushroom body, learning, olfaction, octopamine, sparsity, Hebbian, neural nets

24 Main body has 9950 words, 7 figures.

1 INTRODUCTION

Learning is a vital function of biological neural networks, yet the underlying mechanisms responsible for 25 26 robust and rapid learning are not well understood. The insect olfactory network, and the moth's olfactory network (MON) in particular (e.g. in the Manduca sexta moth), is a comparatively simple biological 27 28 neural network capable of learning Riffell et al. (2008); Daly et al. (2001), and makes an ideal model 29 organism for characterizing the mechanics of learning. It is amenable to interrogation through experimental neural recordings of key, well-understood structural components including the antennal lobe (AL) Wilson 30 31 (2008) and mushroom body (MB) Campbell and Turner (2010). In addition, the AL-MB contain many 32 structural motifs that are widespread in biological neural systems. These motifs include: (i) the use of neuromodulators (octopamine and dopamine) in learning Dacks et al. (2012), (ii) a cascading networks 33 34 structure Masse et al. (2009), (iii) large changes in dimensionality (ie numbers of neurons) between 35 networks Laurent (2002), (iv) sparse encodings of data in high-dimensional networks Honegger et al. 36 (2011), (v) random connections Caron (2013), and (vi) the presence of noisy signals Galizia (2014). 37 Bio-inspired design principles suggest that each of the features has high value to the olfactory system. The mechanism of octopamine/dopamine release during learning is of particular interest, since it is not 38 well-understood how this stimulation promotes the construction of new sparse codes in the MB. 39

In this work, we build a computational model of the moth olfactory network, including both AL and MB, 40 41 that is closely aligned with both the known biophysics of the moth AL-MB and in vivo neural firing rate data, and that includes the dynamics of octopamine stimulation. We then run simulations to investigate how 42 the system components interact to learn new odors. When building our Network Model we have consulted 43 the literature, subject to a caveat: Moths and flies are similar enough that findings in flies (Drosophila) can 44 generally be transferred to the moth; but locusts and honeybees are more complex, and some findings in 45 these insects do not safely transfer, while other findings are general enough to readily apply Riffell et al. 46 (2009a). 47

48 There exist several computational models based on the insect brain Nowotny et al. (2005); Bazhenov et al. (2013); Peng and Chittka (2017); Roper et al. (2017); Mosqueiro and Huerta (2014); Arena et al. (2013); 49 50 Faghihi et al. (2017); Huerta and Nowotny (2009); Nowotny (2009); Jortner et al. (2007); García-Sanchez 51 and Huerta (2003). Because the MB is central to memory, these models focus on the sparsely-firing, high-dimensional MB plus readout neuron(s), leaving aside the AL or treating it as a "pass-through". Some 52 53 of these models incorporate forms of Hebbian plasticity Nowotny (2009); Nowotny et al. (2005); Peng 54 and Chittka (2017); Bazhenov et al. (2013); Huerta and Nowotny (2009) In general these models are not closely tied to a particular organism (though they are usually inspired by locusts or honeybees), so they are 55 56 "top-down" designs, allowing freedom with model parameters in the service of capturing general behaviors (indeed, Peng and Chittka (2017) points out the advantages of this more general approach). 57

Key findings of these models include: The value (for class separation) of the fan-out into the high-58 dimensional, sparsely-firing MB; theoretical calculations of parameters such as optimal AL->MB 59 connectivity García-Sanchez and Huerta (2003); Nowotny (2009); the value of random neural connectivity 60 Nowotny (2009); Peng and Chittka (2017); the ability of the simple MB structure to capture complex 61 behaviors Bazhenov et al. (2013); Roper et al. (2017); Peng and Chittka (2017); robust performance over 62 wide tuning parameter ranges Roper et al. (2017); Huerta and Nowotny (2009); and the generalized learning 63 skills of the insect MB, given a Hebbian update mechanism Peng and Chittka (2017); Huerta and Nowotny 64 (2009); Faghihi et al. (2017); Arena et al. (2013); Mosqueiro and Huerta (2014). 65

Our computational model of learning in the MON is distinct in four key ways from these general, MB-66 67 focused studies. (i) We model the architecture and neural dynamics of the whole system. This includes detailed internal wiring of the AL, the MB, inhibition from the Lateral Horn, octopamine stimulation 68 69 during learning, Hebbian plasticity, and an extrinsic (readout) neuron downstream. Linking careful models 70 of the AL and MB fills a particular gap in the literature called out by Mosqueiro and Huerta (2014). (ii) We include octopamine stimulation in the dynamics equations. The neuromodulator octopamine (similarly 71 dopamine) is essential to learning in the insect olfactory network Hammer and Menzel (1995, 1998), but 72 73 it has not (to our knowledge) been incorporated into a computational model. Thanks to a unique dataset 74 we are able to model this key component, and trace its effects on the AL, the MB, and learning. (iii) We tether our model architecture to a particular insect system (the M. sexta moth). As part of this tethering, 75 (iv) we calibrate the model's AL firing rate behavior to a dataset of in vivo neural recordings of moths 76 77 during learning, ie while exposed to both odors and octopamine. Thus our model is built "bottom-up", with parameters as far as possible determined by a particular organism. For example, we set AL \rightarrow MB 78 79 connectivity and weights based on clues in biophysical studies and calibration to *in vivo* firing rates, rather than using the model to explore theoretically optimal values. This is an opposite, and complementary, 80 approach to the studies cited above, and in combination with our unique *in vivo* octopamine data yields in 81 several new findings, as well as some findings that reinforce those of previous studies but from a different 82 83 angle.

84 We thus create a full, end-to-end neural network model (hereafter "Network Model") that demonstrates 85 robust learning behavior while also tightly matching the structure and behavior of a particular biological 86 system. This approach has three advantages: (i) we can meaningfully compare Network Model simulation 87 output to experimental data in order to tune model parameters; (ii) findings from our simulation results 88 can map back to the original biological system to offer meaningful biophysical insights; and (iii) Network 89 Model simulations allow us to study how the various key elements in the moth's toolkit (eg AL, MB, 90 octopamine, and Hebbian updates) interact to enable learning. We can thus derive bio-inspired insight 91 into the mathematical framework that enables rapid and robust learning in neural nets. Specifically, our 92 experiments elucidate mechanisms for fast learning from noisy data that rely on cascaded networks, sparsity, 93 and Hebbian plasticity.

94 These mechanisms have potential applications to engineered neural nets (NNs). NNs have emerged 95 as a dominant mathematical paradigm for characterizing neural processing and learning, honoring their 96 inspiration in the Nobel prize winning work of Hubel and Wiesel on the primary visual cortex of cats Hubel and Wiesel (1962). These seminal experiments showed that networks of neurons were organized 97 98 in hierarchical layers of cells for processing visual stimulus. The first mathematical model of a neural 99 network, the *Neocognitron* in 1980 Fukushima (1980), had many of the characteristic features of today's 100 deep neural networks (DNNs). However, many of the biological motifs listed above (for insect AL-MBs) 101 are largely absent from engineered NNs, whose principles and building blocks are biologically implausible 102 even as DNNs have achieved great success Schmidhuber (2015); LeCun (2015). For example, the AL-MB 103 interaction with octopamine, Hebbian plasticity, and sparsity operates in a fundamentally different manner than the backprop optimization used in DNNs, and it also succeeds at tasks (eg rapid learning) where 104 105 DNNs struggle. These biological mechanisms thus represent a potential opportunity to expand the set of structural and algorithmic tools available for ML tasks. We seek to characterize an actionable set of 106 biological elements, a "biological toolkit", that can be assembled into complementary NN architectures 107 108 or inserted into engineered NNs, and that are capable of rapid and robust learning from very few training 109 samples, an ability common in biological NNs but challenging for today's DNNs.

To briefly summarize the AL-MB network: It is organized as a feed-forward cascade of five distinct 110 networks, as well as a reward mechanism Martin et al. (2011); Kvello et al. (2009); Hildebrand (1996). 111 Roughly 30,000 noisy chemical receptor neurons (RNs) detect odor and send signals to the Antenna 112 Lobe (AL) Masse et al. (2009). The AL acts as a pre-amplifier, providing gain control and sharpening 113 114 of odor representations Bhandawat et al. (2007); Kuebler et al. (2012). It contains roughly 60 isolated units (glomeruli) Huetteroth and Schachtner (2005), each focused on a single odor stimuli feature Martin 115 et al. (2011); Christensen et al. (1995). Glomeruli laterally inhibit each other, and project odor codes to 116 the Mushroom Body (MB). AL neurons are noisy Galizia (2014); Lei et al. (2011). The MB contains 117 about 4000 Kenyon Cells (KCs). These fire sparsely and encode odor signatures as memories Perisse et al. 118 (2013); Campbell and Turner (2010); Honegger et al. (2011); Balkenius and Hansson (2012). MB sparsity 119 is enforced by global inhibition from the Lateral Horn Bazhenov and Stopfer (2010). Extrinsic Neurons 120 (ENs), numbering ~ 10 's, are believed to be "readout neurons" that interpret the KC codes Campbell et al. 121 (2013); Hige et al. (2015). In response to reward (sugar at the proboscis), a large neuron sprays octopamine 122 globally over the AL and MB, causing generalized stimulation of neurons Riffell et al. (2012); Dacks 123 et al. (2008). Learning does not occur without this octopamine input Hammer and Menzel (1995, 1998). 124 The connections into the KCs (AL \rightarrow KCs) and out of the KCs (KCs \rightarrow ENs) are plastic during learning 125 Cassenaer and Laurent (2007); Masse et al. (2009). Figure 1 gives a system schematic (A) along with 126 typical firing rate (FR) timecourses (from simulation) for neurons in each network (B). More network 127 details are given in Methods. 128

2 RESULTS

We first show the calibration of our Network Model to *in vivo* data. We then describe neural behaviors of the Network Model during learning and give results of learning experiments. Finally, we give results of experiments on MB sparsity.

132 2.1 Calibration of Model

The Network Model was calibrated to behave in a statistically similar way to three sets of *in vivo* data measuring projection neuron (PN) firing rate (FR) activity in the AL (See Methods for details): (i) PN spike counts with odor but without octopamine: 129 units with FR>1 spike/sec, (ii) PN spike counts with odor and with octopamine: 180 units with FR>1 spike/sec, and (iii) PN spike counts with odor, with and without octopamine: 52 units with FR>1 spike/sec.

Due to the limited number of experimental units, only qualitative comparisons of the model and ex-138 periment could be made: Excessive tuning of the model parameters would have served only to overfit 139 the particular data, rather than matching true PN behavior distributions or, more importantly, the general 140 learning behavior of the moth. Figure 2 shows the close match of typical Network Model PN statistics to 141 in vivo PN statistics based on mean (μ) and variance (σ) of spontaneous FRs and odor responses, both 142 without and with octopamine (details of metrics are given in Methods). Importantly, Figure 2 shows 143 significant octopamine-modulated increase in PN FR activity in the Network Model, consistent with in 144 vivo experiments involving octopamine stimulation. 145

There is limited experimental data measuring the FR activity of Kenyon cells (KC) in the MB, and no data to our knowledge measuring KC in response to octopamine stimulation. However, we note that the behavior of KCs during the application of octopamine to the AL, either with or without odor, is not an artifact of parameter tuning. Rather, it follows from the tuning the AL to match *in vivo* data. Specifically, PN FRs at baseline (with no odor or octopamine), with odor alone, with octopamine alone, and with odor 151 and octopamine, are all determined by calibration of the model to *in vivo* data. KCs respond only to PNs 152 and to inhibition from the LH (See Figure 1). Calibrating the KC baseline response in the absence of 153 octopamine to *in vivo* data in Turner et al. (2008) fixes the feed-forward connections from PNs. Assumed 154 in this model, due to lack of biophysical evidence, is that octopamine has no direct stimulative effect on 155 KC FRs (we do posit that it acts as an "on switch" for plasticity. Thus KC behavior with octopamine is 156 fully determined once the model is tuned to PN data. This completes the calibration process of our model

157 parameters. As Figure 2 shows, the model agrees well with *in vivo* experiment.

There are no bulk data, to our knowledge, measuring EN firing rates in response to odors and/or octopamine. However, calibrating EN response is not necessary to demonstrate an ability to learn. The key marker is post-training increase in EN response.

161 2.2 Learning Experiments: PN and KC Behaviors

PN activity in the AL, and KC activity in the MB, from typical Network Model simulations are shown in Figure 1 as heatmaps, evolved over the time course of a simulation in which the system was exposed to two different odors and trained on one of them. The AL is stimulated with octopamine during training. Each row of the heatmap represents a distinct PN or KC as it evolves in time (left to right columns of heat map). All the timescales are aligned. Neural behaviors are as follows:

167 2.2.1 PNs

In the AL heatmap, the light blue region corresponds to PN FRs within $2.5\sigma_s$ of their respective mean 168 169 spontaneous FRs μ_s , warm colors correspond to very high FRs, and dark blues correspond to strongly inhibited FRs. The simulations demonstrate a number of key PN behaviors, including (i) absent odor, PN 170 FRs stay within their noise envelopes (by definition), (ii) the two odors have distinct excitation/inhibition 171 signatures on PNs, (iii) octopamine alone (without odor) results in more PNs being excited beyond their 172 usual noise envelopes, and also results in some PNs being inhibited below their usual envelopes, (iv) 173 octopamine and odor, applied together, result in an overall excitation of PNs, and (v) the AL behavior 174 returns to baseline after octopamine is withdrawn, since AL connection weights do not have (long-term) 175 plasticity Davis (2005). 176

177 2.2.2 KCs

In the MB, the KCs fire sparsely due to global inhibition from the Lateral Horn. The only plastic connections in the AL-MB system involve the KCs: Between PNs and KCs (M^{PK} , M^{QK} in Methods Section); and between KCs and extrinsic readout neurons (ENs) (M^{KE} in Methods Section). Thus the KC odor signatures are modulated with training. Each row in the MB heatmap represents one of 500 randomly selected KCs in the simulation as it evolves in time (left to right columns of heat map). Black regions indicate FRs < 1 spike/sec, white regions indicate FRs > 1 spike/sec. The white regions have been dilated to make the sparsely-firing KCs easier to see.

The simulations demonstrate a number of key KC behaviors, including (i) the baseline KC FR response 185 absent any odor is essentially zero, (ii) the two odors excite distinct sets of KCs with varying consistency 186 187 from noise trial to noise trial, (iii) for a given odor, some KCs fire reliably in response to an odor stimulation 188 and some fire only occasionally, (iv) when subject to octopamine but no odor, KCs are unresponsive, a benefit during learning since it prevents environmental noise from being encoded as meaningful, (v) when 189 190 subject to both octopamine and odor, KCs respond strongly to the odor with high trial-to-trial consistency, 191 and (vi) the global inhibition from the LH controls the level of sparseness in the KCs, both their silence absent any odor (with or without octopamine), and their sparse firing in response to odors. 192

Statistics of KC responses to odors pre-, during, and post-training are shown in Figure 3. Naive moths 193 have low KC response to odors, in both percentage of KCs activated and their consistency of response to 194 odor stimulations (Figure 3, blue dots and curves). During training octopamine induces high KC response, 195 in both percentage and consistency (Figure 3, red dots and curves). After octopamine is withdrawn, 196 KC response is lower than during training, but remains higher than naive levels in both percentage and 197 consistency (Figure 3, green dots and curves) for the trained odor only. Thus the newly-learned importance 198 of the trained odor is encoded as broader and stronger KC responses by means of strengthened synaptic 199 connections. 200

EN (readout neuron) activity is also shown over time at the bottom of Figure 1. Learning is evidenced by the increased EN response to the trained odor even after octopamine has been withdrawn, due to Hebbian growth of synaptic connections into and out of the MB.

The FR activity of the PNs in the AL, the KCs in the MB, and the ENs, as illustrated in Figs. 1 and 3, demonstrate the entire learning process that occurs under the influence of octopamine stimulation. Without octopamine, learning does not occur.

Interestingly, although the AL does not itself experience plasticity changes, it is the AL's increased FR
 activity (induced by octopamine) which enables permanent synaptic weight changes in the MB via Hebbian
 plastic updates.

210 2.3 Learning Experiments: EN Behavior

A key finding of this paper is that the AL-MB model demonstrates robust learning behavior. Here "learning" is defined as permanently modifying synaptic weights in the system so that the reinforced odor yields a significantly stronger response in the readout neuron (EN) post-training, relative to naive (ie pre-training) response to that odor, and also relative to the post-training responses to control odors.

215 2.3.1 Structure of Learning Experiments

Moths were randomly generated from a fixed parameter template, which included randomly-assigned input maps (odor \rightarrow AL) for four odors. The odors projected broadly onto the AL, each odor hitting ~20 out of 60 glomeruli. As a result, their projections onto the AL overlapped substantially. A combinatorial calculation (4 independent draws of "60 choose 20") shows that, on average, a given odor projected uniquely onto about 6 glomeruli, and shared its other 14 glomeruli with other odors. Each generated moth was put through a series of training experiments, with each run in the series structured as follows:

- The moth first received a series of stimulations from each odor, to establish a baseline (naive) EN response. The stimulations were 0.2 seconds long and separated by gaps of several seconds.
- 224 2. The moth was trained on one of the odors for 1 to 4 sessions (one session = 5 odor stimulations), by
 225 applying odor and octopamine concurrently. The MB plastic weights were updated according to a
 226 Hebbian rule.
- 227 3. Upon completion of training, the four odors were each again applied as a series of odor stimulations,
 228 to establish post-training EN response.

For each {odor, #sessions} pair, this experiment was conducted 11 times (ie 11 noise realizations), for a total of 176 experiments on each moth. These results were aggregated to assess the particular moth's learning response.

232 2.3.2 Learning Experiment Results

As a general rule, Network Model moths consistently demonstrated strong learning behavior in terms of EN response: Training increased EN response to the trained odor well beyond naive levels, and also much more than it affected EN response to control odors. Figure 4 summarizes the changes in EN responses in a typical experiment on a moth with four odors. Panel A shows a typical noise realization timecourse, where one odor was reinforced with octopamine and the other three odors were controls. Panel B shows the statistics of EN response modulation, according to {odor, #sessions} pairs.

For ease of interpretation, the moth shown in Figure 4 had naive EN responses of roughly equal magnitude for all four odors. When naive EN response magnitudes were highly uneven (> 3x), robust learning still occurred, but the interpretation of the results is more complex due to scaling issues. A typical experiment using a moth with odor responses of highly unequal magnitude is shown in SI.

- 243 2.3.3 Points of interest (EN responses to learning)
- Because EN response is driven solely by feed-forward signals from KCs, ENs had response ≈ 0 in the absence of odor, with or without octopamine, as expected (since KCs are silent absent any odor). Thus Hebbian growth during training did not increase EN baseline (no-odor) response.
- 247 2. The EN response to odor + octopamine was always very strong, as seen in Figure 4A, where EN
 248 responses to odor + octopamine extend above the top of the figure. Note that this effect follows
 249 automatically from the calibration of the Network Model to *in vivo* data. Its functional value to the
 250 moth is addressed in the Discussion.
- 3. Training consistently increased the EN response to the reinforced odor much more than EN response
 to control odors, measured as percentage of naive odor response.

Since the Network Model did not include a Hebbian decay dynamic (for simplicity, absent clear evidence), this was the key indicator of robust learning. That is, focused learning was expressed by substantially higher increase in EN response to reinforced vs control odors. We assume that an added Hebbian decay term would have knocked smaller increases back, thus returning control odor responses to baseline.

Results of ANOVA analysis for differential effects of training on reinforced vs unreinforced odors shows that when naive odor EN response magnitudes were similar (within 3x of each other) p-values were consistently < 0.01. ANOVA analysis results are given in SI.

261 2.4 MB Sparsity Experiments

Projection into a high-dimensional, sparse layer is a common motif in biological neural systems Ganguli and Sompolinsky (2012); Litwin-Kumar et al. (2017). To explore the role of MB sparsity during learning, we ran Network Model experiments that varied the level of generalized inhibition imposed on the MB (the lateral horn, LH, controls MB sparsity level). Each experiment set a certain level of LH inhibition, then ran simulations (see Methods) that trained moths on one odor with 15 odor stimulations and left one control odor untrained. EN responses to both trained and control odors were recorded, as well as the percentage of KCs active in response to odor.

Too little damping from the LH resulted in a high percentage of KCs being active (low sparsity). This regime gave consistent EN responses to odor. But it also caused EN responses to both control odor and noise to increase significantly during training, reducing the contrast between EN responses to trained and control odors and also increasing spontaneous EN noise. Too much damping resulted in a very low percentage of KCs being active (high sparsity). This ensured that training gains were focused on the trained odor while EN response to control odors and noise were not boosted. However, in this regime EN responses to all odors, both pre- and post-training, were generally unreliable because too few KCs were activated.

Thus sparseness in the high-dimensional MB fulfilled a vital role in the Network Model's learning system. LH inhibition of the MB had an optimal sparsity regime that balanced opposing demands: KC firing had to be sufficiently dense for reliable odor response on one hand, and sufficiently sparse for well-targeted Hebbian growth on the other. Timecourses illustrating the effects of too-little or too-much sparsity are seen in Figure 5A. Figure 5B shows how this trade-off varied with MB sparsity, by plotting two figures-of-merit:

Signal-to-Noise Ratio (SNR)
$$= \frac{\mu(f)}{\sigma(f)}$$
 where $f = \text{EN odor response}$; (1)

282 and

Learning Focus"
$$= \frac{\mu(f_T)}{\mu(f_C)}$$
, where $\mu(f_T) =$ mean post-training EN (2)

283 response to trained odor, $\mu(f_C)$ = mean post-training EN response to control odor.

3 DISCUSSION

"

Because we took a distinct approach to designing our Network Model, and because we had access to unique *in vivo* octopamine data, our experiments yield novel insights into the moth olfactory network and how it learns. This discussion focuses on four areas: (i) predictions about aspects of the AL-MB still unclear in the literature, (ii) the role of sparse layers, (iii) the role of octopamine, and (iv) the value of noise. In addition, we consider these insights in the context of Machine Learning.

290 3.1 Predictions re details of AL-MB structure

Because our Network Model in tethered to a particular system, both the calibration process and simulations offer hints as to some unresolved biophysical aspects of the moth's AL-MB system. Some examples:

293 3.1.1 Do LNs inhibit PNs and LNs as well as RNs

In the AL, LNs have a net inhibitory effect on PNs Olsen et al. (2010); Lei et al. (2002), but the exact 294 means to this end are not clear. In particular, while LNs are known to inhibit RNs Olsen et al. (2010), it 295 is less clear whether or to what degree LNs also directly inhibit PNs and LNs. Efforts to calibrate our 296 Network Model to *in vivo* data indicate that LNs need to inhibit not just RNs, but also (to a lesser degree) 297 LNs and PNs. The model weight strengths for LN \rightarrow RN, \rightarrow LN, and \rightarrow PN are in the ratio of 6:2:1. That 298 LNs would inhibit LNs makes sense when the goal is maximizing PN output of the active glomerulus: By 299 inhibiting the LNs of rival glomeruli, the active glomerulus reduces the amount of inhibition directed at 300 itself. Similarly, that LNs would inhibit PNs makes sense if the goal is to reduce the PN output of rival 301 glomeruli. 302

303 3.1.2 Octopamine's effects on different neuron types

Octopamine increases the responsivity of a neuron to incoming signals. It is unclear how or whether octopamine affects various neuron types (ie RNs, PNs, LNs, KCs). Calibration of the Network Model's AL behavior, and tuning of KC behavior to enable learning, indicate that octopamine needs to target RNs andLNs, but not PNs, KCs, or ENs. Logical arguments support these findings:

RNs: Because RNs initially receive the odor signal, these are logical neurons to stimulate with octopamine,
because it sharpens their response to the exact signature being trained, which in turn sharpen the AL's
output code for that odor.

311 LNs: LNs have the dual roles of inhibiting rival glomeruli and limiting overall PN output in the AL.
312 For the first role, increased LN response to RNs will tend to sharpen AL response to the trained odor, by
313 accentuating inhibition of rival glomeruli PNs. For the second role, increased LN activity mitigates the risk
314 that increased RN activity (due to octopamine) might blow up the overall PN output of the AL.

PNs: Our Network Model simulations suggest that PNs should receive little or no octopamine stimulation. While increasing PN responsivity would benefit RN-induced sharpening of the trained odor's signature, there are three downsides. First, RN input to PNs is intrinsically noisy, so higher PN responsivity amplifies noise as well as signal. Second, since PNs respond to LNs, higher PN activity tends to reduce the impact of LN inhibition, and thus reduces the inhibition-induced sharpening of the AL odor response caused by octopamine. Third, increasing PN responsivity can have an outsize effect on overall PN firing rates, ie it is a "high-gain" knob and therefore risky.

322 **KCs:** Our Network Model simulations indicate that direct octopamine stimulation of KCs greatly reduces sparseness in the MB (given the mechanics of our global inhibition on KCs), which can be disastrous to 323 324 learning. Thus we expect that octopamine has no, or only slight, direct stimulative effect on KCs. However, 325 other forms of global inhibition might admit direct octopamine stimulation of KCs while still preserving sparsity: (i) If the strength of the Lateral Horn's inhibition signal tracks the AL output, then octopamine 326 stimulation of KCs would be offset by increased inhibition from the LH, due to increased AL output to the 327 LH, preserving sparsity; (ii) if the KC population generates the inhibition signal, as in Lin et al. (2014), 328 then octopamine stimulation of KCs would result in a counteractive stronger inhibition, again preserving 329 sparsity. The arguments for why PNs should receive very little direct octopamine stimulation (given above) 330 apply to KCs as well; but it is also possible that direct stimulation of KCs might improve learning by 331 332 enabling random exploration of the odor-coding solution space (as mooted below for octopamine).

333 3.2 Noise filtering role of the sparse, high-dimensional stage

Projection from a dense, low-dimensional coding space (eg the AL) to a sparse, high-dimensional coding space (eg KCs in the MB) is a widespread motif of biological neural systems, with size shifts routinely on the order of 20x to 100x Ganguli and Sompolinsky (2012); Babadi and Sompolinsky (2014); Litwin-Kumar et al. (2017). Some proposed reasons include information capacity, long-range brain communication, and reduced training data needs Ganguli and Sompolinsky (2012), as well as better inherent discrimination ability Peng and Chittka (2017); Litwin-Kumar et al. (2017); Bazhenov et al. (2013).

Our Network Model experiments highlight another key role of sparseness, relevant to learning: It acts as a robust noise filter that prevents the Hebbian growth process from amplifying upstream noise to out-of-control levels. Though noise may be useful (or unavoidable) in upstream networks such as the AL, noise that reaches the neurons on both sides of a plastic synaptic connection will be amplified by Hebbian growth during learning, swamping the system's downstream neurons (eg ENs) with noise.

However, the "fire together, wire together" principle of Hebbian learning is an AND gate. Thus it suffices to remove noise from just one of the two connected neurons to eliminate synaptic growth. Sparsity does precisely this, and is arguably a necessary part of a workable Hebbian learning mechanism. We also find that high sparsity focuses learning (even absent upstream noise), ie it enables better learned separation
of classes, in agreement with Peng and Chittka (2017); Huerta and Nowotny (2009). The negative effect
of high sparsity on SNR (Fig 5) found in our experiments meshes with a similar observation in Nowotny
(2009).

Setting aside the particular demands of Hebbian plasticity, robust noise filtering may be a core function of sparse, high-dimensional stages within any network cascade where noise accumulates due to (beneficial) use in upstream stages.

355 3.3 Roles of octopamine

The levels of octopamine stimulation in our Network Model were calibrated to *in vivo* data on PN responses to octopamine. Thus, our simulations give novel insights into downstream effects of octopamine on plasticity, KC responses, EN responses, and Hebbian learning.

359 3.3.1 Accelerant

Moths can learn to respond to new odors remarkably quickly, in just a few exposures. Our simulations indicates that while Hebbian growth can occur without octopamine, it is so slow that actionable learning, ie in terms of amplified EN responses, does not occur.

This implies that octopamine, through its stimulative effect, acts as a powerful accelerant to learning. Perhaps it is a mechanism that allows the moth to work around intrinsic organic constraints on Hebbian growth of new synapses, constraints which would otherwise restrict the moth to an unacceptably slow learning rate. To the degree that octopamine enabled a moth to learn more quickly, with fewer training samples, it would clearly be highly adaptive.

368 3.3.2 Active learning

Our simulations indicate that octopamine strongly stimulates the EN response to even an unfamiliar odor. Since octopamine is delivered as a reward, this has a beneficial effect in the context of reinforcement learning Sutton and Barto (1998), with the moth as the learning agent. An agent (the moth) can in some cases learn more quickly when it has choice as to the sequence of training samples (Active Learning Settles (2012)).

In particular, when a certain class of training sample is relatively rare, it benefits the agent to actively seek out more samples of that class Attenberg and Provost (2010). Octopamine enforces high EN response to a reinforced odor, ensuring that ENs will consistently exceed their "take action" threshold during training. If the action is to "approach", the moth is more likely to again encounter the odor, thus reaping the benefits predicted by Active Learning theory. This advantage applies in the context of positively-reinforced odors.

In the case of aversive learning, the high EN responses to unfamiliar but objectionable odors, due to dopamine, would cause the moth to preferentially avoid further examples of the odor. This would slow learning of aversive responses (a drawback), but would also minimize the moth's exposure to bad odors (danger avoidance, a benefit).

383 3.3.3 Exploration of optimization space

A limitation of Hebbian growth is that it can only reinforce what already exists. That is, it only strengthens channels that are transmitting signals deemed (by association) relevant to the stimulus being reinforced. Absent a mechanism like octopamine, this constrains growth to channels that are already active. Our simulations indicate that octopamine induces much broader activity, both upstream from and within the 388 plastic layer, thus activating new transmitting channels. This allows the system to strengthen, and bring 389 permanently online, synaptic connections that were formerly silent. This expands the solution space the 390 system can explore during learning. This function may be particularly important given the constraint of 391 sparsity placed on odor codes in the MB.

392 3.3.4 Injury compensation

There is evidence that many forms of injury to neurons result in dropped spikes and thus lower firing rates in response to odors Maia and Kutz (2017). This injury-induced drop in the signals reaching the ENs could induce behavioral consequences, by lowering EN responses to below key behavioral action thresholds. Experiments in Delahunt et al. (2018) suggest that octopamine drives a mechanism to compensate for this type of neural injury.

Suppose that injury has reduced an upstream neural signal, such that a downstream EN can no longer exceed its behavioral action threshold. Octopamine stimulation of the upstream network will temporarily boost the reduced (injured) signal strength, so that the input signals to the ENs are above threshold during training. This in turn allows Hebbian growth to strengthen the synaptic connections to those ENs. Once octopamine is withdrawn, the inputs from the (still-injured) upstream network revert to their reduced level. But due to the newly-strengthened connection weights, these reduced inputs suffice to push EN response above its action threshold, restoring EN-controlled behaviors to their pre-injury baseline.

This mechanism (octopamine stimulation plus Hebbian synaptic growth) might allow neural systems to regain behavioral function lost due to damage to upstream regions, by increasing connection strengths downstream from the point of injury. Indeed, given the vital importance of injury mitigation to survival, it is possible that the "learning" mechanism originally evolved for the purpose of restoring behavioral function impaired by neural damage.

410 3.4 The value of noise

Noise in biological neural systems is believed to add value (besides just being cheap), for example by
encoding probability distributions and enabling neural Baysian estimations of posteriors Ma et al. (2006).
In addition, injection of noise while training engineered NNs can improve trained accuracy An (1996). Our
experiments indicate that in the AL-MB, noise has two other potential benefits, coupled with a caveat.

First, noise in the AL adds an extra dimension to MB odor encoding, increasing the granularity of its odor responses (Figure 3). The MB responds to odors in two ways: (i) by the number of KCs that are responsive, and (ii) by the reliability (eg from 10% to 100%) of their responses. This can be seen in the effect of octopamine on KC odor response, Figure 3 (B). Octopamine boosts MB odor response by increasing the number of active KCs (horizontal shift in response curves) and also by increasing the reliability of responsive KCs (vertical shift in responsivity curves). Both these shifts represent a stronger MB response and translate into stronger EN response.

Taken together, they provide a finer granularity of the response range than does the binary response of a noise-free system. That is, the MB response to noisy inputs from the AL is a concrete example of a mechanism used by a neural system to translate the probability distributions encoded by noisy neurons into actionable signals with high dynamic range and granularity.

Second, the system is also potentially robust to noisy stimuli. In the neural net context, input samples (ie inputs to the feature-reading layer) can be thought of as a *de facto* "first layer" of the neural net. A system that is robust to upstream noise may also be naturally robust to noisy inputs, a further potential advantageof judicially-placed sparse layers.

The caveat is that noise in the AL-MB must be confined to the AL, ie upstream from the encoding layer, in order to protect the readout neurons and Hebbian learning mechanism from noise. The system's success depends on robust noise filtering at the MB layer, via global inhibition from the LH. So the three-stage architecture consisting of: "Noisy pre-amplifier layer \rightarrow Sparse noise-reduction layer \rightarrow Decision layer" is an interdependent system well-suited to nuanced decision-making.

435 3.5 Applications to Machine Learning

The model and simulations in this paper characterize key features of the AL-MB system which might usefully be ported to machine learning algorithms. These features include: Generalized stimulation during training; Hebbian growth; sparse layers to control plastic connections and filter noise; and noisy initial layers. Advantages of this biological toolkit include:

440 3.5.1 Fast learning

Moths can reliably learn a new odor in less than 10 exposures, and biological brains in general can learn given very few training samples. This contrasts by orders of magnitude with the voracious data demands of DNNs, for which assembling sufficient training data can be a serious chokepoint in deployment. Indeed, when learning handwritten digits from very few samples (1 to 10 per class), an insect brain outperforms ML methods including CNNs Delahunt and Kutz (2018). These "fast and rough" biological mechanisms, seen in the moth in their simplest form, thus have potential to act as a complement to the precise but slow learning of DNNs.

448 3.5.2 Robustness to noise

The sparse layer in the AL-MB acts as an effective noise filter, protecting the readout neurons from a noisy upstream layer (the AL). Since the system is designed to accommodate upstream noise, it is possible that it can also readily accommodate noisy input samples. NNs have a troublesome property, that input-output score functions are not locally continuous Szegedy et al. (2013). Biological neural nets seem to avoid this particular fault (or at least have different, complementary discontinuities). The noisy layer \rightarrow sparse layer motif may be one reason for this. It may thus be a useful motif to apply in ML architectures.

455 3.5.3 Novel training mechanism

Hebbian growth, combined with octopamine stimulation and the focusing effect of sparse layers, is a 456 novel (in the context of ML) mechanism to explore a solution space and train a classifier. In particular, it 457 works on a different principle than the backprop algorithm that drives DNNs: It does not minimize a loss 458 function via gradient descent, nor does it punish incorrect answers; rather, it selectively strengthens those 459 connections that transmit meaningful signals, and weakens connections that are inactive. We argue that this 460 biological optimization mechanism is of potential value to ML because it is (i) functionally distinct from 461 the backprop algorithm currently used, and (ii) known to succeed in regimes (eg rapid learning) where 462 backprop struggles. 463

464 3.5.4 Biological plausibility

One characteristic (not criticism) of backprop is its biological implausibility, since it requires a neuron to
have more than local knowledge of the system. A current area of interest, especially in the context of DNNs
Bengio and Fischer (2015), is the search for neural network architectures (for example with recurrent

468 connections to transport non-local information) and variants of backprop which are biologically plausible,
469 which might narrow the gap between biological and engineered NNs. Our experiments demonstrate that
470 the triad of octopamine stimulation + Hebbian growth + sparse layers can efficiently train a NN, and is thus

471 a possible candidate to address the biological plausibility gap.

4 MATERIALS AND METHODS

This section gives a detailed description of the biological moth olfactory network, as well as our Network Model. The biological detail is relevant because our model is built "from the ground up", ie it is based on a particular organism, rather than being a "top-down", more theoretically-based, architecture. We also provide a Glossary, and describe the *in vivo* data used for model calibration.

476 4.1 Moth olfactory system overview

The parts of the AL-MB implicated in learning are organized as a feed-forward cascade of five distinct networks, as well as a reward mechanism Martin et al. (2011); Kvello et al. (2009). Figure 1 gives a system schematic along with typical firing rate (FR) timecourses (from simulation) for neurons in each network.

- 480 1. Antennae. Roughly 30,000 noisy chemical receptors detect odor and send signals to the Antenna Lobe481 Masse et al. (2009).
- 482 2. Antenna Lobe (AL). Contains roughly 60 units (glomeruli), each focused on a single odor feature
 483 Martin et al. (2011). The AL essentially acts as a pre-amp, boosting faint signals and denoising the
 484 antennae inputs Bhandawat et al. (2007). AL neurons are noisy Galizia (2014).
- 485 3. Lateral Horn (LH). Though not fully understood, one key function is global inhibition of the Mushroom
 486 Body to enforce sparseness Bazhenov and Stopfer (2010).
- 487 4. Mushroom Body (MB), here synonymous with the Kenyon Cells (KCs). About 4000 KCs are located in
 488 the calyx of the Mushroom Body (MB). These fire sparsely and are believed to encode odor signatures
 489 Perisse et al. (2013); Campbell and Turner (2010); Honegger et al. (2011). KCs are believed to be
 490 relatively noise-free Perez-Orive et al. (2002).
- 491 5. Extrinsic Neurons (ENs), numbering ~10's, located downstream from the KCs. These are believed
 492 to be "readout neurons" that interpret the KC codes and convey actionable messages (such as "fly
 493 upwind") Campbell et al. (2013); Hige et al. (2015).
- 6. Reward Mechanism. A large neuron sprays octopamine globally over the AL and MB, in response to reward, such as sugar at the proboscis. Learning does not occur without this octopamine input Hammer and Menzel (1995, 1998). The neuromodulator dopamine works similarly, but drives aversive learning Dacks et al. (2012).
- 498 7. Inter-network connections: In the AL-MB these are strictly feed-forward, either excitatory or inhibitory. In particular, Antennae→AL, AL→LH, KCs→ENs are all excitatory. LH→KCs is inhibitory.
 500 AL→KCs have both excitatory and inhibitory channels.
- 8. Plasticity: The connections into the KCs (AL→KCs) and out of the KCs (KCs→ENs) are known to
 be plastic during learning Cassenaer and Laurent (2007); Masse et al. (2009). The AL does not have
 long-term plasticity Davis (2005).

504 4.2 Glossary

- Antenna lobe (AL): A collection of neurons innervated by odor receptors in the antennae. It sends signals
 to the mushroom body via projection neurons. Connections in the AL are not plastic.
- 507 Mushroom body (MB): A collection of neurons (Kenyon cells KCs) downstream from the antenna lobe.
- 508 The MB is believed to store odor codes that serve as a memory, allowing the moth to recognize odors. 509 Connections in the MB are plastic.
- 510 **Lateral horn** (**LH**): A collection of neurons which receives input from the AL and sends inhibitory output 511 to the MB. One of its roles is to enforce sparse firing in MB neurons.
- 512 **Receptor neuron (RN)**: These neurons respond to odors (volatiles) at the antennae and stimulate the 513 antenna lobe. RNs respond to different, distinct odors.
- 514 **Glomerulus**: The antenna lobe is divided into about 60 glomeruli, each of which is a self-contained 515 collection of neurons (projection and lateral), innervated by RNs that respond to particular odors.
- 516 Projection neuron (PN): Each glomerulus contains projection neurons, whose output innervates the KCs
- 517 and also the lateral horn, but not other glomeruli in the AL, ie they are feed-forward only. Most PNs start
- 518 in one glomerulus and are excitatory. A few PNs arborize in several glomeruli and are inhibitory (we refer519 to inhibitory PNs as "QNs"). Each glomerulus initiates about five PNs.
- 520 **Lateral neuron** (**LN**): Each glomerulus contains lateral neurons, which innervate other glomeruli in the 521 AL. LNs are inhibitory. One function is competitive inhibition among glomeruli. Another function is gain 522 control, ie boosting low signals and damping high signals.
- **Kenyon cell (KC)**: Neurons in the calyx of the MB. These have very low FRs, and tend to respond to particular combinations of PNs. KCs respond sparsely to a given odor. There are about 4000 KCs, ie a two-orders-of-magnitude increase over the number of glomeruli. Each KC synapses with about ten PNs.
- 526 Connections into and out of KCs are plastic.
- 527 **Extrinsic neuron (EN)**: A small number of neurons downstream from the KCs. ENs are thought to be 528 "readout" neurons. They interpret the odor codes of the KCs, deciding to eg "ignore", "approach", or 529 "avoid".
- **Firing rate (FR)**: The number of spikes/second at which a neuron fires. Typically FRs are counted using a window (eg 500 ms). The moth's response to odor stimulations is episodic, with FR spikes in FR and rapid return to spontaneous FRs. Neurons respond to relative changes in FR, rather than to raw magnitude changes. A neuron's relative change in FR is scaled by its spontaneous FR (see section 4.5 below).
- **Octopamine**: A neuromodulator which stimulates neural firing. The moth spritzes octopamine on both the AL and MB in response to sugar, as a feedback reward mechanism. Dopamine has a similar stimulating effect on both AL and MB, but it reinforces adverse rather than positive events.

537 4.3 Component networks and their Network Model representations

This subsection offers a more detailed discussion of the constituent networks in the biological AL-MB, and details about how they are modeled in our Network Model.

540 4.3.1 Antennae and receptor neurons

The Antennae receptors, activated by chemical molecules in the air, send excitatory signals to Receptor Neurons (RNs) in the AL. Several thousand antennae converge onto 60 units (glomeruli) in the AL Nagel and Wilson (2011). All the receptors for a given atomic volatile converge onto the same glomerulus in the AL, so the glomeruli each have distinct odor response profiles Deisig et al. (2006). Since natural odors are a blend of atomic volatiles, a natural odor stimulates several units within the AL Riffell et al. (2009b). 546 Our model does not explicitly include antennae. Rather, the first layer of the model consists of the RNs 547 entering the glomeruli. Though \sim 500 RNs feed a given glomerulus, the model assumes one RN. The 548 benefit of many RNs converging appears to be noise reduction through averaging Olsen et al. (2010). This 549 can be simulated by one RN with a smaller noise envelope.

Each glomerulus' RN has a spontaneous FR and is excited, according to random weights, by odor stimuli.

551 4.3.2 Antenna lobe and projection neurons

The AL is fairly well characterized in both structure and dynamics, with a few important gaps. It contains 552 about 60 glomeruli, each a distinct unit which receives RN input and projects to the KCs via excitatory 553 PNs. The same PN signal also projects to the LH Bazhenov and Stopfer (2010). The AL, unique among 554 the networks, has inhibitory lateral neurons (LNs) Wilson and Laurent (2005), the only neurons that are 555 not strictly feed-forward. (There is some evidence of excitatory LNs, eg Olsen et al. (2008); the Network 556 Model excludes this possibility.) The LNs act as a gain control on the AL, and also allow odors to mask 557 each other by inhibiting other glomeruli's RNs Olsen and Wilson (2008); Hong and Wilson (2015). It is 558 not known whether LNs also inhibit PNs and LNs. Based on calibrations to in vivo data, in Network Model 559 LNs inhibit all neuron types (cf section 3.1). Thus each glomerulus contains dendrites (ie outputs) for PNs 560 and LNs, and axons (ie inputs) for RNs and LNs, as shown in Figure 6. 561

Each glomerulus does the following: Receives RN input from the antennae receptors upstream; inhibits
other glomeruli within the AL via LNs; and sends excitatory signals downstream via Projection Neurons
(PNs).

In general, each PN is innervated in a single glomerulus. In moths, there are \sim 5 PNs rooted in each glomerulus (60 glomeruli, \sim 300 PNs). The Network Model assumes all PNs from a given glomerulus carry the same signal (because they share the same glomerulus and therefore inputs, and perhaps also because of ephaptic binding) Sjoholm (2006).

569 Glomeruli also initiate pooled Inhibitory Projection Neurons (QNs) that send inhibitory signals 570 downstream to the KCs.

The AL contains a powerful macro-glomerulal complex (MGC), which processes pheromone. Because pheromone response has fundamentally different dynamics than food odor response Jefferis et al. (2007), the model ignores it. Only the glomeruli associated with non-pheromone (food) odors are modeled.

574 Connections in the AL are not plastic with long-term persistence Dacks et al. (2012). While some 575 evidence of short-term plasticity exists, the Network Model ignores this option.

576 4.3.3 Lateral horn

577 The LH receives input from the PNs. It then sends an inhibitory signal to the KCs. This inhibition from 578 the LH appears to ensure that the KCs fire very sparsely and thus act as coincidence detectors for signals 579 from the AL Sjoholm (2006); Lin et al. (2014); Gruntman and Turner (2013).

580 The LH is also suspected of containing a parallel system for processing certain intrinsically-known 581 odors in short-cut fashion (labeled lines) Luo et al. (2010). Since this parallel system is (by definition) not 582 involved with learning, the Network Model ignores it. The LH is modeled solely as a simple sparsifying 583 inhibition on the KCs.

(Note: The locust and honeybee, which have more complex olfactory systems and different use-cases in
 terms of odor processing, have a time-oscillating depolarization mechanism (local potential fields, LPF)

Perez-Orive et al. (2002) which serves a similar purpose to LH inhibition in the moth. LPF oscillations areabsent in the moth Martin et al. (2011).)

588 4.3.4 Mushroom body and Kenyon cells

The KCs (\sim 4000) in the MB are believed to encode odor memories in a high-dimensional, sparse space Turner et al. (2008). Odors with no meaning to the moth still have non-zero codes in the KCs.

KCs receive excitatory input from the PNs and inhibitory input from QNs, both of which vary greatly 591 592 between KCs, since each KC is innervated by only ~ 10 PNs Martin et al. (2011). The connection map appears to be random Caron et al. (2013). The KCs also receive generalized damping inhibition from the 593 LH. (There is some evidence in drosophila of an MB \rightarrow MB global inhibitory neuron Lin et al. (2014), with 594 595 the same essential effect as LH inhibition; the Network Model excludes this possibility.) KCs fire very sparsely, generally respond to only a single odor, and are silent absent that odor Honegger et al. (2011). 596 KCs are treated as noise-free. Their output is an excitatory signal sent to the extrinsic neurons (ENs) 597 Campbell et al. (2013). 598

In addition to olfactory input, the KCs receive input signals from other parts of the moth (eg hearing) Sjoholm (2006). Because the Network Model targets olfactory learning, it ignores these other inputs and uses a reduced number of KCs (\sim 2000 instead of \sim 4000).

The synaptic connections in the MB (PNs \rightarrow KCs, QNs \rightarrow KCs, and KCs \rightarrow ENs) are plastic, ie they can be modified during training Menzel and Manz (2005). The generalized inhibition from LH \rightarrow KCs is modeled as non-plastic (actual physiology is not known). This LH inhibition is modeled as a global damping term on KCs, giving dynamics equation equivalent to the Pitt- McColloch approximation McCulloch and Pitts (1943) as used in Bazhenov et al. (2013); Peng and Chittka (2017); Mosqueiro and Huerta (2014); Huerta and Nowotny (2009).

608 4.3.5 Extrinsic neurons

Though located in the lobes of the MB, here ENs are not considered part of the MB, which is taken to be synonymous with the KCs. ENs are few in number compared to the KCs (\sim 10s) Campbell et al. (2013); Hige et al. (2015). They are believed to be "readout" neurons, that interpret the KC codes as actionable signals (eg "approach", "avoid") Masse et al. (2009). We assume that ENs trigger actions when their output FRs exceed some threshold.

We define Learning as: Permanently boosting EN responses beyond their naive (untrained) level, so that EN responses to reinforced stimuli can consistently exceed an action-triggering threshold. This is tantamount to modifying the moth's behavior.

617 4.3.6 Octopamine (reward circuit)

A large neuron delivers octopamine to the entire AL and MB, in response to positive stimuli, eg sugar at the proboscis. It acts as a reward feedback to the system. A similar neuron delivers dopamine to the AL and MB in response to negative stimuli, and acts as an aversive feedback signal Dacks et al. (2012). Learning does not occur without octopamine (or dopamine) Hammer and Menzel (1998).

Despite their opposite reward values, both octopamine and dopamine act in the same way when sprayed on a neuron: They increase the neuron's general tendency to fire Riffell et al. (2012). In the Network Model this effect is modeled as making a neuron more responsive to excitatory inputs (eg from odors and RNs) and less responsive to inhibitory inputs (eg from LNs). Details of octopamine's effects if any on particular neural types are not well-characterized. In the Network Model octopamine directly affects RNs and LNs but not PNs in the AL (cf section 3.1); has no direct effect on KCs or ENs (though there are strong indirecteffects); and has no effect on the LH inhibitory signal.

It is unclear whether octopamine delivery to both the MB and AL is necessary and sufficient for learning Hammer and Menzel (1998); Dacks et al. (2012). The Network Model assumes that octopamine controls an "on/off" switch for Hebbian growth, ie there is no plasticity in the MB (and therefore no learning) without octopamine.

633 4.4 Network Model model description

This section describes our Network Model model in detail. It covers the firing rate measure used to compare model output to *in vivo* data; model dynamics; plasticity and other details; model parameters; and moth generation. All coding was done in Matlab. Computer code for the Network Model in this paper will be found at:

638 https://github.com/charlesDelahunt/SmartAsABug

639 4.4.1 Model dynamics

Our Network Model uses standard firing rate dynamics Dayan and Abbott (2005) (chapter 7), evolved 640 as stochastic differential equations Higham. (2001). We use a firing rate model for two reasons. First, 641 642 it is the simplest model able to both capture the key response and learning dynamics of the moth, and also allow calibration to our *in vivo* datasets. Our *in vivo* datasets are spike trains and thus admit use of 643 an integrate-and-fire or Izhikevich model; but at the cost of more parameters, higher complexity, and 644 longer simulation times, without (in our eyes) commensurate benefit. Second, and importantly, we wish to 645 apply our results to engineered NNs, which makes a firing rate model the natural choice due to its close 646 similarities to NNs. 647

648 The firing rate model is formulated as follows:

649 Let x(t) = firing rate (FR) for a neuron. Then

$$\tau \frac{dx}{dt} = -x + s(\Sigma \mathbf{w}_i \mathbf{u}_i) = -x + s(\mathbf{w} \cdot \mathbf{u}), \text{ where}$$
(3)

650 \mathbf{w} = connection weights;

651 **u** = upstream neuron FRs;

652 s() is a sigmoid function or similar.

653

PN dynamics are given here as an example. Full model dynamics are given in SI. PNs are excitatory, and project forward from $AL \rightarrow MB$:

$$\tau \frac{d\mathbf{P}}{dt} = -\mathbf{P} + s(\widetilde{\mathbf{P}}) + \mathbf{dW}^{P} \text{ where}$$
(4)

656 W(t) = brownian motion process;

657 $\widetilde{\mathbf{P}} = -(1 - \gamma o(t)M^{OP})*M^{LP}*\mathbf{u}^{L} + (1 + o(t)M^{OP})*M^{RP}*\mathbf{u}^{R};$

658 M^{OP} = octopamine \rightarrow PN weight matrix (diagonal $nG \times nG$);

659 $M^{LP} = LN \rightarrow PN$ weight matrix $(nG \times nG \text{ with } trM^{LP} = 0);$

660 $M^{RP} = \text{RN} \rightarrow \text{PN}$ weight matrix (diagonal $nG \times nG$);

661 o(t) indicates if octopamine is active (o(t) = 1 during training, 0 otherwise).

662 $\mathbf{u}^L = \text{LN FRs}$, vector $nG \times 1$;

663 $\mathbf{u}^R = \mathbf{RN} \, \mathbf{FRs} \, (nG \times 1);$

664 γ = scaling factor for effects on inhibition.

665 4.4.2 Discretization

666 The discretization uses Euler-Maruyama (E-M), a standard step-forward method for SDEs Higham. 667 (2001).

668 Euler (ie noise-free): $x_{n+1} = x_n + \Delta t f(x_n)$

669 Euler-Maruyama: $x_{n+1} = x_n + \Delta t f(x_n) + \epsilon \operatorname{randn}(0,1)\sqrt{\Delta t}$, where ϵ controls the noise intensity.

670 4.4.3 Convergence

Timestep Δt was chosen such that noise-free E-M evolution gives the same timecourses as Runge-Kutta (4th order), via Matlab's ode45 function. $\Delta t = 10$ mSec suffices to match E-M evolution to R-K in noise-free moths. Values of $\Delta t \leq 20$ mSec gives equivalent simulations in moths with AL noise calibrated to match *in vivo* data. Values of $\Delta t \geq 40$ mSec show differences in evolution outcomes given AL noise.

675 4.4.4 Plasticity

The model assumes a Hebbian mechanism for growth in synaptic connection weights Hebb (1949); Cassenaer and Laurent (2007). That is, the synaptic weight w_{ab} between two neurons a and b increases proportionally to the product of their firing rates ("fire together, wire together"): $\Delta w_{ab}(t) \propto f_a(t) f_b(t)$.

679 Thus, synaptic plasticity is defined by:

$$\Delta w_{ab}(t) = \gamma f_a(t) f_b(t), \text{ where } \gamma \text{ is a growth parameter.}$$
(5)

There are two layers of plastic synaptic weights, pre- and post-MB: AL \rightarrow MB ($M^{P,K}, M^{Q,K}$), and MB \rightarrow ENs ($M^{K,E}$). Learning rate parameters of the Network Model were calibrated to match experimental effects of octopamine on PN firing rates and known moth learning speed (eg 4 - 8 trials to induce behavior modification) Riffell et al. (2012). The Network Model does not decay unused synaptic weights. Training does not alter octopamine delivery strength matrices ($M^{O,*}$). That is, the neuromodulator channels are not plastic (unlike, for example, the case in Grant et al. (2017)).

686 4.4.5 Odor and octopamine injections

Odors and octopamine are modeled as hamming windows. The smooth leading and trailing edges ensures low stiffness of the dynamic ODEs, and allows a 10 mSec timestep to give accurate evolution of the SDEs in simulations.

690 4.4.6 Training

Training on an odor consists of simultaneously applying stimulations of the odor, injecting octopamine, and "switching on" Hebbian growth. Training with 5 to 10 odor stimulations typically produces behavior change in live moths.

694 4.5 Firing rate measure

To compare PN firing rate statistics from *in vivo* experiments and from Network Model simulations (ie model calibration), we use a measure of firing rate (FR) based on Mahalanobis distance, similar to the measure $\frac{DF}{F}$ common in the literature Campbell et al. (2013); Hong and Wilson (2015); Turner et al. (2008); Silbering and Galizia (2007). The premise is that neurons downstream respond to a ±1 std change

in FRs equally (modulo different connection weights), independent of the sometimes large (up to 40x) 699 magnitude differences in the raw spontaneous FRs of different neurons. The FR measure is defined as 700 follows: 701

- 1. Each PN has a spontaneous firing rate (FR) with a gaussian noise envelope. 702
- 2. PNs with FR < 1 spike/sec are ignored, on the assumption that such PNs represent artifacts of 703 experiment (also, the gaussian noise assumption fails). About 10% of PNs in experimental data fall in 704 705 this category.
- 3. Output FR activity of PNs is measured as M(t) = distance from mean spontaneous FR, in units of 706 time-varying std dev of spontaneous FR (ie Mahalanobis distance): Let 707
- F(t) = raw firing rate (spikes per second). 708
- S(t) = spontaneous firing rate (no odor). 709
- $\mu S(t)$ = moving average of S (no odor). 710
- $\bar{\mu}S(t)$ = smoothed estimate of the moving average μS , eg a quadratic or spline fit. 711
- $\sigma_S(t)$ = standard deviation of S, calculated using $S \bar{\mu}S$ values within a moving window centered on 712 713 t.
- $\sigma_S(t)$ and $\mu S(t)$ are typically steady absent octopamine, but are often strongly modulated by oc-714 715 topamine.
- Then the measure of FR activity M is: 716
- 717

$$M(t) = \frac{F(t) - \bar{\mu}S(t)}{\sigma_S(t)}$$
(6)

718

4. *M* is related to the measure $\frac{DF}{F}$: $\frac{DF}{F} = \frac{\Delta F}{F} = \frac{F(t) - \mu S}{\mu S}$, ie $\frac{DF}{F}$ is change in FR, normalized by spontaneous FR. The key difference between *M* and $\frac{DF}{F}$ is whether or how σ_S is estimated, due to varying exigencies of experiment. Our 719 720 experimental data allow reasonable estimates of σ_S and μS . Network Model simulations produce very 721 722 good estimates, since computer models are more amenable to repeated trials than live moths.

4.6 Model parameters 723

There is a risk, when modeling a system, of adding too many free parameters in an effort to fit the system. 724 725 Fewer free parameters are better, for the sake of generality and to avoid overfitting. Conversely, we wish 726 to reasonably match the physiological realities of the system. Because the key goal of this paper was to demonstrate that a simple model, in terms of parameters and structure, can reproduce the learning behavior 727 728 of the AL-MB, we made efforts to minimize the number of free parameters. For example, neuron-to-neuron 729 connections in the model are defined by their distributions, ie two parameters each. These are (usually) distinct for different source-to-target pairs (eg LN → RN, LN → LN, etc). 730

- 731 Some mean and std dev parameters for distributions are shared among different neuron types (eg LNs, PNs, 732 and QNs all share the same variance scaling parameter).
- 4.6.1 Parameter list 733
- 734 The model has in total 47 free parameters:
- 1. Structure: 5 (eg number of neurons in each network) 735
- 2. Dynamics: 12 (noise: 2. decay and sigmoid: 3. Hebbian growth: 6. misc: 1). 736
- 737 3. Spontaneous RN FRs: 3.

4. Connection matrices: 27 (controlling non-zero connection ratios: 5; synaptic weights (eg $M^{P,K}, M^{R,P}$): means 12, std devs 4; octopamine weights (eg $M^{O,R}, M^{O,P}$): means 6, std devs 2).

740 4.6.2 Dynamics parameters

The differential equations of all neuron types share the same decay rate, set to allow return to equilibrium in ~1 second, consistent with *in vivo* data. Neurons also share parameters of the sigmoid function within the differential equation. Noise added via the SDE model is controlled by a single parameter ϵ , the same for all neuron types. It is determined by empirical constraint on $\frac{\sigma_S}{\mu S}$, as shown in column 2 of Figure 2.

745 4.6.3 Connection matrix generation

Connection weight matrices (eg $M^{P,K}$ etc) are generated in a standard way, from Gaussian distributions with std dev σ defined proportional to the mean μ , using a scaling factor v:

748 $M^{*,*} \sim N(\mu_c, \sigma_c^2)$ where μ_c depends on the neuron types being connected, and $\sigma_c = v\mu_c$. Many connection 749 types typically share the same v.

A special feature of the AL is that all the neurons in a given glomerulus share a common environment. For example, all the neurons, of whatever type, in glomerulus *A* will share the same strong (or weak) LN axon from glomerulus *B*. Thus, the RN, LN, and PNs in a given glomerulus are all correlated. In addition, neuron types are correlated. To model this dual set of correlations, connection matrices in the AL are generated as follows. As an example, consider LN connection matrices in the AL:

- 755 1. A glomerulus-glomerulus connection matrix $M^{L,G}$ is created, which defines LN arborization at the 756 glomerular level.
- 2. This connection matrix is multiplied by a neural type-specific value to give $M^{L,P}, M^{L,L}$, and $M^{L,R}$ connection matrices. This is particularly important when tuning the various inhibitory effects of LNs on RNs, PNs (QNs), and LNs.
- 3. Sensitivity to GABA: A separate variance factor determines glomerular sensitivity to GABA (ie sensitivity to inhibition). This is tuned to match data in the literature Hong and Wilson (2015), and applies to LN-to-PN(QN) (ie $M^{L,P}$) connections only.
- The goal of this two-stage approach is to enforce two types of similarity found in the AL: (i) Connections to all neurons within a single glomerulus are correlated; and(ii) connections to all neurons of a certain type (LN, PN, RN) are correlated.
- Due to constraints of the biological architecture there are many zero connections. For example, about 85% of entries in the AL \rightarrow MB weight matrix are zero because MB neurons connect to only \sim 10 projection neurons Caron et al. (2013). All MB \rightarrow EN weights are set equal at the start of training. Training leads rapidly to non-uniform distributions as inactive connections decay and active connections strengthen.
- 770 4.6.4 RN spontaneous firing rates
- RNs in the glomeruli of the AL have noisy spontaneous firing rates Bhandawat et al. (2007). The Network
 Model simulates this by assigning spontaneous firing rates to RNs. These spontaneous firing rates are
 drawn from a gamma distribution plus a bias:

drawn from a gamma distribution plus a bias: $\gamma(x|\alpha,\beta,b) = b + \frac{\beta^{\alpha}}{\Gamma(\alpha)}x^{\alpha-1}e^{-\beta x}$, where α,β are shape and rate parameters, and $\Gamma(\cdot)$ is the Gamma function.

This can be thought of as a source of energy injected into the system, at the furthest upstream point (absent odor). Other energy sources are odor signals and octopamine. The spontaneous firing rates of all other neurons in the Network Model are the result of their dynamics as RN spontaneous FRs propagate throughthe system.

780 4.7 Discrepancies between biology and model

There are some known discrepancies between our Network Model and the moth AL-MB. These are listedbelow.

783 4.7.1 Connection weight distributions

This model version uses gaussian distributions to generate initial connection weights. However, moths used in live experiments are older and thus presumably have modified PN \rightarrow KC and KC \rightarrow EN connection weights. If this modification was strong, we might expect the connection weight distributions to tend towards a scale-free rather than gaussian distribution Barabasi and Albert (1999). This represents an unknown discrepancy between structure parameters of the live moths used in experiments vs the model.

789 4.7.2 Hebbian pruning

The Network Model contains no pruning mechanism to offset, via decay, the Hebbian growth mechanism. Such pruning mechanisms are common in nature, so it is reasonable to suppose that one might exist in the AL-MB. The moth has inhibitory as well as excitatory feed-forward connections from AL to MB. In the Network Model, pruning is functionally replaced by Hebbian growth of $QN \rightarrow KC$ inhibitory connections, which act to inhibit KCs and thus offset the growth of excitatory $PN \rightarrow KC$ connections (this does not directly offset KC \rightarrow EN Hebbian growth). Thus omitting a separate Hebbian decay mechanism is a matter of convenience rather than a match to known biology.

797 4.7.3 Non-olfactory input to KCs

In addition to olfactory input, the KCs receive signals from other parts of the moth, eg hearing. Because this model targets only olfactory learning, it ignores these other inputs to the KCs, and reduces the total number of KCs (from \sim 4000 to \sim 2000).

801 4.7.4 Number of QNs

There are believed to be about 3-6 QNs projecting from the AL to the MB. This model sets their number at about 15. The reason is that, absent a Hebbian pruning system in the model, the QNs function as the brake on runaway increases in KC responses due to Hebbian growth. So the increased number of QNs is a compensation for the lack of a weight-decay system.

806 4.7.5 Number of ENs

This model version has only one EN, since its goal is to demonstrate simple learning. The moth itself possesses multiple ENs.

809 4.7.6 LH inhibition

810 The LH \rightarrow KC inhibitory mechanism used in this chapter is modeled as a time-invariant global signal,

811 delivered equally to all KCs. This simplifies the model parameter space while retaining the essential 812 functionality of the LH. A more refined version of LH \rightarrow KC inhibition might vary in strength according to

813 PN output, since the same PN signals that excite the KCs also excite the LH. The actual dynamics of the

814 $AL \rightarrow LH \rightarrow KC$ linkage are not known, beyond the principle that inhibition from the LH sparsifies the KC

815 codes and makes the individual KCs act as coincidence detectors.

816 4.8 in vivo neural recordings data

Model parameters were calibrated by matching Network Model performance to *in vivo* electrode readings
from the ALs of live moths. The various performance metrics are described in Results.

819 Electrode data was collected by the lab of Prof Jeff Riffell (Dept of Biology, UW). It consists of 820 timecourses of PN firing rates measured via electrode in the AL of live moths, during a variety of regimes 821 including:

- Series of 0.2 sec odor stimulations delivered without octopamine. These experiments gave data re PN
 response to odor relative to PN spontaneous (baseline) FRs, absent octopamine.
- 824 2. Series of 0.2 sec odor stimulations delivered coincident with sugar reward (which delivers octopamine).
 825 This gave data re how PN odor response is modulated by octopamine, relative to octopamine-free
 826 spontaneous FR. See Figure 7A.
- Series of 0.2 sec odor stimulations, delivered first without and then coincident with an octopamine wash
 applied to the AL. This gave data re how PN spontaneous FR and PN odor response are modulated by
 octopamine. See Figure 7B.

In most cases the applied odor consisted of a collection of 5 volatiles, which taken together stimulate many
glomeruli in the AL. It was selected to ensure sufficient odor-responsive PNs, such that inserted electrodes
would detect interesting (ie responsive) PNs. Further details re *in vivo* data collection can be found in
Shlizerman et al. (2014) and in SI. Example timecourses are shown in Figure 7.

834 4.9 Simulation setup

For Network Model learning experiments, the time sequence of events for simulations, shown in Figure 1, is as follows:

- 837 1. A period of no stimulus, to assess baseline spontaneous behavior.
- 838 2. Four odor stimuli are delivered, 16 stimulations each (two odors were used in MB sparseness839 experiments).
- 840 3. A period of control octopamine, ie without odor or Hebbian training.
- 841 4. The system is trained (odor + octopamine + Hebbian mechanism) on one of the odors.
- 842 5. A period of no stimulus, to assess post-training spontaneous behavior.
- 843 6. The odors are re-applied (16 stimulations each), without octopamine, to assess effects of training on844 odor response.

5 AUTHOR CONTRIBUTIONS

- 845 CBD developed the model, conducted experiments, and wrote the draft.
- 846 JAR provided data and domain expertise, and reviewed the draft.
- 847 JNK guided the research, and contributed to and reviewed the draft.

6 FUNDING

848 CBD gratefully acknowledges partial funding from the Swartz Foundation.

7 DATA AVAILABILITY STATEMENT

- 849 The *in vivo* datasets analyzed in this study will be found at:
- 850 https://github.com/charlesDelahunt/SmartAsABug

REFERENCES

- An, G. (1996). The effects of adding noise during backpropagation training on a generalization performance.
 Neural Comput. 8, 643–674. doi:10.1162/neco.1996.8.3.643
- Arena, P., Patané, L., Stornanti, V., Termini, P. S., Zäpf, B., and Strauss, R. (2013). Modeling the insect
 mushroom bodies: Application to a delayed match-to-sample task. *Neural Networks* 41, 202 211.
 doi:https://doi.org/10.1016/j.neunet.2012.11.013. Special Issue on Autonomous Learning
- Attenberg, J. and Provost, F. (2010). Why label when you can search?: Alternatives to active learning for
 applying human resources to build classification models under extreme class imbalance. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (New
 York, NY, USA: ACM), KDD '10, 423–432. doi:10.1145/1835804.1835859
- Babadi, B. and Sompolinsky, H. (2014). Sparseness and expansion in sensory representations. *Neuron* 83, 1213 1226. doi:https://doi.org/10.1016/j.neuron.2014.07.035
- Balkenius, A. and Hansson, B. (2012). Discrimination training with multimodal stimuli changes activity in
 the mushroom body of the hawkmoth manduca sexta. *PLOS ONE* 7, 1–7. doi:10.1371/journal.pone.
 0032133
- Barabasi, A.-L. and Albert, R. (1999). Emergence of scaling in random networks. *Science* 286, 509–512.
 doi:10.1126/science.286.5439.509
- Bazhenov, M., Huerta, R., and Smith, B. H. (2013). A computational framework for understanding
 decision making through integration of basic learning rules. *Journal of Neuroscience* 33, 5686–5697.
 doi:10.1523/JNEUROSCI.4145-12.2013
- Bazhenov, M. and Stopfer, M. (2010). Forward and back: Motifs of inhibition in olfactory processing.
 Neuron 67, 357 358. doi:http://dx.doi.org/10.1016/j.neuron.2010.07.023
- Bengio, Y. and Fischer, A. (2015). Early inference in energy-based models approximates back-propagation.
 arXiv e-prints abs/1510.02777
- Bhandawat, V., Olsen, S. R., Gouwens, N. W., Schlief, M. L., and Wilson, R. I. (2007). Sensory processing
 in the drosophila antennal lobe increases reliability and separability of ensemble odor representations. *Nature Neuroscience* 10, 1474–1482. doi:http://www.nature.com/neuro/journal/v10/n11/full/nn1976.
- 877 html
- Campbell, R., Honegger, K., Qin, H., Li, W., Demir, E., and Turner, G. (2013). Imaging a population
 code for odor identity in the drosophila mushroom body. *Journal of Neuroscience* 33, 10568–81.
 doi:https://doi.org/10.1523/JNEUROSCI.0682-12.2013
- 881 Campbell, R. A. and Turner, G. C. (2010). The mushroom body. *Current Biology* 20, R11 R12.
 882 doi:https://doi.org/10.1016/j.cub.2009.10.031
- Caron, S., Ruta, V., Abbott, L., and Axel, R. (2013). Random convergence of olfactory inputs in the
 drosophila mushroom body. *Nature* 497, 113–7. doi:10.1038/nature12063
- Caron, S. J. C. (2013). Brains don't play dice—or do they? *Science* 342, 574–574. doi:10.1126/science.
 1245982
- Cassenaer, S. and Laurent, G. (2007). Hebbian stdp in mushroom bodies facilitates the synchronous flow
 of olfactory information in locusts. *Nature* 448, 709 EP –

Christensen, T. A., Harrow, I. D., Cuzzocrea, C., Randolph, P. W., and Hildebrand, J. G. (1995). Distinct
projections of two populations of olfactory receptor axons in the antennal lobe of the sphinx moth
manduca sexta. *Chemical Senses* 20, 313–323. doi:10.1093/chemse/20.3.313

B2 Dacks, A., Riffell, J., Martin, J., Gage, S., and Nighorn, A. (2012). Olfactory modulation by dopamine in
the context of aversive learning. *Journal of Neurophysiology* 108, 539–550. doi:10.1152/jn.00159.2012

- 894 Dacks, A. M., Christensen, T. A., and Hildebrand, J. G. (2008). Modulation of olfactory information
- processing in the antennal lobe of manduca sexta by serotonin. *Journal of Neurophysiology* 99,
 2077–2085. doi:10.1152/jn.01372.2007. PMID: 18322001
- Daly, K. C., Durtschi, M. L., and Smith, B. H. (2001). Olfactory-based discrimination learning in the moth,
 manduca sexta. *Journal of Insect Physiology* 47, 375 384. doi:https://doi.org/10.1016/S0022-1910(00)
 00117-7
- Davis, R. L. (2005). Olfactory memory formation in drosophila: From molecular to systems neuroscience.
 Annual Review of Neuroscience 28, 275–302. doi:10.1146/annurev.neuro.28.061604.135651. PMID: 16022597
- Dayan, P. and Abbott, L. F. (2005). *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems* (The MIT Press)
- Deisig, N., Giurfa, M., Lachnit, H., and Sandoz, J.-C. (2006). Neural representation of olfactory mixtures
 in the honeybee antennal lobe. *European Journal of Neuroscience* 24, 1161–1174. doi:10.1111/j.
 1460-9568.2006.04959.x
- Delahunt, C. B. and Kutz, J. N. (2018). Putting a bug in ml: The moth olfactory network learns to read
 mnist. *arXiv*
- Delahunt, C. B., Maia, P. D., and Kutz, J. N. (2018). Built to last: Functional and structural mechanisms in
 the moth olfactory network mitigate effects of neural injury. *arXiv*
- Faghihi, F., Moustafa, A. A., Heinrich, R., and Wörgötter, F. (2017). A computational model of conditioning
 inspired by drosophila olfactory system. *Neural Networks* 87, 96 108. doi:https://doi.org/10.1016/j.
 neunet.2016.11.002
- Fukushima, K. (1980). Neocognitron: A self-organizing neural network model for a mechanism of pattern
 recognition unaffected by shift in position. *Biological Cybernetics* 36, 193–202
- Galizia, C. G. (2014). Olfactory coding in the insect brain: data and conjectures. *European Journal of Neuroscience* 39, 1784–1795. doi:10.1111/ejn.12558
- Ganguli, S. and Sompolinsky, H. (2012). Compressed sensing, sparsity, and dimensionality in neuronal
 information processing and data analysis. *Annual Review of Neuroscience* 35, 485–508. doi:10.1146/
 annurev-neuro-062111-150410. PMID: 22483042
- García-Sanchez, M. and Huerta, R. (2003). Design parameters of the fan-out phase of sensory systems.
 Journal of Computational Neuroscience 15, 5–17. doi:10.1023/A:1024460700856
- Grant, W. S., Tanner, J., and Itti, L. (2017). Biologically plausible learning in neural networks with
 modulatory feedback. *Neural Networks* 88, 32 48. doi:https://doi.org/10.1016/j.neunet.2017.01.007
- Gruntman, E. and Turner, G. C. (2013). Integration of the olfactory code across dendritic claws of single
 mushroom body neurons. *Nature Neuroscience* 16, 1821 EP –. Article
- Hammer, M. and Menzel, R. (1995). Learning and memory in the honeybee. *Journal of Neuroscience* 15, 1617–1630
- Hammer, M. and Menzel, R. (1998). Multiple sites of associative odor learning as revealed by local brain
 microinjections of octopamine in honeybees. *Learn Mem* 5, 146–156. 10454379[pmid]
- 932 Hebb, D. O. (1949). *The organization of behavior : a neuropsychological theory* (Wiley New York)

- Hige, T., Aso, Y., Rubin, G. M., and Turner, G. C. (2015). Plasticity-driven individualization of olfactory
 coding in mushroom body output neurons. *Nature* 526, 258 EP –
- Higham., D. J. (2001). An algorithmic introduction to numerical simulation of stochastic differential
 equations. *SIAM Rev.* 43, 525–546. doi:10.1137/S0036144500378302
- Hildebrand, J. G. (1996). Olfactory control of behavior in moths: Central processing of odor information and
 the functional significance of olfactory glomeruli. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology* 178, 5–19
- 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. *Journal of Neuroscience* 31, 11772–
 11785. doi:10.1523/JNEUROSCI.1099-11.2011
- Hong, E. J. and Wilson, R. I. (2015). Simultaneous encoding of odors by channels with diverse sensitivity
 to inhibition. *Neuron* 85, 573 589. doi:https://doi.org/10.1016/j.neuron.2014.12.040
- Hubel, D. and Wiesel, T. (1962). Receptive fields, binocular interaction, and functional architecture in the
 cat's visual cortex. *Journal of Physiology* 160, 106–154
- Huerta, R. and Nowotny, T. (2009). Fast and robust learning by reinforcement signals: Explorations in the
 insect brain. *Neural Computation* 21, 2123–2151. doi:10.1162/neco.2009.03-08-733
- Huetteroth, W. and Schachtner, J. (2005). Standard three-dimensional glomeruli of the manduca sexta
 antennal lobe : a tool to study both developmental and adult neuronal plasticity. *Cell and Tissue Research*319, 513–524. doi:10.1007/s00441-004-1016-1
- Jefferis, G. S., Potter, C. J., Chan, A. M., Marin, E. C., Rohlfing, T., Jr., C. R. M., et al. (2007).
 Comprehensive maps of drosophila higher olfactory centers: Spatially segregated fruit and pheromone
 representation. *Cell* 128, 1187 1203. doi:https://doi.org/10.1016/j.cell.2007.01.040
- Jortner, R. A., Farivar, S. S., and Laurent, G. (2007). A simple connectivity scheme for sparse coding in an
 olfactory system. *Journal of Neuroscience* 27, 1659–1669. doi:10.1523/JNEUROSCI.4171-06.2007
- Kuebler, L. S., Schubert, M., Kárpáti, Z., Hansson, B. S., and Olsson, S. B. (2012). Antennal lobe
 processing correlates to moth olfactory behavior. *Journal of Neuroscience* 32, 5772–5782. doi:10.1523/
 JNEUROSCI.6225-11.2012
- Kvello, P., Løfaldli, B., Rybak, J., Menzel, R., and Mustaparta, H. (2009). Digital, three-dimensional average shaped atlas of the heliothis virescens brain with integrated gustatory and olfactory neurons. *Frontiers in Systems Neuroscience* 3, 14. doi:10.3389/neuro.06.014.2009
- Laurent, G. (2002). Olfactory network dynamics and the coding of multidimensional signals. *Nature Reviews Neuroscience* 3, 884 EP –. Review Article
- LeCun, Y. (2015). Facebook ai director yann lecun on his quest to unleash deep learning and makemachines smarter. *IEEE Spectrum*
- Lei, H., Christensen, T. A., and Hildebrand, J. G. (2002). Local inhibition modulates odor-evoked synchronization of glomerulus-specific output neurons. *Nature Neuroscience* 5, 557 EP –. Article
- Lei, H., Reisenman, C. E., Wilson, C. H., Gabbur, P., and Hildebrand, J. G. (2011). Spiking patterns and
 their functional implications in the antennal lobe of the tobacco hornworm manduca sexta. *PLOS ONE*6, 1–11. doi:10.1371/journal.pone.0023382
- 972 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. *Nature Neuroscience* 17, 559
 EP –. Article
- Litwin-Kumar, A., Harris, K. D., Axel, R., Sompolinsky, H., and Abbott, L. (2017). Optimal degrees of
 synaptic connectivity. *Neuron* 93, 1153 1164.e7. doi:https://doi.org/10.1016/j.neuron.2017.01.030

- Luo, S. X., Axel, R., and Abbott, L. F. (2010). Generating sparse and selective third-order responses in
 the olfactory system of the fly. *Proceedings of the National Academy of Sciences* 107, 10713–10718.
 doi:10.1073/pnas.1005635107
- Ma, W. J., Beck, J. M., Latham, P. E., and Pouget, A. (2006). Bayesian inference with probabilistic
 population codes. *Nature Neuroscience* 9, 1432 EP –. Article
- Maia, P. D. and Kutz, J. N. (2017). Reaction time impairments in decision-making networks as a diagnostic
 marker for traumatic brain injuries and neurological diseases. *Journal of Computational Neuroscience* 42, 323–347
- Martin, J. P., Beyerlein, A., Dacks, A. M., Reisenman, C. E., Riffell, J. A., Lei, H., et al. (2011). The
 neurobiology of insect olfaction: Sensory processing in a comparative context. *Progress in Neurobiology* 95, 427 447. doi:https://doi.org/10.1016/j.pneurobio.2011.09.007
- Masse, N. Y., Turner, G. C., and Jefferis, G. S. (2009). Olfactory information processing in drosophila.
 Current Biology 19, R700 R713. doi:https://doi.org/10.1016/j.cub.2009.06.026
- McCulloch, W. S. and Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics* 5, 115–133. doi:10.1007/BF02478259
- Menzel, R. and Manz, G. (2005). Neural plasticity of mushroom body-extrinsic neurons in the honeybee
 brain. *Journal of Experimental Biology* 208, 4317–4332. doi:10.1242/jeb.01908
- Mosqueiro, T. S. and Huerta, R. (2014). Computational models to understand decision making and pattern
 recognition in the insect brain. *Current Opinion in Insect Science* 6, 80 85. doi:https://doi.org/10.1016/
 j.cois.2014.10.005. Pests and resistance/Parasites/Parasitoids/Biological control/Neurosciences
- Nagel, K. I. and Wilson, R. I. (2011). Biophysical mechanisms underlying olfactory receptor neuron
 dynamics. *Nature Neuroscience* 14, 208 EP –. Article
- Nowotny, T. (2009). "sloppy engineering" and the olfactory system of insects. In *Biologically Inspired Signal Processing for Chemical Sensing*, eds. S. Marco and A. Gutierrez (Springer), vol. 188. 3–32
- Nowotny, T., Huerta, R., Abarbanel, H. D. I., and Rabinovich, M. I. (2005). Self-organization in
 the olfactory system: one shot odor recognition in insects. *Biological Cybernetics* 93, 436–446.
 doi:10.1007/s00422-005-0019-7
- Olsen, S. R., Bhandawat, V., and Wilson, R. I. (2008). Excitatory interactions between olfactory processing
 channels in the drosophila antennal lobe. *Neuron* 54, 667. doi:10.1016/j.neuron.2007.05.006
- Olsen, S. R., Bhandawat, V., and Wilson, R. I. (2010). Divisive normalization in olfactory population
 codes. *Neuron* 66, 287–299. doi:10.1016/j.neuron.2010.04.009. 20435004[pmid]
- Olsen, S. R. and Wilson, R. I. (2008). Lateral presynaptic inhibition mediates gain control in an olfactory
 circuit. *Nature* 452, 956 EP –. Article
- Peng, F. and Chittka, L. (2017). A simple computational model of the bee mushroom body can explain
 seemingly complex forms of olfactory learning and memory. *Current Biology* 27, 224 230. doi:https:
 //doi.org/10.1016/j.cub.2016.10.054
- Perez-Orive, J., Mazor, O., Turner, G. C., Cassenaer, S., Wilson, R. I., and Laurent, G. (2002). Oscillations
 and sparsening of odor representations in the mushroom body. *Science* 297, 359–365. doi:10.1126/
 science.1070502
- Perisse, E., Burke, C., Huetteroth, W., and Waddell, S. (2013). Shocking revelations and saccharin
 sweetness in the study of drosophila olfactory memory. *Curr Biol* 23, R752–R763. doi:10.1016/j.cub.
 2013.07.060. S0960-9822(13)00921-4[PII], 24028959[pmid]
- Riffell, J. A., Abrell, L., and Hildebrand, J. G. (2008). Physical processes and real-time chemical measurement of the insect olfactory environment. *Journal of Chemical Ecology* 34, 837–853. doi:10.
- 1021 1007/s10886-008-9490-7

- 1022 Riffell, J. A., Lei, H., Abrell, L., and Hildebrand, J. G. (2012). Neural basis of a pollinator's buffet:
 1023 Olfactory specialization and learning in manduca sexta. *Science* doi:10.1126/science.1225483
- Riffell, J. A., Lei, H., Christensen, T. A., and Hildebrand, J. G. (2009a). Characterization and coding of
 behaviorally significant odor mixtures. *Current Biology* 19, 335 340. doi:https://doi.org/10.1016/j.cub.
 2009.01.041
- Riffell, J. A., Lei, H., and Hildebrand, J. G. (2009b). Neural correlates of behavior in the moth manduca
 sexta in response to complex odors. *Proceedings of the National Academy of Sciences* 106, 19219–19226.
 doi:10.1073/pnas.0910592106
- 1030 Roper, M., Fernando, C., and Chittka, L. (2017). Insect bio-inspired neural network provides new
 1031 evidence on how simple feature detectors can enable complex visual generalization and stimulus
 1032 location invariance in the miniature brain of honeybees. *PLOS Computational Biology* 13, 1–23.
 1033 doi:10.1371/journal.pcbi.1005333
- Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural Networks* 61, 85 117.
 doi:https://doi.org/10.1016/j.neunet.2014.09.003
- Settles, B. (2012). *Active Learning*. Synthesis Lectures on Artificial Intelligence and Machine Learning
 (Morgan & Claypool)
- Shlizerman, E., Riffell, J. A., and Kutz, J. N. (2014). Data-driven inference of network connectivity
 for modeling the dynamics of neural codes in the insect antennal lobe. *Frontiers in Computational Neuroscience* 8, 70. doi:10.3389/fncom.2014.00070
- Silbering, A. F. and Galizia, C. G. (2007). Processing of odor mixtures in the drosophila antennal
 lobe reveals both global inhibition and glomerulus-specific interactions. *Journal of Neuroscience* 27,
 11966–11977. doi:10.1523/JNEUROSCI.3099-07.2007
- Sjoholm, M. (2006). Structure and function of the moth mushroom body. *Swedish Univ of Agricultural Sciences, Alrarp* PhD thesis
- 1046 Sutton, R. S. and Barto, A. G. (1998). Reinforcement Learning (MIT Press)
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I. J., et al. (2013). Intriguing
 properties of neural networks. *CoRR* abs/1312.6199
- Turner, G. C., Bazhenov, M., and Laurent, G. (2008). Olfactory representations by drosophila mushroom
 body neurons. *Journal of Neurophysiology* 99, 734–746. doi:10.1152/jn.01283.2007
- Wilson, R. I. (2008). Neural and behavioral mechanisms of olfactory perception. *Current Opinion in Neurobiology* 18, 408 412. doi:https://doi.org/10.1016/j.conb.2008.08.015. Sensory systems
- 1053 Wilson, R. I. and Laurent, G. (2005). Role of gabaergic inhibition in shaping odor-evoked spatiotemporal
- patterns in the drosophila antennal lobe. *Journal of Neuroscience* 25, 9069–9079. doi:10.1523/
- 1055 JNEUROSCI.2070-05.2005



Figure 1. AL-MB overview. On the left is a system schematic: Chemical sensors (RNs) excite a noisy pre-amp network (AL), which feeds forward to a plastic sparse memory layer (MB), which excites readout (decision) neurons (ENs). Green lines show excitatory connections, red lines show inhibitory connections (LH inhibition of the MB is global). Light blue ovals show plastic synaptic connections into and out of the MB.

On the right are neuron timecourse outputs from each network (typical simulation) with time axes aligned vertically. Timecourses are aligned horizontally with their regions-of-origin in the schematic. The AL timecourse shows all responses within ± 2.5 std dev of mean spontaneous rate as medium blue. Responses outside this envelope are yellow-red (excited) or dark blue (inhibited). MB responses are shown as binary (active/silent). Timecourse events are as follows: (1) A period of no stimulus. All regions are silent. (2) Two odor stimuli are delivered, 3 stimulations each. AL, MB, and ENs display odor-specific responses. (3) A period of control octopamine, ie without odor or Hebbian training. AL response is varied, MB and EN are silent. (4) The system is trained (octopamine injected) on the first odor. All regions respond strongly. (5) A period of no stimulus. All regions are silent, as in (1). (6) The stimuli are re-applied. The AL returns to its pre-trained activity since it is not plastic. In contrast, the MB and EN are now more responsive to the trained odor, while response to the untrained odor is unchanged. Green dotted line in the EN represents a hypothetical "action" threshold. The moth has learned to respond to the trained odor.

This is a provisional file, not the final typeset article



Figure 2. In vivo firing rate data and model calibration: Comparison of PN firing rate activity from *in vivo* data and simulations. A: Histograms and CDFs of *in vivo* data and simulations. Col a: mean spontaneous FRs μ_s .

Col b: σ_s/μ_s of spontaneous FRs, a measure of noisiness of a PN. **Col c**: odor response, measured as distance from μ_s in σ_s units. Distance $> 2\sigma_s$ implies a strong activation/inhibition. **Col d**: odor response during octopamine, in σ_s units distance from μ_s . Note that PN responses are broadened (ie more PNs are strongly activated or inhibited). The dotted line in the CDF inset is the same as the CDF of the odor response without octopamine, to show the broadening towards both extremes. **Col e**: change in mean spontaneous FRs due to octopamine, measured in σ_s units distance from (non-octopamine) μ_s . Some PNs are excited, some are inhibited.

B: Activity of PNs indexed by increasing spontaneous FR. Blue lines = mean spontaneous FRs μ_s (cf col 1). Shaded regions = σ_s , $2\sigma_s$ envelopes (cf col 2). Solid red dots = odor response FRs (cf col 3). Hollow red dots = odor response FRs during octopamine (cf col 4). Red lines show the change in odor response FRs due to octopamine (cf broadened response). Black stars (*) = spontaneous FRs during octopamine (cf col 5). In panel A cols 3, 4, 5, the x-axes are expressed in units of σ_s , while in panel B the y-axis measures raw spikes/sec FR.



Figure 3. KC responses to odor during training: KCs respond sparsely to odor pre- and post-training, ie absent octopamine (blue and green dots and curves). Octopamine induces transient increased responsivity (red dots and curves). Training results in permanent increases in response to the trained odor, but no increase in response to control odor (green dots and curves).

A: KC response to an odor before, during, and after training. x-axis: indexed KCs (500 shown). y-axis: consistency of response (in %). The plots are for odor 1 as the trained odor (ie same data as panel B). Blue = pre-training (no octopamine). Red = during training (with octopamine); note the heightened transient response. Green = post-training (no octopamine). There is a permanent increase in the number of KCs that respond to the trained odor.

B: Response rate vs. percentage of active KCs for trained and control odors before, during, and after training.

x-axis: percentage of KCs responding at the given rate. y-axis: consistency of response (in %). Blue = pre-training. Red = during octopamine (transient). Green = post-training. The LH plot shows odor 1 as the reinforced odor. The scatterplots in (A) correspond to the three curves in this plot. Note that the permanent KC response curve shifts up and to the right (blue \rightarrow green) in the trained odor, ie more KCs respond to the odor (right shift) and they respond more consistently (upward shift). The RH plot shows odor 2 as a control. The control's permanent KC response curve does not shift.

C: As (B) above, but in this experiment odor 1 is now the control (LH plot), and odor 2 is reinforced (RH plot). In this case, the response curve of odor 2 (reinforced) shifts to the right (blue \rightarrow green), while the response curve of odor 1 (control) is unchanged.



Figure 4. Effect of training on EN FRs: A: Typical timecourse of EN responses from an experiment with a single moth. First, 16 stimulations of each odor were delivered, to establish naive odor responses. Note EN response variability due to noise in the system, especially in the AL. Next, the moth was trained on the first (blue) odor trained over 2 sessions (10 stimulations), by delivering odor and octopamine concurrently This timecourse corresponds to the {odor, #sessions} pair in the first column in panel B, at index 2 on the x-axis. Octopamine was then withdrawn, and the four odors were again delivered in series of stimulations, to establish post-training changes in EN response. The long green line represents a hypothetical trigger threshold, such that EN response > threshold would induce a distinct behavior.

B: EN response changes due to training, aggregated results with 11 noise realizations for each {odor, #sessions} pair. Each column shows results of training a given odor, color coded: blue, purple, red, green. x-axis = number of training sessions.

First row: The y-axis measures percent change in EN FR. The line shows mean percent change. The error bars show $\pm 1, 2$ stds.

Second row: The y-axis measures percent changes in EN response, relative to the trained odor (ie subtracting the trained odor's change from all odors). This shows how far each control odor lags behind the trained odor. The line shows mean percent lag. The error bars show ± 1 , 2 stds.



Figure 5. Effects of sparsity on learning and EN reliability Results for a typical experiment on a moth with two odors. **A:** EN responses timecourses for two odors, at varying levels of KC activation (a, b: <1%. c, d: 5 to 15%. e, f: 20 to 45%. Order of events: 3 stimulations of each odor as baseline, train on first odor (only one session shown), then 3 stimulations each post-training. At very sparse levels (a, b) training is focused but odor response is not reliable. At low sparsity levels (e, f) training is unfocused, boosting EN response to control odor and to background noise.

B: Two Figures of Merit (FoMs) plotted against MB sparsity. Low KC activation (high sparsity) correlates with well-focused learning, but low odor response SNR. High KC activation (low sparsity) correlates with poorly-focused learning, but high odor response SNR. The FoMs are each normalized for easier plotting. y-axis: Blue data: $\frac{\mu(f)}{\sigma(f)}$, a measure of odor EN response SNR, where f = EN odor response. Red data: $\frac{\mu(f_T)}{\mu(f_C)}$, a measure of learning focus, where $\mu(f_T) = \text{mean}$ EN post-training response to reinforced odor; $\mu(f_C) = \text{mean}$ EN post-training response to control odor (values are thresholded at 1 for plotting). A high value indicates that increases in EN response due to training were focused on the trained odor; low values indicate that irrelevant signal (F_C) was also boosted by training. The points are experimental data, the curves are cubic fits. Vertical green lines indicate the 5 - 15% sparsity region, typical in biological neural systems.



Figure 6. Schematic of AL glomeruli. Detail of neural connections within a glomerulus. Red = inhibitory, green = excitatory, blue = increases responsiveness. RNs enter from the antennae. LNs enter from other glomeruli; one full LN is shown. It is not known if octopamine modulates LNs and PNs (see section 3.1).



Figure 7. Time series of PN firing rates from *in vivo* experiments. x-axis = time, y-axis = FR. Blue lines = mean spontaneous rate, shaded regions = ± 1 and 2 std. Red dots are odor responses. Green dots are response to control (mineral oil).

A: PN response, given odor plus coincident sugar reward, ie plus octopamine (time series for PNs with odor only are similar, but with less strong odor responses). Top row: unresponsive to odor. Middle row: excited response to odor. Bottom row: inhibited response to odor.

B: PNs with octopamine wash added in mid-experiment, then rinsed away (duration shown by black line). Octopamine can alter (up, down, or not at all) the spontaneous FR and/or the odor response, so there are 9 possible modulation regimes. This grid of timecourses shows a typical PN from each regime. Top row: spontaneous FR in unaffected. Middle row: spontaneous FR is boosted. Bottom row: spontaneous FR is inhibited. First column: odor response is unaffected. Second column: odor response is boosted. Third column: odor response is inhibited.