

1 University of Sheffield

2 Modeling olfactory processing and insights on  
3 optimal learning in constrained neural networks:  
4 learning from the anatomy of the *Drosophila*  
5 mushroom body.



7 Nada Abdelrahman

8 *Supervisor:* Dr. Andrew C.Lin, Prof. Eleni Vasilaki

9 This dissertation is submitted for the degree of  
10 Doctor of Philosophy

11 *in the*

12 Department of Biomedical Science

13 July 25, 2022



14

## Declaration

15 I hereby declare that except where specific reference is made to the work of others, the contents of this dissertation  
16 are original and have not been submitted in whole or in part for consideration for any other degree or qualification  
17 in this, or any other university. This dissertation is my own work and contains nothing which is the outcome of  
18 work done in collaboration with others, except as specified in the text and Acknowledgements. This dissertation  
19 contains fewer than 65,000 words including appendices, bibliography, footnotes, tables and equations and has fewer  
20 than 150 figures.

21

Nada Abdelrahman

22

July 25, 2022



## Acknowledgments

24 First, I would like to thank my primary supervisor, Dr. Andrew Lin for his continuous support and guidance  
25 throughout my learning process as a researcher and a thinker. Especially in the challenging times of Covid-19  
26 when it was hard to stay motivated and focused, he was constantly there providing support and great mentorship  
27 to get the best out of my PhD learning opportunity. Andrew helped me to improve my scientific writing skills,  
28 in particular the ways I would structure and order the flow of my thoughts while writing a manuscript to deliver  
29 clearly and concisely my work goals to the readers .

30 In addition, I am greatly thankful for the help and support of my secondary supervisor, Prof. Eleni Vasilaki,  
31 for their wise guidance and great push to become an independent scientist. Eleni helped me to better understand  
32 how I can formalise my intuitive ideas into rigorous mathematical approaches, and to become more intuitive and  
33 clear in the ways I present my work with in formal and informal meetings.

34 I am more than grateful for both of my supervisors to the great turnout I always wanted for my PhD, topped  
35 by a published work from my PhD in a top-tier journal.

36 I would also like to thank my labmates and friends: Dr. Luca Manneschi, Dr. Avgoustinos Vouros, Dr. Chao  
37 Han, and Dr. Hoger Amin, who are great friends and scientists, from whom I learnt a lot and discussed some of  
38 my work ideas with.

39 I also would like to thank Melanie Jegu, who is a fellow PhD student in faculty of engineering for being a great  
40 friend, fellow runner and supportive throughout my PhD journey.

41 Last but not least, I want to thank my parents, Yehia Elhussieny and Mona Dessouky, for their kind continuous  
42 love, support, and believing in me and my abilities.

43 P.s. I want to thank myself for staying persistent and always dreaming big.



## Abstract

44

45 Animals adapt their systems to optimise for different competing goals at the same time. Ideally, they will reach an  
46 optimal state of equilibrium where the outcome from any goal cannot get better without at the same time making  
47 another worse off, similar to the state of Pareto optimality (Mock 2011). Animals can seek different goals like, to  
48 maintain their systems' stability and robustness, or improving their performances in a given computational task,  
49 which is reflected in their memory capacity and ability to make more rewarding decisions.

50 Many species are capable of forming associative memories, they can learn to contextualise sensory stimuli as  
51 good, bad or neutral, when they are associated by a shortly upcoming salient outcome and bias their behaviours  
52 to approach or avoid these cues in the future. In this work I will focus on modelling the associative learning in the  
53 mushroom body circuit of the fruit fly, its center of olfactory associative learning.

54 Despite of the small size of the mushroom body circuit, flies can learn to associate an odor (sensory experience)  
55 with an appetitive or aversive outcome. They do so by modifying the connections between the mushroom body  
56 intrinsic neurons, called Kenyon cells (KCs), and their downstream mushroom body output neurons (MBONs).  
57 The fly motor behaviour was found to be biased by the activity of the MBONs to either approach or avoid an odor  
58 (Aso et al. 2014a; Aso et al. 2014b).

59 Although many studies uncovered the molecular mechanisms and the neurons underpinning associative learning  
60 in different species, there has been no work done to answer some specific questions: (a) Why do the neurons in  
61 the same circuit within the same animal exhibit variability among each others in their intrinsic properties? It  
62 is unknown how variability among the same types of neurons in the same circuit and animal would eventually  
63 affect the animal's optimal behaviour in a computational task. Even previous studies that tackled inter-neuronal  
64 variability were trying to study its effect on circuits stability and were dealing with inter-neuronal variability across  
65 animals and not within an individual circuit (Marder and Goaillard 2006; Golowasch et al. 2002; Schulz, Goaillard,  
66 and Marder 2006; Schulz, Goaillard, and Marder 2007). Can the observed inter-neuronal variability be a result  
67 of some optimisation protocol that enhances the circuit computational performance, for example, memory or data  
68 performance? Or has it just happened at random?

69 (b) Learning in the cerebellum (and its alike structures in other animals like the fruit fly mushroom body)  
70 happen by long term depression (weakening) between its intrinsic neurons -encoding the sensory input- and the  
71 downstream neurons that guide the animal's motor behaviour (Ito 1989). Like in (a), I ask if this learning rule has  
72 been conserved across species for optimising some computational aspects of learning.

73 I studied these questions in a model of the mushroom body, the center of olfactory associative learning in the  
74 fruit fly. The well-detailed anatomy of the mushroom body, the existence of its great genetic toolkit and the fly  
75 connectome makes it easier to model the learning mechanisms underpinning some behaviours and map them onto  
76 neurons in the fly mushroom body. Besides, the striking similarity in the circuit structure between the mushroom  
77 body in the fruit fly and the mammalian cerebellum means our work might provide computational insights relevant  
78 across species.

79 In this 3 Chapters thesis, I will present a computational model of associative learning in the fruit fly mushroom  
80 body using realistic input odors statistics, as well as putting some constraints on the model network that were  
81 observed experimentally in the real mushroom body (e.g. the level of KCs sparse coding, the level of KCs sparse  
82 coding when their inhibitory inputs are silenced).

83 In Chapter 2, I will answer the first question, the first aim, of this thesis and show that random variability  
84 between the KCs in their intrinsic parameters will impair the fly's memory performance. I find that the random  
85 inter-KCs variability will result in a high variability among the neurons in their sparsity values, which results in very  
86 few neurons being specifically active for some odors whilst the vast majority are activated by all incoming odors,  
87 that reduces the fly's ability to distinguish between odors and their identity as 'good' rewarded or 'bad' punished  
88 odors. However, I show that compensatory variability mechanisms will rescue the memory performance. I present  
89 4 different models (activity-independent and activity-dependent rules) for how this compensatory variability can  
90 take place in real neurons. Last but not least, I show that the data from the newly released fly connectome actually  
91 reveal compensatory variability in the KCs which agree with my models' predictions.

92 In Chapter 3, I will answer the second question in this thesis and show that, under some conditions, learning  
93 by depression can be more optimal than by potentiation. I will show that if the fly's decision making policy  
94 integrates the information from the MBONs in a divisive normalisation like manner (I explain more about divisive  
95 normalisation in Chapter 3), then learning by depression will lead to a higher memory performance. I also suggest  
96 a biologically plausible implementation for this normalisation decision policy using a winner-take-all (WTA) circuit  
97 model. I predict that in a WTA circuit that integrates the MBONs outputs, the fly's memory performance will be  
98 higher under learning by depression than under potentiation if the noise in the MBONs responses is of multiplicative  
99 nature (that is, if the noise in the MBONs responses across different trials is higher at higher MBONs firing rates).



# Contents

100		
101	<b>List of Figures</b>	<b>10</b>
102	<b>1 General introduction and Background</b>	<b>23</b>
103	1.1 Olfactory associative memory in the Fruit fly <i>Drosophila</i> . . . . .	25
104	1.2 Anatomy of olfactory processing in the fruit fly . . . . .	28
105	1.2.1 The Kenyon cells . . . . .	28
106	1.2.2 The MB output neurons . . . . .	31
107	1.2.3 Dopaminergic neurons . . . . .	32
108	1.2.4 Previous computational models of memory and learning . . . . .	34
109	1.3 Focus of my thesis . . . . .	35
110	<b>2 Compensatory Variability in network parameters enhances memory performance in the <i>Drosophila</i></b>	
111	<b>mushroom body</b>	<b>37</b>
112	2.1 Introduction . . . . .	37
113	2.2 Methods . . . . .	39
114	2.2.1 Modelling KC activity . . . . .	39
115	2.2.2 Modelling olfactory associative learning . . . . .	42
116	2.2.3 Metrics for evaluating Kenyon cell odour representations . . . . .	43
117	2.2.4 Models for compensatory variability . . . . .	43
118	2.2.5 Robustness analysis . . . . .	52
119	2.2.6 Connectome analysis . . . . .	53
120	2.2.7 Code availability . . . . .	53
121	2.3 Results . . . . .	54
122	2.3.1 Performance depends on KC lifetime sparseness . . . . .	57
123	2.3.2 Compensatory tuning of KC parameters rescues memory performance . . . . .	61
124	2.4 Connectome reveals compensatory variation of input strength and numbers . . . . .	67
125	2.5 Discussion . . . . .	70

126	<b>3 Computational benefits of learning by depression emerge in the model fruit fly under a divisive</b>	
127	<b>normalisation decision making policy</b>	<b>75</b>
128	3.1 Introduction . . . . .	75
129	3.2 Normalisation as a canonical operation in neural circuits . . . . .	76
130	3.3 Mathematical analysis: Learning by depression is optimal under a divisive normalisation decision	
131	making policy. . . . .	79
132	3.4 Simulation results: Learning by depression enhances the memory performance in the fruit fly model	
133	under a divisive normalisation decision policy . . . . .	86
134	3.4.1 Learning by depression enhances the memory performance under a divisive normalisation	
135	decision policy and using real odour input responses . . . . .	88
136	3.5 Winner-take-all (WTA) circuit model to approximate divisive normalisation in different neural circuits	93
137	3.6 Discussion . . . . .	100
138	<b>4 General Discussion and Future Work</b>	<b>101</b>
139	4.1 Compensatory variability rescues the memory performance . . . . .	101
140	4.1.1 Extending compensatory variability models to the cerebellum and other MB-like neural circuits	102
141	4.1.2 Compensatory variability in other models . . . . .	105
142	4.1.3 Inter-neuronal variability is beneficial in dense coding regimes . . . . .	106
143	4.2 Significance of learning by long term depression in the mushroom body and its alike circuits . . . . .	108
144	4.3 Possible extensions and improvements to the models . . . . .	111
145	<b>References</b>	<b>112</b>

# List of Figures

146

- 147 1.1 Schematic of the classical conditioning paradigm. Learning phase in the paradigm is shown on the  
148 left of the dashed line. The conditioned stimulus (CS) can be any arbitrary odour. The temporal  
149 coupling of the CS with either the reward of food or punishment with an electric shock (unconditioned  
150 stimulus) induces appetitive or aversive memory formation in the fly's brain and biases its behaviour.  
151 Testing of the learned memory is shown on the right of the dashed line. The fly's learned memory is  
152 tested when it encounters the same CS again whence it will either approach or avoid it. . . . . 26
- 153 1.2 An illustration of the aversive memory formation by long term depression (LTD). On the top panel  
154 the fly has the same odour drive to both of the mushroom body output neurons (MBONs) which  
155 direct the fly motor behaviour. Synaptic strengths between the KCs (black circles) onto the approach  
156 (green circle) MBON are equal to these onto the avoidance (dark red circle) MBON.  
157 In the bottom panel synaptic plasticity (LTD) is induced between the active KCs and the approach  
158 MBON (avoidance MBON) in an aversive (appetitive) learning experiment. The coincident activation  
159 of subset of the KCs (orange circles) by the CS and the punishment DAN by the US like an electric  
160 shock weakens the output synapses from the active KCs to the approach promoting MBON. Weakened  
161 synapses are shown by dashed lines. The connections between the active KCs onto the avoidance  
162 MBON remains unchanged in aversive learning. In appetitive learning the activation of a reward  
163 DAN and subset of KCs induces LTD between the KCs and the avoidance MBON (not shown in this  
164 schematic). . . . . 27
- 165 1.3 Schematic of the olfactory pathway in the fruit fly. Image from [Aso et al. 2014]. AL: antennal lobe,  
166 PN: projection neurons, MB: mushroom body, LH: lateral horn . . . . . 28
- 167 1.4 **(A)** Kenyon cell (KC) is tagged with mCD8-GFP (magenta) and colabeled with  $\alpha$ -synaptotagmin  
168 (green) which show the KCs claw-like endings and its presynaptic sites, respectively. KC claws receive  
169 all of its presynaptic inputs from the boutons on the projection neurons (PNs) axon termini.  
170 **(B-D)**: Magnifications of the boxed region in A. Images from (Leiss et al. 2009). . . . . 29

171	1.5	Schematic diagram for the MB lobes and intrinsic neurons. Image taken from (Aso et al. 2014a). KCs receive their inputs from around 150 PNs in the input layer shown on the left end of the diagram. KCs axons run through the MB lobes in a highly compartmentalised fashion. Axons from different KCs types innervate different lobes of the MB. The MB lobes are subdivided into 15 compartments based on the innervation patterns of the MBONs (also see Fig.1.6) . . . . .	30
176	1.6	The MB has 21 MBONs cell types which are classified based on their innervation patterns. MBONs receive their inputs from the KC axons in the MB lobes. The MBONs compartmentalised innervation patterns subdivide the MB lobes into 15 subunits (gray rectangles). MBONs 11, 5 and 6 (encircled in black) send feedforward inputs back into the MB lobes to other MBONs. Typically MBONs send their outputs (solid arrows) to 5 main neuropils (CRE, SMP, SIP, SLP and LH) which are upstream to motor guiding circuits. MBONs somas are shown by filled circles, inputs synapses by the half circles, and axons (outputs) by directed solid arrows (see inline legend). . . . .	31
185	1.7	The MB circuit overlaid with the reward and punishment DANs. Most of the PAM DANs (innervating the lobes within the green rectangle) are reward DANs except for the $\gamma 3$ DAN. PPL1 DANs are punishment DANs and innervate the MB lobes overlaid with the red rectangles. DANs valences are opposite to the valences of MBONs found in the lobes which they innervate. PAM DANs (PPL1 DANs) innervate the lobes with the avoidance (approach) promoting MBONs. . . . .	33
190	1.8	An early work to model the visual cortex neurons using a multi perceptron network. Image taken from (Albus 1971) . . . . .	34
192	2.1	Schematic for the mushroom body network model. Projection neurons in the input layer relay the odour responses, $x_i$ , downstream to the Kenyon cells ( $y_j$ ). Kenyon cells connect randomly to the projection neurons with synaptic weights $w_{ji}$ and receive global inhibition from the APL neuron with weight $\alpha_j$ . Learning occurs when dopaminergic neurons (DANs) carrying punishment (reward) signals from the environment depress the synapses ( $v_j$ ) between the active Kenyon cells and the mushroom body output neurons (MBONs) that lead to approach (avoidance) behavior. . . . .	54

2.2 Inter-KC variability in  $w$ ,  $N$  and  $\theta$  degrades the model fly's memory performance. **(A)** Histograms of the experimentally measured distributions for: **(A1)**  $w$  (amplitude of spontaneous excitatory postsynaptic potentials in KCs, mV; data from (Turner, Bazhenov, and Laurent 2008)), **(A2)**  $N$  (number of PN inputs per KC, measured as the number of dendritic 'claws'; data from (Caron et al. 2013)), **(A3)**  $\theta$  (spiking threshold minus resting potential, mV; data from (Turner, Bazhenov, and Laurent 2008)). The overlaid black curves show log-normal ( $w$ ) and Gaussian ( $N$ ,  $\theta$ ) fits to the data.

**(B)** The model fly's memory performance (given 100 input odours), varying the parameters step by step. Fixed and variable parameters are shown by empty and filled circles, respectively. The homogeneous model (all parameters fixed,  $N = 6$ ; black) performs the best and the random model (all parameters variable; red) performs the worst. All bars are significantly different from each other unless they share the same letter annotations (a, b, etc.),  $p < 0.05$  by Wilcoxon signed-rank test (for matched models with the same PN-KC connectivity) or Mann-Whitney test (for unmatched models with different PN-KC connectivity, i.e., fixed vs. variable  $N$ ), with Holm-Bonferroni correction for multiple comparisons.  $n = 30$  model instances with different random PN-KC connectivity.

**(C)** The performance trend is consistent over a range of different conditions: **(C1)** number of input odours, **(C2)** the learning rate used to update KC-MBON weights, **(C3)** amount of noise in PN activity (half, the same, or double the noise measured in (Bhandawat et al. 2007)), **(C4)** the indeterminacy in the decision making, quantified by  $\log(c)$ , where  $c$  is the constant in the soft-max function (SI Appendix, Eq. 21). The vertical dotted lines indicate the conditions used in panel B (each condition used the best learning rate).

**(D)** As KCs receive more inputs (thus more similar inputs), inter-KC variability becomes helpful, not harmful, to memory performance, especially when all KCs receive the same inputs ( $N = 24$ ). Blue: KCs vary in excitatory weights ( $w$ ); red: KCs vary in both  $w$  and thresholds ( $\theta$ ). Data for  $N = 6$  equivalent to panel B.  $n=30$ .

**(E)** Inter-KC variability improves performance in dense coding regimes (coding levels 0.7 - 0.9) at classifying 100 odours (a hard task) or 20 odours (easy task). Left of dashed line: equivalent to panel B, for comparison. Right of dashed line: increasing coding levels, in each case without inhibition (because inhibition is constrained to decrease coding level by half, which is impossible if coding level  $> 0.5$ ).  $n=50$ . \*  $p < 0.05$ , Wilcoxon signed-rank test (D) or Mann-Whitney test (E) with Holm-Bonferroni correction for multiple comparisons. Error bars show 95% confidence intervals. . .

228 2.3 Performance depends on KC lifetime sparseness. **(A1,B1)** Diagrams of angular distance between  
 229 odours (i.e., between centroids of clusters of noisy trials; A1) and dimensionality of a system with 3  
 230 variables (B1). The system with its states scattered throughout 3D space (green) has dimensionality  
 231 3 while the system with all states on a single line (magenta) has dimensionality 1. **(A2,B2)** The  
 232 homogeneous model has higher angular distance and dimensionality than the random model ( $p <$   
 233  $0.05$ , Mann-Whitney test), matching the performance difference when coding level = 0.1, but the  
 234 opposite trend to performance when coding level = 0.9. **(C-D)** Cumulative distribution function (cdf)  
 235 of the lifetime sparseness (C) or valence specificity (D) of KCs in the homogeneous (black) and random  
 236 (red) models, across 50 model instantiations. The gap between 1.0 and the top of the cdf represents  
 237 silent KCs (lifetime sparseness and specificity undefined). At coding level 0.1, the random model has  
 238 many more silent KCs, non-sparse KCs, and non-specific KCs than the homogeneous model, but at  
 239 coding level 0.9, the random model has more KCs with high lifetime sparseness and more KCs with  
 240 high valence specificity. **(E)** High lifetime sparseness enables high valence specificity, although many  
 241 sparse KCs have low valence specificity because of random valence assignments (data here from  
 242 single model instances). **(F)** Removing the sparsest or most valence-specific KCs (corresponding  
 243 to the dashed horizontal lines in C,D) removes the performance advantage of the random model  
 244 under dense coding.  $n=50$  network instantiations; \*  $p < 0.05$ , Mann-Whitney test; error bars, 95%  
 245 confidence interval (horizontal error bars in A2,B2 are smaller than the symbols). These results are  
 246 from the 20-odour task in Fig. 2.2E. . . . . 58

247 2.4 Additional metrics supporting Figure 3. **(A)** Dimensionality can be estimated numerically using  
248 (2.23) given sufficient simulated inputs (dashed line = 50,000, the number used here). Calculations  
249 here on the homogeneous model, coding level = 0.1, with inhibition. **(B-C)** As in Fig. 3A,B,  
250 except models trained to discriminate 100 odours instead of 20 odours. The homogeneous model has  
251 higher angular distance and dimensionality than the random model ( $p < 0.05$ , Mann-Whitney test),  
252 matching the performance difference when coding level = 0.1, but the opposite trend to performance  
253 when coding level = 0.9. **(D-E)** The random model has greater standard deviation of lifetime  
254 sparseness across KCs, compared to the homogeneous model, in all conditions tested (coding level  
255 0.1 or 0.9; with or without inhibition) using 20 odours in (D) and 100 odours in (E). Note: Inhibition  
256 was omitted for comparing coding level 0.1 vs. 0.9 because our model was constrained to have the  
257 coding level without inhibition be double the coding level with inhibition, which is impossible when  
258 the coding level with inhibition is 0.9. The results in (D) are from the 20-odour task in Fig. 2.2E.  
259 **(F-G)** As in Fig. 3C,D, except models trained to discriminate 100 odours instead of 20 odours.  
260 Cumulative distribution function (cdf) of the lifetime sparseness (C) or valence specificity (D) of  
261 KCs in the homogeneous (black) and random (red) models, across 50 model instantiations. The  
262 gap between 1.0 and the top of the cdf represents silent KCs (lifetime sparseness undefined). At  
263 coding level 0.1, the random model has many more silent KCs, non-sparse KCs, and non-specific  
264 KCs than the homogeneous model, but at coding level 0.9, the random model has more KCs with  
265 high lifetime sparseness and more KCs with high valence specificity. **(H)** Reproduced from Fig.  
266 3F for comparison: Removing the sparsest or most valence-specific KCs removes the performance  
267 advantage of the random model under dense coding. \*  $p < 0.05$ , Wilcoxon signed-rank test (see Table  
268 S1). **(I)** Removing the sparsest or most valence-specific KCs generally reduces angular distance and  
269 dimensionality, but not in a way that matches the effect on performance shown in (H). Conditions  
270 are significantly different (Mann-Whitney test) unless they share a letter anotation. n=50 network  
271 instantiations; error bars, 95% confidence interval (where error bars cannot be seen, they are smaller  
272 than the symbols); performance data in B,C from the 100-odour task in Fig. 2E. . . . . 59

273 2.5 Compensation in network parameters rescues memory performance. **(A)** Schematics of different  
 274 compensation methods. **(A1)** Activity-independent compensation. Lognormal fit of experimental  
 275 distribution of the synaptic weights (Exp., red), and its component distributions for different  $N$   
 276 and  $\theta$ , for high  $N = 7$  (dashed) or low  $N = 2$  (solid). Shadings of gray indicate different values  
 277 of  $\theta$ . **(A2-4)** Mechanisms for activity-dependent homeostatic compensation. Overly active KCs  
 278 weaken excitatory input weights ( $w_{ji}$ , A2), strengthen inhibitory input weights ( $\alpha_j$ , A3), or raise  
 279 spiking thresholds ( $\theta_j$ , A4). Inactive KCs do the reverse. **(B1)** Compensation rescues performance,  
 280 alleviating the defect caused by inter-KC variability in the random model (red) compared to the  
 281 homogeneous model (black), whether compensation occurs by setting  $w$  according to  $N$  and  $\theta$  (cyan;  
 282 A1), using activity-dependent homeostatic compensation to adjust excitatory weights (blue; A2),  
 283 inhibitory weights (green; A3) or spiking thresholds (magenta; A4). **(B2)** Differences between models  
 284 are more apparent when the task is more difficult due to more stochastic decision-making ( $c = 1$   
 285 instead of  $c = 10$  in the softmax function). **(C)** Compensation reduces variability in KC lifetime  
 286 sparseness.  $n = 20$  model instances with different random PN-KC connectivity; error bars, 95%  
 287 confidence interval. All bars are significantly different from each other unless they share the same  
 288 letter annotations,  $p < 0.05$ , by Wilcoxon signed-rank test (for matched models with the same PN-  
 289 KC connectivity) or Mann-Whitney test (for unmatched models with different PN-KC connectivity,  
 290 i.e., fixed vs. variable  $N$ ), with Holm-Bonferroni correction for multiple comparisons. Annotations  
 291 below bars indicate whether parameters were fixed (empty circle), variable (filled circle), or variable  
 292 following a compensation rule ('H' for homeostatic tuning,  $f(N, \theta)$  for activity-independent tuning).  
 293 Results here are for 100 synthetic odours; see Fig. 2.6B for similar results with odours from (Hallem  
 294 and Carlson 2006). **(D)** KC excitatory input synaptic weights ( $w$ ) after tuning to equalize average  
 295 activity (blue) follow a similar distribution to experimental data (black, from Fig. 2.2A1) **(E)** KC  
 296 spiking thresholds ( $\theta$ ) after tuning to equalize average activity (magenta) have wider variability than  
 297 the experimental distribution (black, from Fig. 2.2A3). **(F)** Tuning KC inhibitory weights ( $\alpha$ ) to  
 298 equalize average activity requires many inhibitory weights to be negative, unless the coding level  
 299 without inhibition is as high as 99%. . . . . 62

300 2.6 Similar analyses to Fig. 2 and 4, using the original 110 odour responses from (Hallem and Carlson  
 301 2006). **(A)** Inter-KC variability degrades the memory performance. **(B)** Compensation as in Fig.  
 302 4 improves memory performance.  $n = 30$  (A) or 20 (B) model instances with different random  
 303 PN-KC connectivity; error bars, 95% confidence interval. The indeterminacy constant  $c$  from the  
 304 softmax equation was set to 10. Bars within a graph that do not share the same letter annotation  
 305 are significantly different,  $p < 0.05$ , Mann-Whitney or Wilcoxon test as in Fig.2.2,2.5. . . . . 63



306 2.7 Robustness of pre-tuned compensations with novel odours. **(A)** For each model fly, network param-  
 307 eters are tuned as in Fig. 2.5, on a subset of odours. At this stage, no rewards or punishments are  
 308 given, and KC output weights are not modified. Then, the model is trained to classify rewarded  
 309 and punished odours that are the same as or different from the odours used for tuning. Finally, the  
 310 model is tested on new noisy variants of the odours used for training. **(B)** Empty symbols (‘novel  
 311 environment’): models were tuned on odours from one chemical group ( $G_i$ : acids - circles, terpenes -  
 312 triangles, esters - diamonds, or alcohols - squares), then trained and tested on odours from the other  
 313 three groups ( $G_{i \neq j}$ ). Each empty symbol is paired with a matched control (filled symbols) showing  
 314 how that model would have fared in a ‘familiar’ environment: a model tuned, trained, and tested all  
 315 on the same three groups of odours as the matched ‘novel’ model was trained and tested on ( $G_{i \neq j}$ ).  
 316 **(C)** Models with activity-dependent compensation (blue, magenta, green) performed significantly  
 317 worse in novel environment than familiar environments (matching indicated by connecting lines)  
 318 ( $p < 0.05$ , Wilcoxon signed-rank test with Holm-Bonferroni correction). In contrast, models with  
 319 no compensation (black, red), or activity-independent compensation (cyan), performed similarly in  
 320 novel and familiar environments ( $p > 0.05$  except for homogeneous (black), acids, and random (red),  
 321 terpenes). Mean of 20 model instantiations, where each instantiation received a different permuta-  
 322 tion of odours (see SI Appendix). Annotations below graph indicate whether parameters were fixed  
 323 (empty circle), variable (filled circle), or variable following a compensation rule (‘H’ for homeostatic  
 324 tuning,  $f(N, \theta)$  for activity-independent tuning). . . . . 65

325 2.8 Alternative update rules for tuning KCs’ input excitatory weights. **(A)** Performance of different  
 326 models at different indeterminacy constants (A1:  $c = 10$ ; A2:  $c = 1$ ). Blue, left: the method in  
 327 the main figures, (2.47), where a given KC’s input weights are all adjusted equally (‘H’); dark blue,  
 328 middle: (2.46), where a given KC’s input weights are adjusted individually according to the average  
 329 activity of the PN (‘H<sub>indiv</sub>’); light blue, right: (2.44), where only non-silent KCs adjust their input  
 330 weights (‘H<sub>active</sub>’).  $n = 20$  model instances with different random PN-KC connectivity. Error bars  
 331 show 95% confidence interval. Bars with the same letter annotations are not significantly different  
 332 from each other; all other comparisons are significant  $p < 0.05$ , by Wilcoxon signed-rank test with  
 333 Holm-Bonferroni correction for multiple comparisons. **(B,C)** Probability distribution of the tuned  
 334 excitatory weights (compare to Fig. 4E). **(D)** The ‘H<sub>indiv</sub>’ model performs worse than the ‘H’ model in  
 335 novel environments (see legend of Fig. 5; the drop in performance from familiar to novel environments  
 336 is significantly greater for the ‘H<sub>indiv</sub>’ model,  $p < 0.05$  by Wilcoxon signed-rank test). . . . . 66

337	2.9 Variants of activity-dependent compensation models.	
338	(A) Tuning inhibitory weights to equalize KC average activity improves performance more when	
339	we remove the constraint that the coding level without inhibition be double (0.2) the coding level	
340	with inhibition (0.1). Coding level without inhibition was 0.2 (left, light green), 0.5 (middle, medium	
341	green) or 0.99 (right, dark green).	
342	(B) Better performance when spiking thresholds are tuned to equalize KC average activity (magenta)	
343	rather than KC response probability (dark magenta), under both more ( $c = 10$ , B1) and less ( $c = 1$ ,	
344	B2) deterministic decision-making.	
345	(D) Probability distribution of spiking thresholds ( $\theta$ ) after tuning them to equalize KCs' response	
346	probabilities (compare to Fig. 4E).	
347	$n = 20$ model instances with different random PN-KC connectivity. Error bars show 95% confidence	
348	interval. * $p < 0.05$ , by Mann-Whitney test with Holm-Bonferroni correction (A) or Wilcoxon	
349	signed-rank test (B). . . . .	66

2.10 Connectome analysis reveals compensatory variation in excitatory and inhibitory input strengths. **(A)** Example  $\alpha\beta$ -c KC (bodyId 5901207528) with inputs from 3 PNs (yellow/green/blue dots) and 7 dendritic APL-KC synapses (red circles). The magenta circle shows the posterior boundary of the peduncle. Line widths not to scale. **(B,C)** Mean synaptic weight ( $w$ ) per PN-KC connection is inversely related to the number of input PNs in models that tune input weights given  $N$  and  $\theta$  **(B)**, or that tune input weights to equalize average activity levels across KCs **(C)**. **(D)** In the model that tunes input inhibitory synaptic weights ( $\alpha$ ) to equalize average activity levels across KCs, inhibitory weights are directly related to the sum of excitatory weights per KC (i.e.,  $wN$ ). Note the negative values of  $\alpha$  (discussed in text). **(E,F)** Probability distributions of the number of synapses per PN-KC connection **(E)** and the number of input PNs per KC **(F)** in the different KCs subtypes ( $\alpha\beta, \gamma, \alpha'\beta'$ ). Dashed line in (E) shows our threshold for counting connections as genuine. **(G)** Mean number of input synapses per PN-KC connection (averaged across PNs for each KC) is inversely related to the number of input PNs per KC, in  $\gamma$ -main KCs (see SI Appendix, Fig. S5 for other KC types). **(H)** Mean distance of PN-KC synapses to the posterior boundary of the peduncle (presumed spike initiation zone) is directly related to the number of input PNs per KC. **(I)** The number of APL-KC synapses per KC is directly related to the total number of PN-KC synapses per KC. **(J)** Four  $\alpha\beta$ -c KCs, one from each neuroblast clone. The posterior boundary of the peduncle (magenta circles) lies where the KC axons begin to converge. **(K)** Grids show Pearson correlation coefficients ( $r$ ) between various KC parameters for all KC subtypes tested (red: positive; blue: negative). Dots indicate  $p < 0.05$  (Holm-Bonferroni corrected). Coloured outlines indicate predictions of models (cyan/blue: models tuning  $w$  **(G,H)**; green: model tuning  $\alpha$  **(I)**). Number of KCs for each subtype, left to right: 588, 222, 350, 220, 127, 119. In (B,C,G,H), red dots are medians and the widths of the violin plots represent the number of KCs in each bin. Trend lines in (D,G,H,I) show linear fits to the data. Scale bars in **(A,J)**: D, dorsal, P, posterior, M, medial. . . . . 68

2.11 Connectome analysis on all KC subtypes ( $\gamma$ -main,  $\alpha\beta$ -s, -m and -c;  $\alpha'\beta'$ -ap2 and -m). **(A-D)** Probability distributions of the number of synapses per PN-KC connection (A,C) and the number of input PNs per KC (B,D) in  $\alpha\beta$  and  $\alpha'\beta'$  KCs separated out by subtype (compare to Fig. 6E,F). **(E)** Mean number of input synapses per PN-KC connection is inversely related to the number of input PNs per KC. **(F)** Mean distance of PN-KC synapses to the posterior boundary of the peduncle (presumed spike initiation zone) is directly related to the number of input PNs per KC in  $\gamma$  and  $\alpha\beta$ -c KCs. **(G)** The number of APL-KC synapses per KC is directly related to the total number of PN-KC synapses per KC. **(H)** The number of PN-KC synapses per KCs grows sublinearly with the number of PN inputs per KC. Red dots: medians. Red lines: linear fits. Blue dashed lines: linear fits through the origin (if every PN-KC connection had the same number of synapses). Note that the red dots follow a concave function relative to both linear fits. . . . . 69

385	3.1	Types of normalisation. The red curve corresponds to the original neuron's (f-I) curve before gain control. (A) the blue curve is the result of modulating the red curve with a subtractive normalisation.	
386		(B) the blue curve is the result of modulating the red curve with a divisive normalisation. . . . .	77
387			
388	3.2	Normalisation in the olfactory neurons in the fruit fly. (A) The response of an antennal lobe neuron to input odour concentration as in Eq.(3.1). $\gamma$ and $\sigma$ values are labeled on the graphs. . . . .	78
389			
390	3.3	Each of odour A and B activates unique KCs shown in red and green respectively. Both odours activate an overlapping KC shown in blue. . . . .	81
391			
392	3.4	Changes in the output weights from KCs to MBONs after learning by depression. Learning by penalising the wrong action. . . . .	82
393			
394	3.5	Learned bias in depression is always higher than potentiation for any $n < 1$ : (A) $n = 0.8$ , (B) $n = 0.2$ . Learned bias by depression remains higher than potentiation for different conditions of $z$ : results are shown for $z=0.1, 0.4$ and $0.6$ . The gap between the learned biases under depression and potentiation closes as the magnitude of the synaptic plasticity, $z$ , increases. . . . .	84
395			
396			
397			
398	3.6	Learned bias by depression is better than potentiation for high values of $\frac{a}{k}$ , when $n > 1$ . A learning by potentiation outperforms depression for low values of the ratio $\frac{a}{k}$ . The order switches as $z$ increases. As $z$ increases, learning by depression crosses the learning curve of potentiation at lower values of $c$ . B learning by depression outperforms potentiation under more conditions of $z$ and $c$ when $\frac{a}{k} = 1$ . C Depression outperforms potentiation under all conditions of $z$ when $k$ is less than $a$ , $\frac{a}{k} > 1$ , equals 1.6. Learned biases for different values of $z$ are shown by the solid, dashed and dotted curves respectively.	85
399			
400			
401			
402			
403			
404	3.7	Toy example to show the effect of the ratio between $k$ and an option value $\{Z_i \text{ or } Z_j\}$ on shifting the probability of picking the bigger outcome from a softmax function to a divisive normalisation one. Each pixel in every heat map in this 3x4 grid is the probability of picking the bigger outcome between $Z_i$ and $Z_j$ . The heat maps with higher $k$ and/or $n$ reveals a decision function that depends only on the gap between the given options values $[\Delta = Z_i - Z_j]$ . For lower $k$ , $k = 0.1$ or $k = 1$ , and for the same gap between both options the probability of picking the bigger option grows as the values of $Z_i$ and $Z_j$ drops. This is depicted by the isolines divergence in a fan shaped beam from the left bottom corner in these maps. Whereas the isolines in the heat maps of higher $k$ run in parallel.	87
405			
406			
407			
408			
409			
410			
411			
412	3.8	Learning by potentiation and depression are equivalent under a softmax policy. . . . .	88
413	3.9	Learning by depression is better than potentiation when the difference from the opposing MBONs outputs is normalised by their sum. $n=20$ random fly networks instantiations. . . . .	89
414			

- 415 3.10 The cumulative distribution function of valence specificity in the different models. odour inputs and  
 416 flies networks instantiations ( $n=20$ ) used here are similar to those in Chapter 2 Fig. 2.5.  
 417 The random, homogeneous and activity-independent equalisation models (top panels; black, red and  
 418 cyan curves) have higher average valence specificity values than the activity dependent models in the  
 419 bottom panels (blue, green, magenta curves).  
 420 The homogeneous, random and cyan models have few perfectly specific KCs (specificity =1) while  
 421 the activity dependent equalisation models have none. The sharp rise at the end of the black, red  
 422 and cyan curves in the top panels account for the mass of the perfectly specific KCs. . . . . 90
- 423 3.11 The cumulative distribution function (cdf) of the differences between the MBONs weights for the  
 424 different models, plotted for different learning rates and under learning by depression (left half of  
 425 the figure) and potentiation (the right half). Odor inputs and flies random networks instantiations  
 426 ( $n=20$ ) are the same as in Chapter 2, Fig. 2.5. The black arrow shows the direction of increase in  
 427 the learning rate. All models have the same cdf of the differences between their MBONs weights  
 428 (at their peak learning rates) after learning by depression or potentiation under a softmax policy.  
 429 This is shown by the dashed curves (black, red and cyan dashed curves) under potentiation and  
 430 depression for the homogeneous, random and activity-independent equalisation models, and by the  
 431 dashed and dotted curves for the activity-dependent equalisation models (blue, green, and magenta  
 432 curves) under depression and potentiation, respectively. The peak learning rates in the activity-  
 433 dependent equalisation models are roughly the same in learning by potentiation and under a softmax  
 434 or a divisive normalisation policy (the dotted curves in the bottom panels under potentiation). . . . 92
- 435 3.12 The random, Homogeneous and activity-independent equalisation models have their peak perfor-  
 436 mances at different learning rates under each of the decision making policies: softmax and divisive  
 437 normalisation. The peak scores in the random, homogeneous and activity-independent equalisation  
 438 models, shown by the red, black and cyan curves respectively, are shifted more to the right under  
 439 divisive normalisation (B) than in a softmax decision policy (A),  $n=20$ . The gray shading highlights  
 440 the region of maximum performance in the red, black and cyan models under each decision rule. . . 93
- 441 3.13 Schematic of the WTA canonical circuit . . . . . 94

442	3.14 Toy example similar to Fig. 3.7 in a WTA circuit to show the effect of the noise on shifting the	
443	probability of picking the bigger outcome from a softmax like to a divisive normalisation function.	
444	Each pixel in every heat map in this 1x3 grid is the probability of picking the bigger outcome between	
445	$Z_i$ and $Z_j$ . For each pair of options values, the probability of picking the bigger option is averaged	
446	over (n=20) random trials. In each random trial (under a certain noise regime) the noise values were	
447	randomly added to the WTA circuit's neurons.	
448	The panels show 3 different noise regimes, (left panel): additive Gaussian white noise, (middle panel):	
449	Gaussian noise with zero mean and variance equal $\sqrt{Z_{i/j}}$ , (right panel): Gaussian noise with zero	
450	mean and variance equal $Z_{i/j}$ . The heat map with additive noise (left) reveals a decision function	
451	that depends only on the gap between the given options [ $\Delta = Z_i - Z_j$ ]. Similar to the maps in (Fig.	
452	3.7) with high $k$ values. The multiplicative noise regimes (middle and right) show a behaviour similar	
453	to a divisive normalisation function. Compare to the maps in (Fig. 3.7) with low values of $k$ .	
454	The WTA circuit parameters in this simulation were set as: $\alpha=0.7$ , $\beta=0.9$ , $\tau_\eta=0.5$ , $\Delta t=\frac{\tau_\eta}{50}$ , $\tau=1$ . $Z_i$	
455	and $Z_j$ run from 1 to 20 in steps of 1. . . . .	97
456	3.15 Learning by depression and potentiation are equivalent in a WTA circuit with additive noise. . . . .	98
457	3.16 performance in learning under depression is higher than in potentiation in a WTA circuit with	
458	multiplicative noise. Noise variance is the square root of the MBONs firing rate. . . . .	99
459	3.17 Performance versus the power of the multiplicative noise in the MBONs in the WTA model. Learning	
460	by depression is better than potentiation under various scales of the noise variance. Data shown for	
461	the peak performances in two models only (Homogeneous and Random) for simplicity. . . . .	99
462	4.1 Similarity between the circuit structure of the MB and the cerebellum. Image courtesy of Modi et	
463	al.,2020 . . . . .	103
464	4.2 Activity of the winner (loser) LIP neuron - shown in black and orange traces respectively- ramps up	
465	(down) as the evidences accumulate to decide the dots motion direction in a visual discrimination	
466	task. The duration of the dots motion stimulus presentation is indicated by the gray box. Decision	
467	is made by saccade eye movements when the wining neuron reaches the decision threshold shown by	
468	the dashed line. Image courtesy of (KF et al. 2007) . . . . .	110

## Chapter 1

# General introduction and Background

Organisms continually evolve to maximize their chances of survival. As the environment pushes towards the survival of certain traits, species evolution also takes place under some constraints. These constraints could be due to the physical design of their biological systems (e.g. number of neurons, synaptic capacity or the neural circuits anatomy) or other energy constraints. For example, the human cortex has evolved to accommodate for learning new tasks and motor acts, which will require the cortex layer to expand as well as the brain to grow in size (Barton and Venditti 2014). To maximise learning capacity while constraining the brain size to remain moderate, the cortex cells layer grew by increasing its surface area and folding of the cortical layer, while keeping its thickness the same as in other mammals (Rakic 2009; Hofman 2014).

Another form of constrained evolution is the observed negative correlation between the brain size and the digestive tract size in anthropoid primates. A widely accepted hypothesis in palaeoanthropology and other fields, coined as the ‘expensive-tissue’ hypothesis (Aiello and Wheeler 1995), was used to explain how humans adapted to satisfy the high energy demands of their bigger brains by minimising the energy requirements in other organs, their guts. Moreover, the number of neurons in the olfactory circuits in mammals and invertebrates was found to have evolved non-trivially. It depends on the expected lifetime of the animal and evolved in order to optimise its associative learning performance and survival (Hiratani and Latham 2022). The biological systems have also evolved mechanisms to be always ready for any noisy changes in their gene expressions. Taking neurons as an example, while neurons live for appreciable long times in humans and animals, their constituent biochemical molecules and proteins are turned over every few days and even hours. To maintain their desired functionality, however, neurons undergo homeostatic mechanisms to compensate for any inevitable failure (Marder and Goaillard 2006).

Most of the species share the amazing premise of synaptic plasticity. No matter how small neural circuits are, all beings have to learn from and memorise past experiences. The process of memory formation has been studied many years ago. Vertebrates and invertebrates share the basic common mechanisms for memory formation, like long-term depression and potentiation. The coincidence of activity in the presynaptic and postsynaptic neurons, with or in the absence of neuromodulatory chemical release, will induce long term changes that will potentiate or

495 depress the synaptic strengths between the pre and post-synaptic partners.

496 It has been thought for long that plasticity in vertebrates is Hebbian like, that is the synaptic modulation  
497 only depends on the correlation between the presynaptic and postsynaptic neurons activities; 'neurons that fire  
498 together, wire together' (Hebb 2005; Bliss and Collingridge 1993; Caporale and Dan 2008; Glanzman 2010) without  
499 the need for neuromodulatory chemical release. On the other hand, in invertebrates, it was firmly established  
500 that the release of neuromodulatory chemicals and the sensory stimulus-evoked activity presynaptically, CS, were  
501 necessary for plasticity formation; exclusively due to long term presynaptic changes (Glanzman 2010; Castellucci  
502 and Kandel 1974; Castellucci and Kandel 1976). Some recent studies, however, have proposed that there might  
503 be some exceptions to these long standing hypotheses, which increase the resemblance between memory formation  
504 in vertebrates and invertebrates. For example, the release of neuromodulatory chemicals was found necessary for  
505 inducing the long term freezing responses (aversive memory) in the lateral amygdala pyramidal neurons in mice  
506 (Johansen et al. 2014). Also, new data has suggested the possibility of the existence of Hebbian like plasticity in  
507 invertebrates, for e.g. in the sensorimotor learning in aplysia (Lin and Glanzman 1994; Li, Roberts, and Glanzman  
508 2005), in honeybees (Menzel and Manz 2005), *Drosophila* (Xia et al. 2005), and leech (Burrell and Sahley 2004).

509 The conservation of the mechanisms underlying memory formation between vertebrates and invertebrates makes  
510 the invertebrate models a unique opportunity to study fundamental questions in neuroscience. Thanks to their  
511 simple brains and the advancements in the genetic tools we can disentangle the complicated circuits behaviours in  
512 these small animals, and try to understand the similar circuits in more complicated brains as in humans.

513 In my thesis I tried to bridge the gap in understanding the relationship between neuronal variability and its  
514 computational effects on the networks' data encoding and memory performance, in particular in sparse networks  
515 where a small percentage of the neurons are active at any point of time given an input. Sparse coding regime is  
516 an ubiquitous feature of many neural circuits across different species, it is found in the human cerebellum, fruit  
517 fly mushroom body, Honey bees, piriform cortex in rodents and many other circuits (Modi, Shuai, and Turner  
518 2020). Sparse coding is beneficial for memory capacity (Brunel et al. 2004; Földiak and Young 1995) and speed of  
519 learning (Schweighofer, Doya, and Lay 2001). Sparse encoding of densely overlapping inputs representations was  
520 first studied by Marr and Albus (Marr 1969; Albus 1971). It was also referred to as expansion recording (Albus  
521 1971) because the densely responding neurons transfer their recordings of the sensory inputs to an expanded layer  
522 of neurons, where each responds for few number of times. Later, numerous studies showed numerical and analytical  
523 evidences for the benefits of sparse coding to enhance the speed of learning, memory reliability and capacity in  
524 various simulations of the brain circuitry (Brunel et al. 2004; Babadi and Sompolinsky 2014; Schweighofer, Doya,  
525 and Lay 2001; Memmesheimer et al. 2014).

526 The computational implications of inter-neuronal variability in these circuits remained elusive; in this work I will  
527 present a computational framework and draw predictions from it about these effects. I will attempt these questions  
528 using a model of the fruit fly center of olfactory learning, the mushroom body. As we will see in Chapter 3 and in the  
529 sections below, the structure of the mushroom body circuit and its role in olfactory associative learning resembles



530 to a great extent the mammalian cerebellum. The mushroom body receives inputs stimuli through a layer of  $\approx 150$   
531 neurons (see below) which then fan out to a greater number of neurons intrinsic to the mushroom body (MB), called  
532 the Kenyon cells (KCs). The MB has  $\approx 2000$  KCs which respond very sparsely given any input. This makes them  
533 act like data encoders in the field of computer science and machine learning. Every input odour will be encoded by  
534 a unique subset of KCs; the more different these subsets are, the more will be the fly's memory performance and  
535 capacity to distinguish between different input odours. The important question I am exploring here is, given the  
536 experimental data by other studies, why do KCs show different intrinsic properties and connectivity parameters  
537 (see chapter 3)? From an information and data efficiency point of view, how would this impact the MB's memory  
538 performance, i.e. is it better for the computational nodes in the same circuit to vary in their connection parameters  
539 (e.g. number of inputs, connections strengths) or to be all identical in their properties?

540 First, let's introduce the MB circuit and olfactory associative memory formation in the fruit flies.

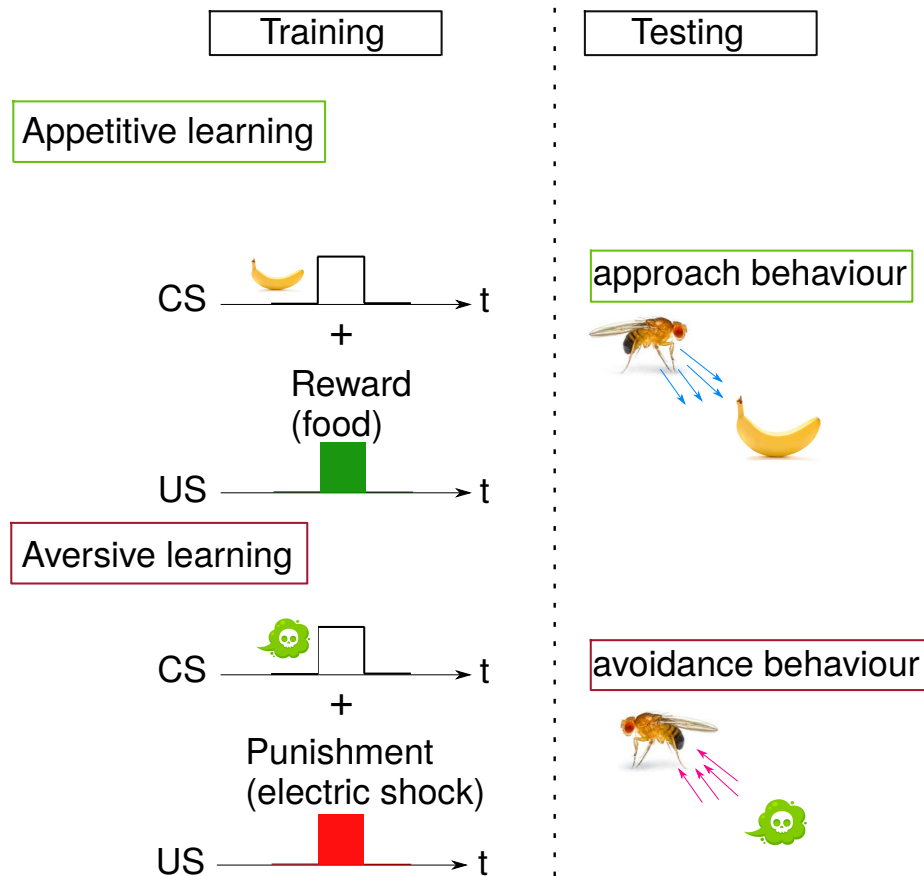
## 541 1.1 Olfactory associative memory in the Fruit fly *Drosophila*

542 *Drosophila melanogaster*, fruit fly, can develop behavioural associations with sensory stimuli. Their learning capa-  
543 bilities are powerful compared to their simple brains which have  $\approx 10^5$  neurons. Associative olfactory learning is  
544 one type of learning where the fly will experience an input odour, termed as the conditioned stimulus (CS), followed  
545 by a a sugar reward or an electric shock, the unconditioned stimulus (US), as depicted in (Fig.1.1). Flies will then  
546 learn a behaviour (approach or avoidance) which they will do upon experiencing the same CS again in the future.

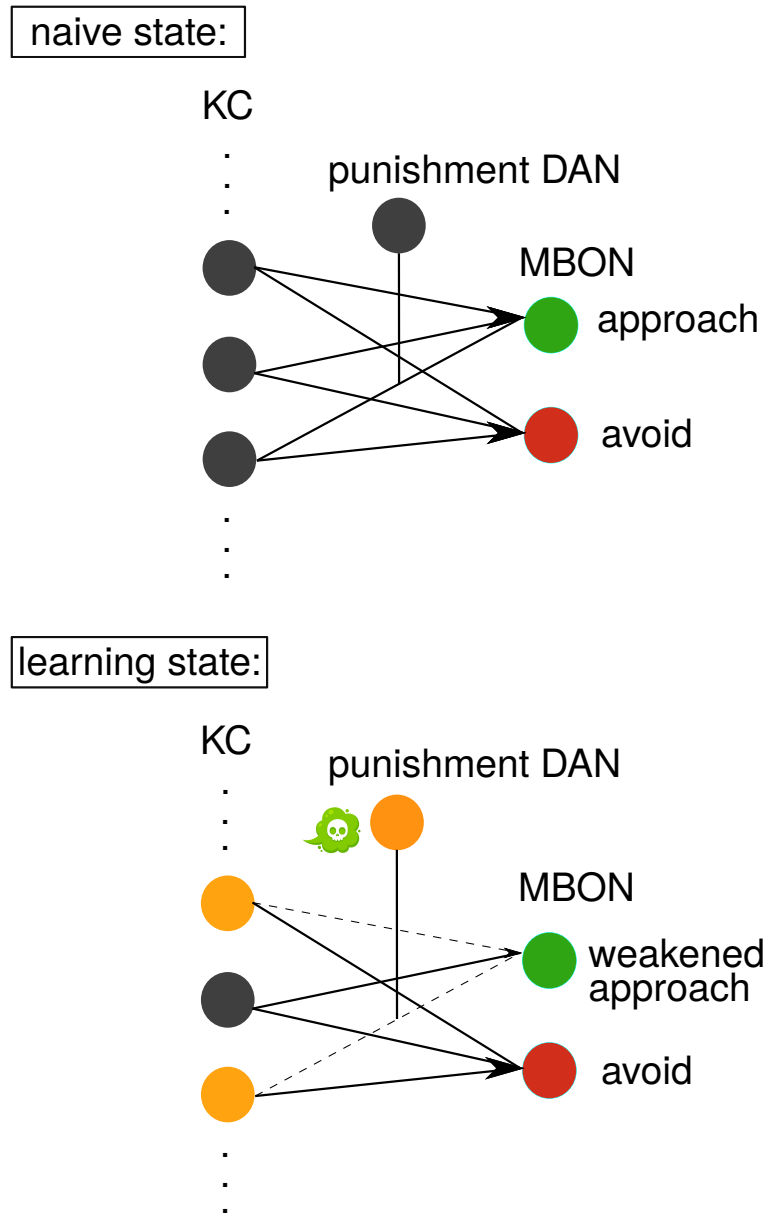
547 Olfactory learning and memory is formed in flies in a structure called, the mushroom body (MB). The simultane-  
548 ous activation of the MB intrinsic neurons (called Kenyon Cells, KCs) and dopamine release from the dopaminergic  
549 neurons (DANs) within a specific compartment (see next sections for more details) will induce long term plasticity  
550 in the synapses between the active KCs and the mushroom body output neurons (MBONs), which are upstream to  
551 intent and motor circuits.

552 At all times, the output to the 'wrong' behavior is depressed: for example, pairing an odour with electric shock  
553 weakens the output synapses from the active KCs unto MBONs that promote an approach behavior (Aso et al.  
554 2014a; Hige et al. 2015; Cohn, Morantte, and Ruta 2015; Handler et al. 2019) (reviewed in (Amin and Lin 2019)),  
555 as illustrated in (Fig.1.2).

556 MBONs Behavioural valences, which were observed by the MBONs optogenetic activation, are opposite to the  
557 DAN induced memory's type formed in their compartments (Aso et al. 2014a). In (Aso et al. 2014a), they found  
558 that the DANs activated by bitter taste or punishment reinforcement signals have mostly innervated the lobes  
559 which had MBONs with an approach valence, and vice versa. The reverse relation between the valences of the  
560 DANs memories and the MBONs in the same compartment can in part explain why learning in the MB happens  
561 by depression; as the only way to reinforce the right behaviour would be by weakening the conditioned odour drive  
562 to the MBON that promotes the wrong behaviour.




**Figure 1.1:** Schematic of the classical conditioning paradigm. Learning phase in the paradigm is shown on the left of the dashed line. The conditioned stimulus (CS) can be any arbitrary odour. The temporal coupling of the CS with either the reward of food or punishment with an electric shock (unconditioned stimulus) induces appetitive or aversive memory formation in the fly's brain and biases its behaviour. Testing of the learned memory is shown on the right of the dashed line. The fly's learned memory is tested when it encounters the same CS again whence it will either approach or avoid it.



**Figure 1.2:** An illustration of the aversive memory formation by long term depression (LTD). On the top panel the fly has the same odour drive to both of the mushroom body output neurons (MBONs) which direct the fly motor behaviour. Synaptic strengths between the KCs (black circles) onto the approach (green circle) MBON are equal to these onto the avoidance (dark red circle) MBON.

In the bottom panel synaptic plasticity (LTD) is induced between the active KCs and the approach MBON (avoidance MBON) in an aversive (appetitive) learning experiment. The coincident activation of subset of the KCs (orange circles) by the CS and the punishment DAN by the US like an electric shock weakens the output synapses from the active KCs to the approach promoting MBON. Weakened synapses are shown by dashed lines. The connections between the active KCs onto the avoidance MBON remains unchanged in aversive learning. In appetitive learning the activation of a reward DAN and subset of KCs induces LTD between the KCs and the avoidance MBON (not shown in this schematic).



**Figure 1.3:** Schematic of the olfactory pathway in the fruit fly. Image from [Aso et al. 2014]. AL: antennal lobe, PN: projection neurons, MB: mushroom body, LH: lateral horn

## 1.2 Anatomy of olfactory processing in the fruit fly


Fruit flies have around 1300 Olfactory Receptor Neurons (ORNs), in their antenna, with each ORN expressing a single, sometimes two, odourant receptors (Couto, Alenius, and Dickson 2005; Fishilevich and Vosshall 2005). These ORNs enter the antennal lobe where they send axons into one of the 50 glomerular targets, ORNs expressing the same odourants receptors converge on the same glomerulus (Couto, Alenius, and Dickson 2005; Gao, Yuan, and Chess 2000). The downstream population of  $\approx 150$  neurons, called Projection neurons (PNs), then receive their inputs from only one of the 50 glomeruli. Most of the glomeruli provide inputs to an average of  $2 \pm 1$  PNs except for 5 glomeruli which provide inputs to average of  $6 \pm 2$  PNs (Grabe et al. 2016).

Odour responses in the second order olfactory neurons (PNs) are dense, which means a large number of them will be active for any given odour. These dense (i.e. non sparse) odour responses are then relayed to the mushroom body (MB) and lateral horn neurons, termed as the third or high order olfactory centers, as illustrated by the schematic in (Fig.1.3). The lateral horn and MB were found responsible for directing the fly's innate and acquired (learned) behaviours to olfactory stimuli, respectively (Masse, Turner, and Jefferis 2009). However, new studies has challenged this distinction and found that the MB output was essential in modulating some innate behaviours (Lewis et al. 2015).

### 1.2.1 The Kenyon cells

PNs project onto the mushroom body intrinsic neurons, Kenyon cells (KCs). There are  $\approx 2000$  KCs in the MB per hemisphere. KCs originate from 4 neuroblasts (Ito et al. 1997; Lee, Lee, and Luo 1999; Zhu, Chiang, and Lee 2003). They extend their dendritic trees in the MB calyx, where their dendrites form claw-like structures which contact boutons in the PNs axons termini (Yasuyama, Meinertzhagen, and Schürmann 2002; Leiss et al. 2009), as shown in

583 (Fig. 1.4); This is also reminiscent of the granule cells to mossy fibres connections in the cerebellum (Huang et al.  
584 2013).



**Figure 1.4:** (A) Kenyon cell (KC) is tagged with mCD8-GFP (magenta) and colabeled with  $\alpha$ -synaptotagmin (green) which show the KCs claw-like endings and its presynaptic sites, respectively. KC claws receive all of its presynaptic inputs from the boutons on the projection neurons (PNs) axon termini.  
(B-D): Magnifications of the boxed region in A. Images from (Leiss et al. 2009).

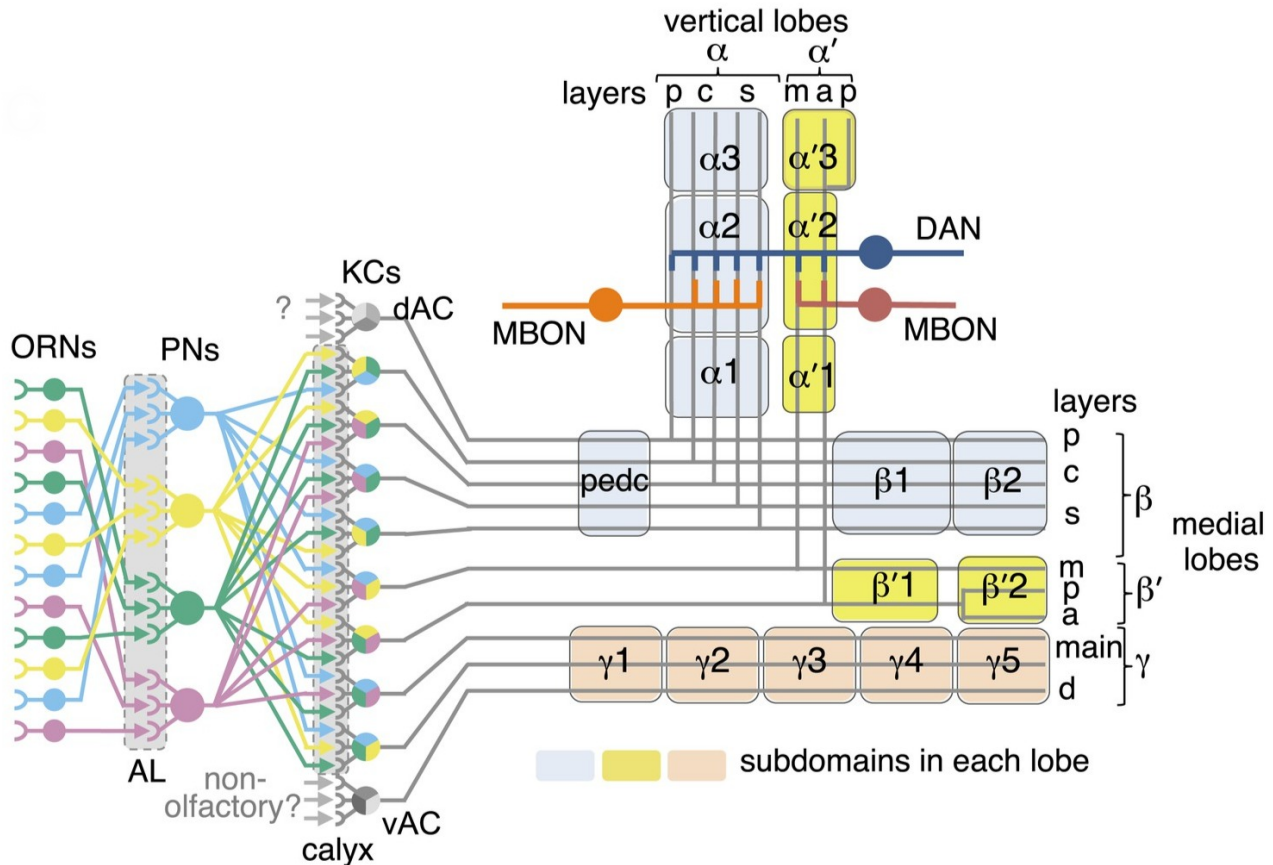
585 KCs sample their inputs from around 5.6 PNs on average (Gruntman and Turner 2013). The KCs-PNs wiring  
586 was thought to happen randomly in both the adult fly (Caron et al. 2013) and the fly larva (Eichler et al. 2017).  
587 However, new data set analyses has revealed a mild underlying structure in the wiring (Zheng et al. 2020). They  
588 showed that there is a subset of the PNs (food-odour-responsive PNs) that project more frequently on individual  
589 KCs more than expected. Their analysis revealed that these over-convergent PNs send their buttons near each  
590 other in the MB calyx, which makes the local downstream KCs more probable to receive inputs from them. We  
591 judged that attempting to model this non-randomness would not add to the realism of our model given that we  
592 modeled only 24 (out of  $\approx 50$ ) glomeruli (see Methods; Chapter 2).

593 One distinct feature about KCs is that they respond sparsely to any input. Only 10% of KCs are activated on  
594 average by any input odour (Honegger, Campbell, and Turner 2011). This percentage is also referred to as the  
595 coding level or the population sparsity of the KCs (Gruntman and Turner 2013). The sparsity of KCs responses  
596 is mostly maintained by global inhibitory feedback from the anterior paired lateral (APL) neuron (Amin and Lin  
597 2019; Scheffer et al. 2020). There is a single APL neuron per hemisphere (Scheffer et al. 2020). APL and KCs form  
598 reciprocal synaptic contacts across the entire length of the mushroom body. However, it was found that both the  
599 activity in the APL and its inhibitory effect on KCs are spatially restricted (Amin et al. 2020), which allows it to  
600 differentially inhibit different compartments in the mushroom body.

601 KCs axons exit the calyx and run in parallel where they cross the pedunculus to terminate in the respective MB  
 602 lobes. KCs have been categorized into 3 distinct classes (Lee, Lee, and Luo 1999; Crittenden et al. 1998) based on  
 603 their innervations to the eponymous lobes, as shown by the MB circuit diagram in Fig(1.5). These 3 classes are  
 604  $\alpha/\beta$ ,  $\alpha'/\beta'$  and  $\gamma$  cells. The  $\beta$ ,  $\beta'$ , and  $\gamma$  cells constitute the horizontal lobes, while the  $\alpha$  and  $\alpha'$  cells make up the  
 605 vertical lobes.

606 The three KCs classes then divide into 7 cell types based on their axonal projection patterns in the lobes [Aso  
 607 et al.,2014]. Each type occupy a specific layer in the ( $\alpha/\beta$ ), ( $\alpha'/\beta'$ ), and ( $\gamma$ ) lobes. Two KC types split the  $\gamma$  lobe,  
 608 into the main and dorsal (d) layers, two types split the ( $\alpha'/\beta'$ ) lobe into the middle (m) and anterior-posterior (ap)  
 609 layers, and three KC types divide the ( $\alpha/\beta$ ) lobe into the posterior (p), core (c), and surface (s) layers.

610 Five KC cell types have their dendritic arbors mainly in the main calyx, where they exclusively receive olfactory  
 611 inputs. The other two types:  $\gamma(d)$  and  $\alpha/\beta(p)$  receive non-olfactory inputs exclusively from the ventral and dorsal  
 612 accessory calyces (Tanaka, Tanimoto, and Ito 2008). Different KCs innervating each lobe form en passant synapses  
 613 along their axons with the MB Output Neurons (MBONs), which provide great number of connections for each  
 614 MBON.



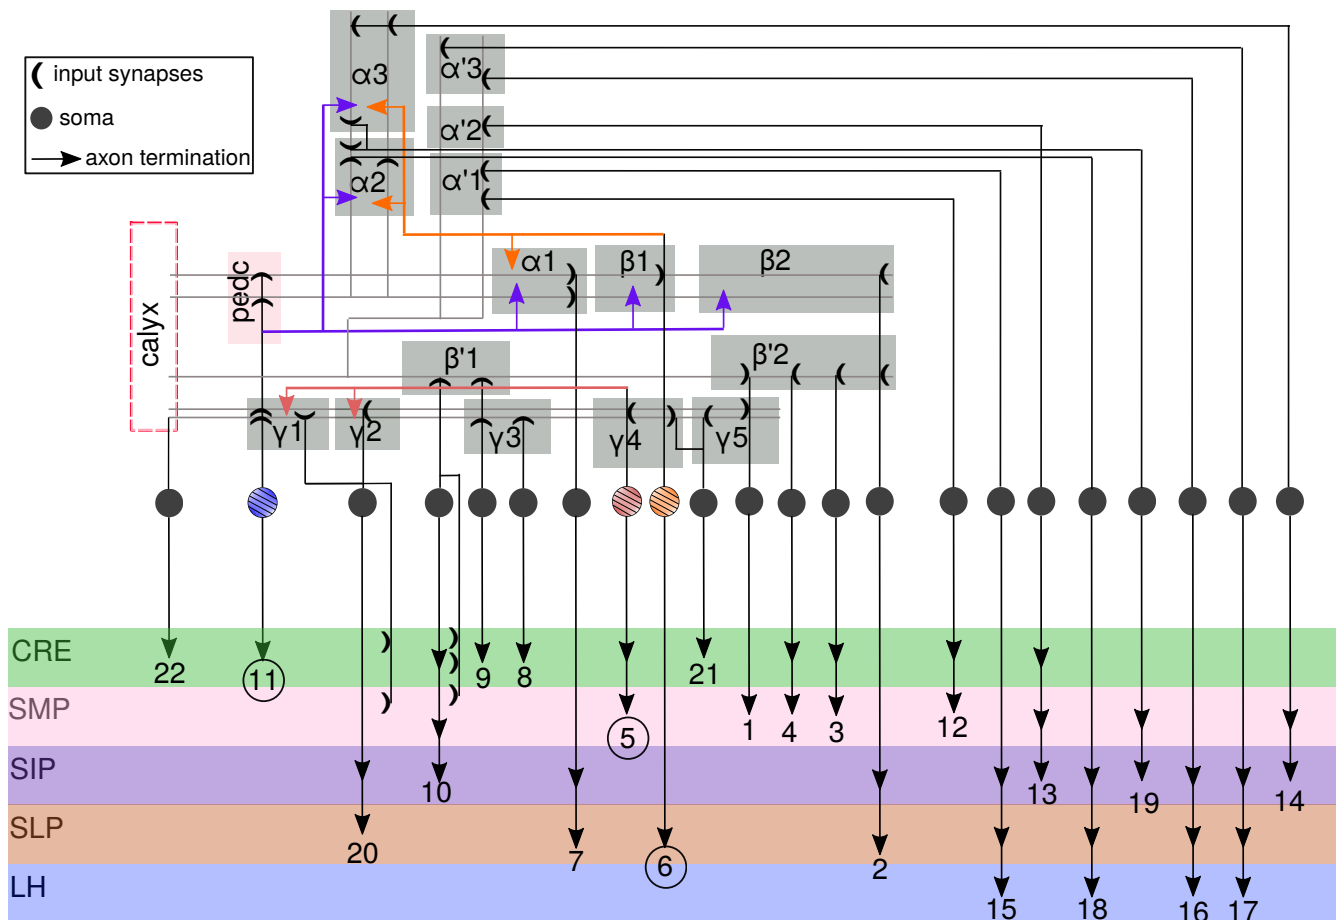
**Figure 1.5:** Schematic diagram for the MB lobes and intrinsic neurons. Image taken from (Aso et al. 2014a). KCs receive their inputs from around 150 PNs in the input layer shown on the left end of the diagram. KCs axons run through the MB lobes in a highly compartmentalised fashion. Axons from different KCs types innervate different lobes of the MB. The MB lobes are subdivided into 15 compartments based on the innervation patterns of the MBONs (also see Fig.1.6)

### 1.2.2 The MB output neurons

The information about an odour converges from the high dimensional representation in 2000 KCs to only 34 MBONs per hemisphere. The high convergence ratio of the MB outputs means that the MBONs do not encode odour identity, rather they convey an abstract representation about the odours (a behavioural bias).

MBONs are classified into 21 cell types based on their innervations to the MB lobes. The MBONs dendrites are thought to subdivide each lobe into 5 compartments, i.e. 15 compartments in all the lobes as shown in Fig.(1.5) [figure is inspired by (Aso et al. 2014a)].

MBONs dendrites tile the 15 compartments where 13 of the MBONs cell types have their inputs restricted to only one compartment and 8 project to 2 compartments, as in Fig.(1.6). While many of the MBONs receive their inputs from KCs in all the layers within a compartment, 8 types restrict their inputs to specific layers (Tanaka, Tanimoto, and Ito 2008; Aso et al. 2014b).



**Figure 1.6:** The MB has 21 MBONs cell types which are classified based on their innervation patterns. MBONs receive their inputs from the KC axons in the MB lobes.

The MBONs compartmentalised innervation patterns subdivide the MB lobes into 15 subunits (gray rectangles). MBONs 11, 5 and 6 (encircled in black) send feedforward inputs back into the MB lobes to other MBONs.

Typically MBONs send their outputs (solid arrows) to 5 main neuropils (CRE, SMP, SIP, SLP and LH) which are upstream to motor guiding circuits.

MBONs somas are shown by filled circles, inputs synapses by the half circles, and axons (outputs) by directed solid arrows (see inline legend).

A few MBONs types synapse back inside the MB lobes and provide inputs onto other MBONs, which form one layer feedforward loops (Li et al. 2020). Following the naming convention in (Aso et al. 2014a), these are MBON11, MBON05 and MBON06, labeled with black circles around them in Fig.(1.6). MBONs axons terminates mainly in 5 major neuropils outside the MB: CRE, SMP, SIP, SLP and the lateral horn (Aso et al. 2014a). To drive behavioural changes, MBONs will need to provide their outputs to motor neurons in the Ventral Nerve Cord (VNC). The brain connects to the motor neurons in VNC via hundreds of descending neurons (DNs) (Namiki et al. 2018). Although no direct connections were reported between MBONs and DNs (Li et al. 2020; Namiki et al. 2018), the optogenetic activation of the MBONs drive changes in the fly behaviour (Aso et al. 2014a), with some MBONs driving an approach behaviour and others biasing the fly to avoid the stimulus.

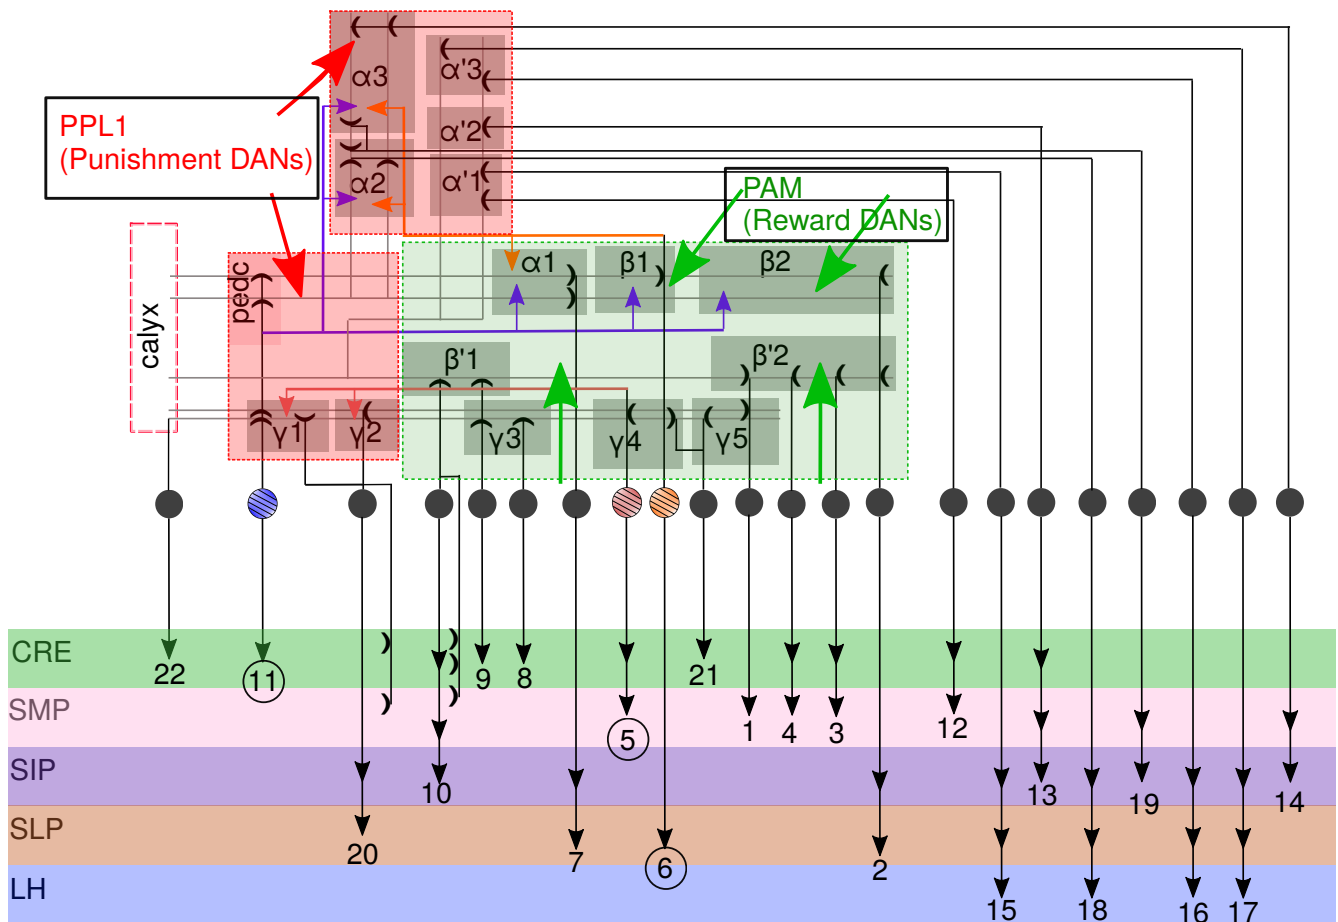
The neat compartmentalisation of the KCs-MBONs connections provides the circuit motif needed to form memory associations. In associative memory formation though the animal needs neuromodulatory signals to induce plasticity. In flies the teaching signal is provided by another type of neurons which are the Dopaminergic neurons (DANs).

### 1.2.3 Dopaminergic neurons

DANs transmit the information about reward or punishment to the MB. There are around 20 types of DANs, similar to the MBONs dendrites, each DAN has its output confined to one or 2 compartments. DANs have been categorised into 2 families, PAM and PPL1, which mostly convey information about reward and punishment respectively as depicted below in (Fig.1.7); with the exception of few PAM neurons like PAM12 ( $\gamma$ 3 DANs) which convey punishment signals (Schwaerzel et al. 2003; Claridge-Chang et al. 2009; Mao and Davis 2009; Aso et al. 2010; Burke et al. 2012; Liu et al. 2012; Hige et al. 2015).

The mushroom body (MB) has 5 PPL1 (punishment) DANs and 150 PAM (reward) DANs. DANs activation along with KCs modulate the synaptic efficacy between the active KCs and their downstream MBON in the same compartment. Besides the external sources of rewards or punishments, DANs receive direct inputs from MBONs themselves. This provides feedback loops within the MB lobes where the DANs activity is modulated by the learned odour's value as predicted by the MBONs: it is a 'good' ('bad') odour which promotes an approach (avoidance) behaviour (Li et al. 2020). In effect, this can allow the fly to learn more complex paradigms beyond pure classical conditioning like, second order conditioning and reinforcement learning. Second order conditioning was indeed observed in flies (Tabone and Belle 2011), where the value of a previously learned odour can act as a pseudo-reinforcement when coupled with another stimulus; this learning motif resembles a class of algorithms in machine learning called the actor-critic models (see Chapter 6 in (Sutton and Barto 2018) for more details). Essentially, these MBONs-DANs feedback loops were shown to be necessary for various memory processes like, nutrient-dependent consolidation of long term appetitive memory (Ichinose et al. 2015), maintenance of short term courtship memory (Zhao et al. 2018), or to 'transfer' short term memories to long term ones (Séjourné et al. 2011; Jacob and Waddell 2020; Awata et al. 2019).





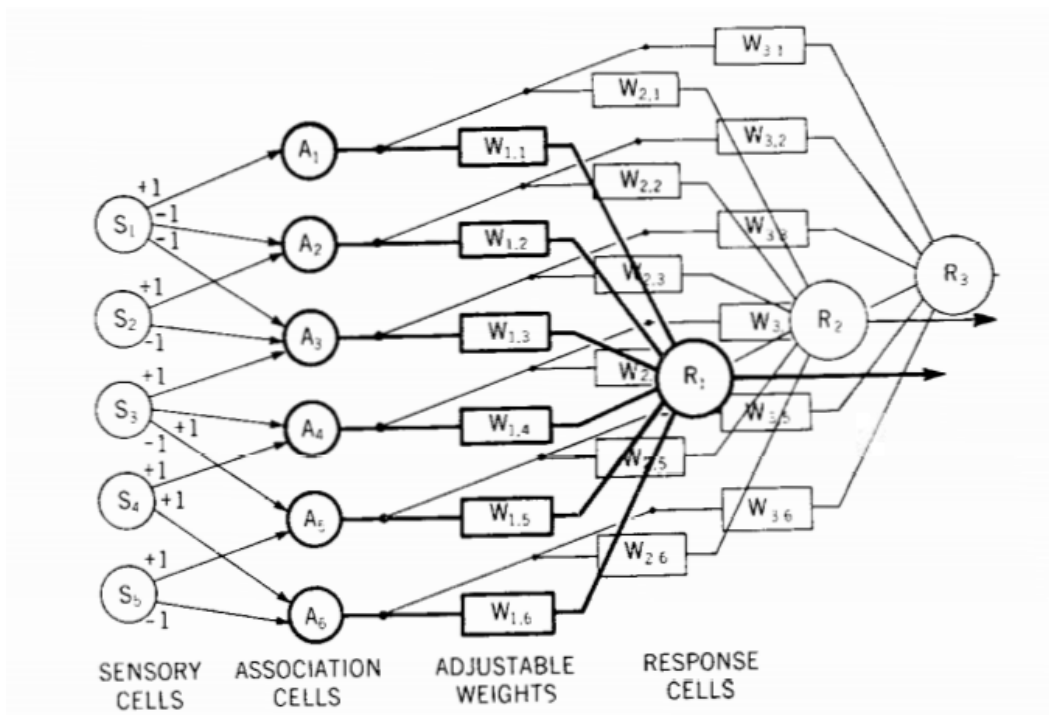
**Figure 1.7:** The MB circuit overlaid with the reward and punishment DANs. Most of the PAM DANs (innervating the lobes within the green rectangle) are reward DANs except for the  $\gamma3$  DAN. PPL1 DANs are punishment DANs and innervate the MB lobes overlaid with the red rectangles.

DANs valences are opposite to the valences of MBONs found in the lobes which they innervate. PAM DANs (PPL1 DANs) innervate the lobes with the avoidance (approach) promoting MBONs.

## 1.2.4 Previous computational models of memory and learning

Olfactory associative learning in flies and sensorimotor learning in an arm reaching task in mammals overlap in many aspects with each other. First, the end goal in both systems is to learn a motion sequence or associate an input with a class of output behaviour. Secondly, learning in the cerebellum and the mushroom body involves integration of sensory inputs, encoding these inputs in a layer of sparse neurons and most importantly receiving a neural signal that encode the error in the animal's action or an unconditioned stimulus which will induce plasticity at the outputs of the sparse nodes onto pre-motor neurons (Marr 1969; Ito 1989; Albus 1971; Aso et al. 2014a; Modi, Shuai, and Turner 2020).

Marr, Albus and Ito (Marr 1969; Albus 1971; Ito 1989; Kawato et al. 2021; Ito 2006; Ito 1972) studies were pioneering in the field of computational neuroscience, their seminal work has inspired numerous studies after to test their hypotheses in vivo (Kawato et al. 2021). Marr and Albus had formulated the objective of cerebellar learning to be pattern recognition, whilst Ito had suggested it to be a regression problem (motor control learning). Their models on learning and memory formation were the first to view the brain as a typical control circuit problem in engineering. They modeled different brain components as circuit blocks which might share feedback signals between each other. They used the classical perceptron developed by Rosenblatt in 1958 (Rosenblatt 1958) to model a neural circuit. A perceptron is made up of artificial neurons where each receives weighted sum of inputs from upstream neurons, as in Fig.1.8. The perceptron goal is to approximate (learn) the function that maps a set of inputs, e.g. sensory inputs, to desired output responses or distinct classes.



**Figure 1.8:** An early work to model the visual cortex neurons using a multi perceptron network. Image taken from (Albus 1971)

678 Besides the growing body of literature in modeling the mammalian cerebellar learning, great advances have been  
679 made in modelling the memory and learning in smaller animals. In the rest of this subsection I will focus on models  
680 done in the fruit fly.

681 Most of the computational models of the fruit fly learning and memory were done to show the significance  
682 behind the MB circuit evolution. For e.g., Litwin-Kumar et al, (Litwin-Kumar et al. 2017) have used empirical and  
683 analytical methods to prove the optimality behind the observed average number of PN inputs per KC, 7. They  
684 explained how indeed it helped the fly to achieve the best memory performance. Other work showed the significance  
685 of the circuitry that transforms the ORNs responses to the PNs and its importance to decorrelate odour responses in  
686 the higher order olfactory circuits (Luo, Axel, and Abbott 2010). In addition, some computational models underwent  
687 to elucidate the significance of the mildly non-random structure observed in the PNs-KCs connections, and show  
688 that a given structure pattern was beneficial to prioritize learning of some important odours and generalizes learning  
689 better to more similar odour groups (Zavitz et al. 2021).

690 Other computational work have used the fruit fly model as an inspiration to create better artificial networks and  
691 algorithms. In (Shen, Dasgupta, and Navlakha 2021), they showed that the fruit fly has evolved a simple continual  
692 learning algorithm that minimizes catastrophic forgetting, a huge problem in the field of machine learning and  
693 artificial neural networks. Moreover, many computational studies modeled the surprisingly advanced computational  
694 features and capabilities arising from the simple MB network like, sensory habituation (Shen, Dasgupta, and  
695 Navlakha 2020) or encoding route memory in complex natural environments (Ardin et al. 2016), or the important  
696 role of the MB sparse encoding in the rapid sensorimotor control of foraging flies (Rapp and Nawrot 2020); beside  
697 other non-elemental learning abilities in the fruit fly MB (Wessnitzer et al. 2012) and other MB networks as in the  
698 honeybee (Peng and Chittka 2017).

### 699 **1.3 Focus of my thesis**

700 In my thesis I address some unanswered gaps in neuroscience using a computational model of the fruit fly. In  
701 contrast to the previous work, I aspired to create a realistic model of the fruit fly. For that it was important to use  
702 real odours inputs to the MB intrinsic neurons (KCs), thanks to the published data in (Hallem and Carlson 2006).  
703 In addition, I modeled the intrinsic properties of the MB neurons using the published data from (Caron et al. 2013)  
704 and (Turner, Bazhenov, and Laurent 2008), I even used the statistics from real odours recordings in (Bhandawat  
705 et al. 2007) to create the noisy trials of odours responses, see Chapter 3 for more details.

706 Using my computational framework I tried to answer some main questions. First, I study the effects (advantages  
707 or disadvantages) of inter-neuronal variability from a computational point of view. To this end, I will model the  
708 inter-neuronal variability among the intrinsic neurons in the fruit fly's center of olfactory learning, the Mushroom  
709 Body, and its effects on the model fly's memory and data performance. To my knowledge, there is no other study  
710 (to the date of this thesis) which explored the consequences of inter-neuronal variability among the same type of  
711 neurons in the same circuit on the data efficiency and performance.

712 In Chapter 2 I introduce the concept of neuronal variability, and elucidate to what extent (and under which  
713 conditions) can the variability be useful for the model fly's memory performance. As we will see later, random  
714 variability is undesirable though and degrades the fly's memory performance. However, I suggest that compensatory  
715 variability models will both rescue the memory performance and keep the realistic variability between neurons  
716 intact. In addition, I found that the correlations between the neuronal parameters that are predicted from these  
717 compensatory models also exist in real neurons in the fly connectome.

718 In the last chapter of the thesis, I present a computational framework to study the potential benefits of learning  
719 by depression (weakening the wrong behaviour) over potentiation (strengthening the right behaviour). In this  
720 quest, I used the fruit fly mushroom body model, where learning has been also observed to happen by depression,  
721 to compare the model's memory performances in both cases of plasticity rules: learning by depression versus  
722 potentiation. I found that learning by depression can outperform potentiation if the noise in the mushroom body  
723 output neurons (MBONs) is multiplicative; I will explain later in Chapter 3 what I mean by multiplicative versus  
724 additive noise. In flies, the MBONs are upstream to motor guiding behaviour circuits. The results I present  
725 suggests that the plasticity rule in flies might have evolved in a non-random manner, that it has evolved such that  
726 the fly's memory performance is immune to the inevitable noise in the MBONs responses, specifically if it is of a  
727 multiplicative nature.

## 728 Chapter 2

# 729 Compensatory Variability in network 730 parameters enhances memory 731 performance in the *Drosophila* 732 mushroom body

### 733 2.1 Introduction

734 How does variability between neurons affect neural circuit function? How might neurons behave similarly despite  
735 having different underlying features? We addressed these questions in neurons called Kenyon cells, which store  
736 olfactory memories in flies. Kenyon cells differ among themselves in key features that affect how active they are,  
737 and in a model of the fly's memory circuit, adding this inter-neuronal variability made the model fly worse at  
738 learning the values of multiple odours. However, memory performance was rescued if compensation between the  
739 variable underlying features allowed Kenyon cells to be equally active on average, and we found the hypothesized  
740 compensatory variability in real Kenyon cells' anatomy. This work reveals the existence and computational benefits  
741 of compensatory variability in neural networks.

742 Noise and variability are inevitable features of biological systems. Neural circuits achieve consistent activity  
743 patterns despite this variability using homeostatic plasticity: because neural activity is governed by multiple intrinsic  
744 and network parameters, variability in one parameter can compensate for variability in another to achieve the same  
745 circuit behaviour (Golowasch et al. 2002; Achard and DeSchutter 2006; Tobin and Calabrese 2006; Taylor, Goillard,  
746 and Marder 2009; Marder and Goillard 2006). This phenomenon of compensatory variability has typically been  
747 addressed from the perspective of consistency of neural activity across individual animals (Schulz, Goillard, and  
748 Marder 2006; Schulz, Goillard, and Marder 2007) or over an animal's lifetime, in the face of circuit perturbations

749 (MacLean et al. 2003; MacLean et al. 2005; O’Leary and Marder 2016; Parrish et al. 2014). However, less attention  
750 has been paid to potential benefits of maintaining consistent neuronal properties across a population of neurons  
751 within an individual circuit.

752 Indeed, previous work has emphasized the benefits of neuronal variability/heterogeneity rather than neuronal  
753 homogeneity (Gjorgjieva, Drion, and Marder 2016; Marsat and Maler 2010; Zeldenrust, Gutkin, and Denève 2019).  
754 (Here we follow (Marder and Goaillard 2006) in using ‘heterogeneity’ to refer to qualitative differences, e.g., between  
755 cell types, and ‘variability’ to refer to quantitative differences in parameter values.) Of course, different neuronal  
756 classes encode different information (e.g., visual vs. auditory neurons, or ON vs. OFF cells). Yet even in populations  
757 that ostensibly encode the same kind of stimulus, like olfactory mitral cells, variability of neuronal excitability can  
758 increase the information content of their population activity (Padmanabhan and Urban 2010; Padmanabhan and  
759 Urban 2014; Tripathy et al. 2013). In addition, variability in neuronal time scales can improve learning in neural  
760 networks (Manneschi et al. 2021; Perez-Nieves et al. 2020). In what contexts and in what senses might the opposite  
761 be true, i.e., when does neuronal similarity provide computational benefits over neuronal variability? And what  
762 mechanisms could enforce neuronal similarity in the face of inter-neuronal variability?

763 Here we address these questions using olfactory associative memory in the mushroom body of the fruit fly  
764 *Drosophila*. Flies learn to associate specific odours with salient events (e.g., food or danger). These olfactory  
765 associative memories are stored in the principal neurons of the mushroom body, called Kenyon cells (KCs), as  
766 modifications in KCs’ output synapses (Owald et al. 2015; Handler et al. 2019; Hige et al. 2015) (reviewed in (Amin  
767 and Lin 2019)). Because learning occurs at the single output layer, the nature of the odour representation in the  
768 KC population is crucial to the fly’s ability to learn to form distinct associative memories for different odours. In  
769 particular, the fact that KCs respond sparsely to incoming odours ( $\approx 10\%$  per odour) (Honegger, Campbell, and  
770 Turner 2011) allows different odours to activate unique, non-overlapping subsets of KCs and thereby enhances flies’  
771 learned discrimination of similar odours (Lin et al. 2014).

772 A potential problem for this sparse coding arises from variability between KCs. KCs receive inputs from second-  
773 order olfactory neurons called projection neurons (PNs), with an average of  $\approx 6$  PN inputs per KC, and typically  
774 require simultaneous activation of multiple input channels in order to spike (Gruntman and Turner 2013), thanks  
775 to high spiking thresholds and feedback inhibition (Turner, Bazhenov, and Laurent 2008; Lin et al. 2014). However,  
776 there is substantial variation across KCs in the key parameters controlling their activity, such as the number of PN  
777 inputs per KC (Caron et al. 2013), the strength of PN-KC synapses, and KC spiking thresholds (Turner, Bazhenov,  
778 and Laurent 2008). Intuitively, such variation could lead to a situation where some KCs with low spiking thresholds  
779 and many or strong excitatory inputs fire indiscriminately to many different odours, while other KCs with high  
780 spiking thresholds and few or weak excitatory inputs never fire; KCs at both extremes are effectively useless for  
781 learning to classify odours, even if overall only 10% of KCs respond to each odour. However, it remains unclear  
782 whether biologically realistic inter-KC variability would affect the mushroom body’s memory performance, and  
783 what potential strategies might counter the effects of inter-KC variability.

784 Here we show in a rate-coding model of the mushroom body that introducing experimentally-derived inter-KC  
 785 variability into the model substantially impairs its memory performance. This impairment arises from increased  
 786 variability in average activity among KCs, which means fewer KCs have sparse enough activity to be specific to  
 787 rewarded vs. punished odours. However, memory performance can be rescued by compensating away variability in  
 788 KC activity while preserving the experimentally observed variation in the underlying parameters. This can occur  
 789 through activity-dependent homeostatic plasticity or direct correlations between key parameters like number vs.  
 790 strength of inputs. Finally, we analyze the hemibrain connectome to show that indeed, the number of PN inputs  
 791 per KC is inversely correlated with the strength of each input, while the strength of inhibitory inputs is correlated  
 792 with the total strength of excitatory inputs. Thus, we show both the existence and computational benefit of  
 793 compensatory variability in mushroom body network parameters.

## 794 2.2 Methods

### 795 2.2.1 Modelling KC activity

796 PN activity was simulated using the odour responses of 24 olfactory receptors (Hallem and Carlson 2006), passed  
 797 through an equation proposed by (Olsen, Bhandawat, and Wilson 2010). For an ORN and PN innervating the  $i$ th  
 798 glomerulus, their responses to the  $k$ th odour can be described using  $ORN_i^k$  (ORN activity) and  $x_i^k$  (PN activity):

$$799 \quad x_i^k = R_{max} \frac{(ORN_i^k)^{1.5}}{(ORN_i^k)^{1.5} + (s^k)^{1.5} + \sigma^{1.5}} \quad (2.1)$$

800 where  $s^k = m \sum_i ORN_i^k / 190$ ,  $m = 10.63$ , representing the gain of lateral inhibition in the antennal lobe,  $R_{max} =$   
 801 165, representing the maximum PN response, and  $\sigma = 12$ , representing the non-linearity of the ORN-PN response  
 802 function. We added noise to PN activity using:

$$803 \quad (x_i^k)_{trial} = x_i^k (1 + CoVN) \quad (2.2)$$

804 where  $CoV$  is the coefficient of variation of PN activity across trials taken from Fig. 2E of (Bhandawat et al. 2007)  
 805 and  $\mathcal{N}$  is a random sample drawn from a Gaussian distribution with mean 0 and standard deviation 1. For Fig.  
 806 2C3, the CoV was scaled by a factor of 0.5, 1 or 2. To increase the number of stimuli beyond the 110 recorded  
 807 odours in (Hallem and Carlson 2006), we generated odour responses in which the activity of each PN was randomly  
 808 sampled from that PN's activity across the 110 odours used in (Hallem and Carlson 2006), i.e.,  $x_i^k = x_i^a$  where  
 809  $k = 1 \dots K$ ,  $K$  being the number of simulated odours, and  $a$  is randomly sampled from integers from 1 to 110 for  
 810 each PN and each odour.

811 We modeled 2000 KCs. The  $j$ th KC received  $N_j$  inputs from randomly selected PNs, where  $N_j$  was either fixed  
 812 at 6 or sampled from a Gaussian distribution with mean 6 and standard deviation 1.7 (rounded to the nearest  
 813 integer; minimum 2, maximum 11), based on experimental measurements (Caron et al. 2013; Li et al. 2020). KC  
 814 claws sample PNs with replacement (Li et al. 2020; Zheng et al. 2020), so the number of unique PNs sampled by a  
 815 KC could be lower than  $N_j$ . Although more recent results show that PN-KC connectivity is not entirely random, as

816 KCs that receive inputs from a certain group of food-odour-responsive glomeruli are slightly more likely to receive  
 817 other inputs from that same group (Li et al. 2020; Zheng et al. 2020), we judged that attempting to model this  
 818 non-randomness would not add to the realism of our model given that we modeled only 24 (out of  $\approx 50$ ) glomeruli.

819 The connection from the  $i$ th PN to the  $j$ th KC had strength  $w_{ji}$ , which was 0 for non-connected neurons, and  
 820 for connected neurons was either fixed at 1, sampled from a log-normal distribution ( $\mu = -0.0507$  and  $\sigma = 0.3527$ ,  
 821 based on (Turner, Bazhenov, and Laurent 2008)), or tuned by one of the methods described below. Weights were  
 822 added for duplicate connections (i.e., KCs connected more than once to the same PN). KCs received inhibition  
 823 from APL (modeled as pseudo-feedforward for simplicity), with a gain that was either constant across all KCs  
 824 ( $\alpha$ ) or tuned individually as described below ( $\alpha_j$ ). The KCs' spiking thresholds  $\theta_j$  were either constant across all  
 825 KCs, or sampled randomly from a Gaussian distribution with coefficient of variation 0.26, based on experimental  
 826 measurements of the difference between spiking threshold and resting potential in 17 KCs (Turner, Bazhenov, and  
 827 Laurent 2008). These spiking thresholds were subject to a scaling factor  $C_\theta$  to achieve the correct average coding  
 828 level (see below). Thus, the activity of the  $j$ th KC for the  $k$ th odour,  $y_j^k$ , was

$$829 \quad y_j^k = Relu\left(\sum_{i=1}^{24} w_{ji}x_i^k - \alpha \sum_{j=1}^M \sum_{i=1}^{24} w_{ji}x_i^k - C_\theta\theta_j\right) \quad (2.3)$$

where  $M = 2000$  is the number of KCs and  $Relu$  is a rectified linear unit:

$$Relu(x) = \begin{cases} 0 & x \leq 0 \\ x & x > 0 \end{cases}$$

830 The coding level, or fraction of KCs active for each odour, averaged across odours, was defined as:

$$831 \quad CL = \frac{1}{K} \sum_{k=1}^K \left[ \frac{1}{M} \sum_{j=1}^M H(y_j^k) \right] \quad (2.4)$$

where  $K$  and  $M$  are the number of odours and KCs, respectively and  $H(x)$  is the Heaviside function:

$$H(x) = \begin{cases} 0 & \text{if } x \leq 0 \\ 1 & \text{if } x > 0 \end{cases}$$

832 Experimental data suggest that coding level is around 0.1 normally, and approximately double that (0.2) when  
 833 inhibition is blocked (Lin et al. 2014). To match these constraints, we minimized this error function with respect  
 834 to  $C_\theta$  (thus preserving the coefficient of variation of thresholds across KCs, i.e.,  $C_\theta\theta_j$ ):

$$835 \quad \epsilon_{CL|\alpha=0} = \min_{C_\theta} \left( \frac{1}{2} [CL|_{\alpha=0} - CL_{target|\alpha=0}]^2 \right) \quad (2.5)$$

836 where  $CL_{target|\alpha=0} = 0.2$ .

837 Then, we minimized the error function below with respect to  $\alpha$ :

$$838 \quad \epsilon_{CL} = \min_{\alpha} \left( \frac{1}{2} [CL - CL_{target}]^2 \right) \quad (2.6)$$

839 where  $CL_{target} = 0.1$ .



We tuned  $C_\theta$  and  $\alpha$  using gradient optimization, using the update equations:

$$\Delta C_\theta = -\eta \frac{d\epsilon_{CL}|_{\alpha=0}}{dC_\theta} \quad (2.7)$$

$$\Delta \alpha = -\eta \frac{d\epsilon_{CL}}{d\alpha} \quad (2.8)$$

To derive the update rule for  $\Delta C_\theta$ , we differentiate (2.5) with respect to  $C_\theta$ :

$$\frac{d\epsilon_{CL}|_{\alpha=0}}{dC_\theta} = [CL|_{\alpha=0} - CL_{target}|_{\alpha=0}] \frac{dCL|_{\alpha=0}}{dC_\theta} \quad (2.9)$$

To differentiate  $CL$  with respect to  $C_\theta$ , we need to replace the discontinuous Heaviside function with a continuous approximation. Similar to (Han and Kloeden 2020) a sigmoid function approximates a Heaviside at the limit  $\sigma \rightarrow 0$ ,

$$H(x) \approx S(x) = \frac{1}{1 + e^{-\frac{x}{\sigma}}} \quad (2.10)$$

Hence, assuming  $\sigma = 1$ , we can define the coding level as:

$$CL = \frac{1}{K} \sum_{k=1}^K \left[ \frac{1}{M} \sum_{j=1}^M S(y_j^k) \right] \quad (2.11)$$

Given the derivative of a sigmoid is:

$$\begin{aligned} S'(x) &= \frac{dS(x)}{dx} = \frac{e^{-\frac{x}{\sigma}}}{[1 + e^{-\frac{x}{\sigma}}]^2} \\ &= S(x)(1 - S(x)) \end{aligned} \quad (2.12)$$

Thus,

$$\begin{aligned} \frac{dCL|_{\alpha=0}}{dC_\theta} &= \frac{1}{K} \sum_{k=1}^K \left[ \frac{1}{M} \sum_{j=1}^M \left[ S'(y_j^k|_{\alpha=0}) \frac{dy_j^k|_{\alpha=0}}{dC_\theta} \right] \right] \\ &= -\frac{1}{K} \sum_{k=1}^K \left[ \frac{1}{M} \sum_{j=1}^M [S'(y_j^k|_{\alpha=0})H(y_j^k|_{\alpha=0})\theta_j] \right] \end{aligned} \quad (2.13)$$

combining (2.9) and (2.13), and plugging in (2.7) we can get the update equation for  $C_\theta$  as

$$\Delta C_\theta = \eta [CL|_{\alpha=0} - CL_{target}|_{\alpha=0}] \frac{1}{K} \sum_{k=1}^K \left[ \frac{1}{M} \sum_{j=1}^M [S'(y_j^k|_{\alpha=0})H(y_j^k|_{\alpha=0})\theta_j] \right] \quad (2.14)$$

For simplicity, this can be re-written using the average operator notation  $\langle \rangle$  across odours (indexed by  $k$ ) and KCs (indexed by  $j$ ),

$$\Delta C_\theta = \eta [CL|_{\alpha=0} - CL_{target}|_{\alpha=0}] \langle S'(y_j^k|_{\alpha=0})H(y_j^k|_{\alpha=0})\theta_j \rangle_{j,k} \quad (2.15)$$

Similarly, for  $\Delta \alpha$  we differentiate (2.6) with respect to  $\alpha$ ,

$$\frac{d\epsilon_{CL}}{d\alpha} = [CL - CL_{target}] \frac{dCL}{d\alpha} \quad (2.16)$$

862 Similarly,

$$\begin{aligned}
 \frac{dCL}{d\alpha} &= \frac{1}{K} \sum_{k=1}^K \left[ \frac{1}{M} \sum_{j=1}^M \left[ S'(y_j^k) \frac{dy_j^k}{d\alpha} \right] \right] \\
 &= -\frac{1}{K} \sum_{k=1}^K \left[ \frac{1}{M} \sum_{j=1}^M \left[ S'(y_j^k) H(y_j^k) \sum_j \sum_i w_{ji} x_i^k \right] \right]
 \end{aligned}
 \tag{2.17}$$

864 combining (2.16) with (2.17) then putting in (2.8),

$$\Delta\alpha = \eta [CL - CL_{target}] \frac{1}{MK} \sum_{k=1}^K \sum_{j=1}^M \left[ S'(y_j^k) H(y_j^k) \sum_j \sum_i w_{ji} x_i^k \right]
 \tag{2.18}$$

866 and using the  $\langle \rangle$  notation:

$$\Delta\alpha = \eta [CL - CL_{target}] \left\langle S'(y_j^k) H(y_j^k) \sum_j \sum_i w_{ji} x_i^k \right\rangle_{j,k}
 \tag{2.19}$$

868 These update equations were used to adjust values of  $\theta$  and  $\alpha$  in any random instantiation of the fly's network  
 869 to match the experimentally observed coding levels. Note that because the update equation for  $\alpha$  is the same for  
 870 all  $j$ , the same equation applies when  $\alpha_j$  is tuned for each KC (see below). In Fig. 2E and part of Fig. 3 and 2.4,  
 871  $CL_{target}$  was set to values  $> 0.1$  and  $\alpha$  was set to 0 because for  $CL_{target} > 0.5$ , it is impossible for  $CL_{target|\alpha=0}$  to  
 872 be  $2CL_{target}$ .

## 873 2.2.2 Modelling olfactory associative learning

874 Learning occurred through synaptic depression at the output synapse from KCs onto MBONs according to this  
 875 exponential decay rule:

$$\Delta v_j = v_j (e^{-\eta y_j^k} - 1)
 \tag{2.20}$$

877 where  $v_j$  is the synaptic weight between the  $j$ th KC and the MBON of the 'wrong' valence and  $\eta$  is the learning  
 878 rate. Thus, KCs active for a punished odour weaken their synapses to the approach MBON while KCs active for  
 879 the rewarded odour weaken their synapses to the avoid MBON. This can be seen as the model fly learning from  
 880 'mistakes' during its training phase (Chialvo and Bak 1999; Albus 1971).

881 The behavior of the fly was determined by a softmax equation:

$$P(\text{approach}) = \frac{e^{cMBON_{\text{approach}}}}{e^{cMBON_{\text{avoid}}} + e^{cMBON_{\text{approach}}}}
 \tag{2.21}$$

883 where the constant  $c$  governs how probabilistic or deterministic the decision-making is. At high  $c$ , the model  
 884 approaches a completely deterministic model where the fly will approach the odour 100% of the time whenever the  
 885 approach MBON's activity is higher than the avoid MBON's activity; at very low  $c$ , the model approaches random  
 886 chance; in between, the fly's behavior is probabilistic but biased by the imbalance between the activity of the two  
 887 MBONs.

888 We trained the model on 15 noisy trials of the odours (no repetitions) and tested it on 15 unseen noisy trials of  
 889 the same odours, and calculated the accuracy as the fraction of trials in which the model behaved correctly (i.e.,

890 avoided punished odours and approached rewarded odours).

### 891 2.2.3 Metrics for evaluating Kenyon cell odour representations

892 Angular distance between two vectors  $A$  and  $B$  (here,  $A$  and  $B$  are the centroids of each cluster of noisy trials of  
893 KC responses to two odours) was calculated using:

$$894 \quad \phi = \frac{2}{\pi} \arccos \frac{A \cdot B}{\|A\| \|B\|} \quad (2.22)$$

895 Dimensionality was calculated according to the equation in (Litwin-Kumar et al. 2017):

$$896 \quad \dim(\mathbf{y}) = \frac{(\sum_{i=1}^m \lambda_i)^2}{\sum_{i=1}^m \lambda_i^2} \quad (2.23)$$

897 where  $\lambda_i$  are the eigenvalues of the covariance matrix of  $\mathbf{y}$ . Whereas Litwin-Kumar et al. calculated dimensionality  
898 analytically given inputs with defined distributions, we calculated it numerically given simulated PN inputs. Because  
899 dimensionality cannot be accurately calculated with a small number of inputs (Fig. 2.4A), we simulated KC activity  
900 for 50,000 input odours for dimensionality calculations.

901 Sparseness was calculated according to (Lin et al. 2014; Willmore and Tolhurst 2001). Using the notation of  
902 this paper, the lifetime sparseness of the  $j$ th KC for a set of  $K$  odours is:

$$903 \quad S_j = \frac{1}{1 - \frac{1}{K}} \left( 1 - \frac{\left( \sum_{k=1}^K \frac{y_j^k}{K} \right)^2}{\sum_{k=1}^K \frac{(y_j^k)^2}{K}} \right) \quad (2.24)$$

904 If a cell is completely silent, firing to no stimuli,  $y_j^k = 0$  for all  $k$  and sparseness is undefined due to division by zero.

905 We defined the ‘valence specificity’  $VS$  of a KC as the degree to which it is more active for the set of rewarded  
906 odours ( $R$ ) than punished odours ( $P$ ), or vice versa:

$$907 \quad VS_j = \left| \frac{\sum_{k \in R} y_j^k - \sum_{k \in P} y_j^k}{\sum_{k \in R} y_j^k + \sum_{k \in P} y_j^k} \right| \quad (2.25)$$

### 908 2.2.4 Models for compensatory variability

909 In this section we hypothesise different mechanisms for implementing the compensatory variability given the known  
910 facts about the synapses types in the mushroom body: for example, the PNs-KCs are excitatory synapses but the  
911 KCs-APL are inhibitory synapses (Aso et al. 2010; Turner, Bazhenov, and Laurent 2008; Gruntman and Turner  
912 2013).

#### 913 Parametric tuning of excitatory input weights

914 We approximated the probability distribution of PN-KC synaptic weights ( $w$ ) using the distribution of amplitudes  
915 of spontaneous excitatory post-synaptic potentials (mini-EPSPs) in KCs, measured by (Turner, Bazhenov, and  
916 Laurent 2008). This experimental distribution was approximately log-normal, as has been described for cortical  
917 synapses (Song et al. 2005; Buzsáki and Mizuseki 2014), so we modeled  $w$  as following a log-normal distribution.  
918 We simulated values of  $w$  such that the overall distribution of  $w$  would follow this log-normal distribution, yet

individual KCs would sample  $w$  from different log-normal distributions depending on  $N$  and  $\theta$ , such that KCs with lower  $N$  or higher  $\theta$  would have higher  $w$ , i.e., sampling from a log-normal distribution shifted to the right (Fig. 4A1).

The probability of PN-to-KC synaptic weights could be estimated from the probability summation rule,

$$P(w) = \int_{\theta} \int_N P(w | N, \theta) P(N) P(\theta) dN d\theta \quad (2.26)$$

where  $P(w | N, \theta)$  is the conditional probability distribution of the input synaptic weights for a KC that has  $N$  claws and spiking threshold  $\theta$ , sampled from probability distributions  $P(N)$  and  $P(\theta)$ , respectively. We approximated  $P(N)$  and  $P(\theta)$  as the Gaussian distributions described above (see Fig. 2), and we approximated integration over  $\theta$  as summation at small intervals ( $\Delta\theta = 2.5$ ).

We modeled the constituent conditional probability distributions  $P(w | N, \theta)$  as also being log-normal, based on previous studies which approximate the sum of log-normal distributions as another log-normal variable by matching the first two moments of the power sum and its individual log-normal contributors (Fenton 1960; Schwartz and Yeh 1982; Dufresne 2008). This approximation holds in our case (the Kullback-Leibler Divergence metric (KLD) converged to less than 0.001).

To get the posterior lognormal distributions  $P(w | N, \theta)$ , we minimized the distance metric Kullback-Leibler Divergence (KLD) between  $P(w)$  and  $\int_{\theta} \int_N P(w | N, \theta) P(N) P(\theta) dN d\theta$ . To implement compensatory tuning in these conditional probabilities, such that a KC with fewer inputs (lower  $N$ ) or higher spiking threshold (higher  $\theta$ ) would have stronger inputs (higher median  $w$ ), we parameterized the medians  $\tilde{\mu}$  of each conditional distribution in  $N$  and  $\theta$  as:

$$\tilde{\mu} = \exp(\mu) = k \sqrt{\frac{\theta}{N}} \quad (2.27)$$

Thus,

$$\mu = \ln \left( k \sqrt{\frac{\theta}{N}} \right) \quad (2.28)$$

$$P(w | N, \theta) = \frac{1}{w\sigma\sqrt{2\pi}} \exp \left( - \frac{\left( \ln(w) - \ln \left( k \sqrt{\frac{\theta}{N}} \right) \right)^2}{2\sigma^2} \right) \quad (2.29)$$

We used gradient descent optimization to find the values of  $\sigma$  and  $k$  in Eq. 2.29 that would minimize the fitting error:

$$\begin{aligned} \epsilon &= KLD[P(w), \bar{P}(w)] \\ &= \int P(w) \ln \left[ \frac{\bar{P}(w)}{P(w)} \right] dw \end{aligned} \quad (2.30)$$

where

$$\bar{P}(w) = \int_{\theta} \int_N P(w | N, \theta) P(N) P(\theta) dN d\theta \quad (2.31)$$

948 First, we found the optimal  $\sigma$  by gradient optimisation:

$$949 \quad \Delta\sigma = -\eta_1 \frac{d\epsilon}{d\sigma} \quad (2.32)$$

950 The derivative of the fitting error with respect to  $\sigma$  is:

$$951 \quad \frac{d\epsilon}{d\sigma} = - \int \frac{d\bar{P}(w)}{d\sigma} \frac{P(w)}{\bar{P}(w)} dw \quad (2.33)$$

952 with,

$$953 \quad \frac{d\bar{P}(w)}{d\sigma} = \int_{\theta} \int_N \frac{dP(w | N, \theta)}{d\sigma} P(N)P(\theta) dN d\theta \quad (2.34)$$

954 where  $\frac{dP(w|N,\theta)}{d\sigma}$  is:

$$955 \quad \frac{dP(w | N, \theta)}{d\sigma} = \frac{1}{w\sigma^2\sqrt{2\pi}} \exp - \frac{\left(\ln w - \ln \left(k\sqrt{\frac{\theta}{N}}\right)\right)^2}{2\sigma^2} \left( \frac{1}{\sigma^2} \left(\ln w - \ln \left(k\sqrt{\frac{\theta}{N}}\right)\right)^2 - 1 \right) \quad (2.35)$$

956 Similarly for  $k$ ,

$$957 \quad \Delta k = -\eta_2 \frac{d\epsilon}{dk} \quad (2.36)$$

$$\frac{d\epsilon}{dk} = - \int \frac{d\bar{P}(w)}{dk} \frac{P(w)}{\bar{P}(w)} dw$$

958 such that,

$$959 \quad \frac{d\bar{P}(w)}{dk} = \int_{\theta} \int_N \frac{dP(w | N, \theta)}{dk} P(N)P(\theta) dN d\theta \quad (2.37)$$

960 with  $\frac{dP(w|N,\theta)}{dk}$  given by:

$$961 \quad \frac{dP(w | N, \theta)}{dk} = \frac{1}{kw\sigma^3\sqrt{2\pi}} \exp - \frac{\left(\ln w - \ln \left(k\sqrt{\frac{\theta}{N}}\right)\right)^2}{2\sigma^2} \left( \ln w - \ln \left(k\sqrt{\frac{\theta}{N}}\right) \right) \quad (2.38)$$

962 Starting from arbitrary values for  $k$  and  $\sigma$  and using small learning rates  $\eta_1$  and  $\eta_2$ , at each iteration, the  
 963 gradient descent algorithm alternated between using  $\sigma$  to update  $k$  and using  $k$  to update  $\sigma$ . We stopped the  
 964 gradient descent (i.e., the algorithm converged) at  $\epsilon < 0.001$ .

## 965 **Tuning KC input excitatory weights to equalize KC activity**

966 In this model, we reduce the high variance in KCs' average activity levels by tuning their input synaptic weights,  
 967 such that each  $j$ th KC adjusts its input synaptic weights ( $w_{ji}$ ) to make its average activity level  $\bar{y}_j$  reach a certain  
 968 desired level  $A_0$ . Although we ended up using a simple synaptic scaling rule in the main figures ((2.47)), we also  
 969 explored other rules based on gradient descent and describe here the mathematical relation between them. We  
 970 initially analyzed this problem using an error function:

$$971 \quad \epsilon = \frac{1}{2} [\bar{y}_j - A_0]^2 \quad (2.39)$$

$$\bar{y}_j = \frac{1}{K} \sum_{k=1}^K y_j^k$$

972 where  $y_j^k$  is the  $j$ th KC's response to the  $k$ th odour calculated as in (2.3) and  $K$  is the number of odours. Finding  
973 the weights to minimize the error in (2.39) can be found by gradient optimisation,

$$974 \quad \Delta w_{ji} = -\eta \frac{d\epsilon}{dw_{ji}} \quad (2.40)$$

975 with,

$$976 \quad \frac{d\epsilon}{dw_{ji}} = [\bar{y}_j - A_0] \frac{1}{K} \sum_{k=1}^K \frac{dy_j^k}{dw_{ji}} \quad (2.41)$$

977 Taking the derivative of  $y_j^k$  w.r.t.  $w_{ji}$  yields:

$$978 \quad \frac{dy_j^k}{dw_{ji}} = H(y_j^k)(x_i^k - \alpha x_i^k) \quad (2.42)$$

979 Plugging (2.42) in (2.41) gives:

$$980 \quad \begin{aligned} \frac{d\epsilon}{dw_{ji}} &= [\bar{y}_j - A_0] \frac{1}{K} \sum_{k=1}^K H(y_j^k)(x_i^k - \alpha x_i^k) \\ &= [\bar{y}_j - A_0] \langle H(y_j^k)(1 - \alpha)x_i^k \rangle_K \end{aligned} \quad (2.43)$$

981 Hence,  $w_{ji}$  will be updated as follows:

$$982 \quad \Delta w_{ji} = -\eta [\bar{y}_j - A_0] \langle H(y_j^k)(1 - \alpha)x_i^k \rangle_K \quad (2.44)$$

983 The equation above means that a KC with an average activity  $\bar{y}_j$  higher (lower) than  $A_0$  will scale down (up)  
984 its input synaptic weights,  $w_{ji}$ , proportional to both the difference  $(\bar{y}_j - A_0)$  and the average input activity from  
985 the  $i$ th PN. Note that in this derivation a KC must have non-zero average activity, i.e.,  $H(y_j^k) = 1$  for at least  
986 one odour, for its weights to be updated. We believe such a rule would be biologically implausible, as there should  
987 not be a discontinuity between a silent KC and a nearly silent KC. To allow totally silent KCs (which have only  
988 subthreshold activity) to update their weights in the same way as active KCs, we heuristically apply the following  
989 rule:

$$990 \quad \Delta w_{ji} = -\eta [\bar{y}_j - A_0] \langle (1 - H(y_j^k))(1 - \alpha)x_i^k \rangle_K \quad (2.45)$$

991 Adding (2.44) and (2.45) we obtain:

$$992 \quad \Delta w_{ji} = -\eta [\bar{y}_j - A_0] \langle (1 - \alpha)x_i^k \rangle_K \quad (2.46)$$

993 The rule has a fixed point  $\bar{y}_j = A_0$  since  $\langle (1 - \alpha)x_i^k \rangle_K > 0$ . Note that we apply the constraint  $w_{ji} \geq 0$ . How  
994 updates for  $w_{ji} = 0$  are treated depends on the reason why  $w_{ji} = 0$ : if the  $i$ th PN and  $j$ th KC are not connected,  
995 then the update is not applied. But if they were originally connected and the update rule pushed  $w_{ji}$  to zero, the  
996 update rule will continue to be applied.

997 To test whether performance is affected by adding the heuristic term to allow silent KCs to update their  
998 weights, we compared the performance using update rule (2.44) vs. (2.46). The rule without the heuristic performed  
999 significantly worse than the rule with the added heuristic for activating silent KCs (Fig. 2.8A). This means that a  
1000 formally derived update rule for  $w$  was not enough, since it would not equalize activity for all KCs (silent KCs will  
1001 remain silent) and would not enhance the population coding as in the heuristic rule.

1002 We further noted that (2.46) contains a factor  $x_i^k$  meaning that the update to  $w_{ji}$  depends on the average  
 1003 input activity from the  $i$ th PN. As this rule makes the biological interpretation more complex (the synaptic update  
 1004 depends on both pre- and post-synaptic activity), we also tested a simplified rule where synaptic changes depend  
 1005 only on the average KC activity:

$$1006 \quad \Delta w_{ji} = -\eta [\bar{y}_j - A_0] \quad (2.47)$$

1007 This simplification did not affect memory performance or the tuned distribution of weights (Fig. 2.8A-C), but it  
 1008 improved the robustness of the model to novel odour environments (Fig. 2.8D). This improvement in the model  
 1009 robustness might be because including the extra factor  $x_i^k$  in the learning rule caused the model to be overfitted  
 1010 to the tuning environment. Therefore, we used (2.47) for the results presented in the main figures, as it is simpler  
 1011 and produces better performance, despite not being formally derived from an error function. As with (2.46), this  
 1012 update rule has a fixed point  $\bar{y}_j = A_0$ .

1013 Because KC claws sampled PNs with replacement, some KCs had ‘duplicate’ inputs from the same PN. For these  
 1014 weights, we initialised  $w_{ji}$  at double the normal level before beginning optimization. When plotting the distribution  
 1015 of values of  $w$  in Fig. 4D, we split these ‘duplicate’ weights into two weights of half the strength, on the basis that  
 1016 we were comparing our  $w$  values to amplitudes of spontaneous EPSPs from (Turner, Bazhenov, and Laurent 2008),  
 1017 and in a KC with two claws connected to different boutons of the same PN, spontaneous EPSPs from the two claws  
 1018 would likely occur at different times and thus be counted separately.

### 1019 Tuning KC input inhibitory weights to equalize average KC activity

1020 In this model, we model each KC as adjusting its individual input inhibitory synaptic weights from APL, to match  
 1021 its average activity level  $\bar{y}_j$  to a certain desired level  $A_0$ . We minimize the error function in (2.39) by adjusting  $\alpha_j$   
 1022 instead of  $w_{ji}$ :

$$1023 \quad \Delta \alpha_j = -\eta \frac{d\epsilon}{d\alpha_j} \quad (2.48)$$

$$1024 \quad \frac{d\epsilon}{d\alpha_j} = [\bar{y}_j - A_0] \frac{1}{K} \sum_{k=1}^K \frac{dy_j^k}{d\alpha_j} \quad (2.49)$$

1026 Differentiating  $y_j^k$  with respect to  $\alpha_j$  yields

$$1027 \quad \frac{dy_j^k}{d\alpha_j} = H(y_j^k) \left[ -\sum_{j=1}^M \sum_{i=1}^{24} w_{ji} x_i^k \right] \quad (2.50)$$

1028 Plugging (2.50) in (2.49) gives,

$$1029 \quad \begin{aligned} \frac{d\epsilon}{d\alpha_j} &= [\bar{y}_j - A_0] \frac{1}{K} \sum_{k=1}^K H(y_j^k) \left[ -\sum_{j=1}^M \sum_{i=1}^{24} w_{ji} x_i^k \right] \\ &= [\bar{y}_j - A_0] \left\langle H(y_j^k) \left( -\sum_{j=1}^M \sum_{i=1}^{24} w_{ji} x_i^k \right) \right\rangle_K \end{aligned} \quad (2.51)$$

1030 Therefore,

$$1031 \quad \Delta\alpha_j = \eta [\bar{y}_j - A_0] \left\langle H(y_j^k) \left( \sum_{j=1}^M \sum_{i=1}^{24} w_{ji} x_i^k \right) \right\rangle_K \quad (2.52)$$

1032 Similar to the previous section, we assume that weight changes for silent neurons happen in the same way as  
1033 for active neurons:

$$1034 \quad \Delta\alpha_j = \eta [\bar{y}_j - A_0] \left\langle (1 - H(y_j^k)) \left( \sum_{j=1}^M \sum_{i=1}^{24} w_{ji} x_i^k \right) \right\rangle_K \quad (2.53)$$

1035 .  
1036 Adding (2.52) and (2.53) we obtain the inhibitory plasticity rule allowing KCs to achieve equal average activity:

$$1037 \quad \Delta\alpha_j = \eta [\bar{y}_j - A_0] \left\langle \sum_{j=1}^M \sum_{i=1}^{24} w_{ji} x_i^k \right\rangle_K \quad (2.54)$$

1038 Given that  $\left\langle \sum_j \sum_i w_{ji} x_i^k \right\rangle_K$  is a constant as  $w_{ji}$  is not updated in this model, this term can be subsumed into  
1039 the learning rate, so this equation reduces to:

$$1040 \quad \Delta\alpha_j = \eta [\bar{y}_j - A_0] \quad (2.55)$$

1041 Besides the homeostatic tuning of the APL inhibitory feedback values, these individual values of  $\alpha_j$  also have  
1042 to satisfy the sparsity constraint in (2.5). Therefore, the learning rule for these inhibitory weights requires simulta-  
1043 neously optimizing both error functions, (2.5) and (2.39). Thus combining (2.55) and the derivative of the sparsity  
1044 constraint (CL=10%) with respect to each value of  $\alpha_j$ ,

$$1045 \quad \Delta\alpha_j = \eta_1 [\bar{y}_j - A_0] - \eta_2 \frac{d\epsilon_{CL}}{d\alpha_j} \quad (2.56)$$

$$1046 \quad \Delta\alpha_j = \eta_1 [\bar{y}_j - A_0] - \eta_2 [CL - CL_{target}] \frac{dCL}{d\alpha_j} \quad (2.57)$$

1047 where

$$1048 \quad \frac{dCL}{d\alpha_j} = -\frac{1}{MK} \sum_{k=1}^K \left[ S'(y_j^k) H(y_j^k) \sum_{j=1}^M \sum_{i=1}^{24} w_{ji} x_i^k \right] \quad (2.58)$$

1049 Combining (2.57) with (2.58),

$$1050 \quad \Delta\alpha_j = \eta_1 [\bar{y}_j - A_0] + \eta_2 [CL - CL_{target}] \left\langle S'(y_j^k) H(y_j^k) \sum_{j=1}^M \sum_{i=1}^{24} w_{ji} x_i^k \right\rangle_k \quad (2.59)$$

1051 We tested re-parameterizing  $\alpha_j$  into  $C_\alpha \alpha_j$  where  $C_\alpha$  is tuned across all KCs to adjust coding level while  $\alpha_j$  is  
1052 tuned individually to equalize KC activity levels, but this had no effect on memory performance, so we kept the  
1053 simpler model formulation.



### 1054 **Tuning KC spiking thresholds to equalize average KC activity**

1055 In this compensatory technique, we tune individual KCs' spiking thresholds  $\theta_j$  to achieve equal average activity  
 1056 across the KC population. Starting with arbitrary initial values, each KC adjusts its spiking threshold so its  
 1057 average activity across  $K$  odours reaches a target level,  $A_0$ , by minimizing the error in average activity as in (2.39)  
 1058 by gradient optimization:

$$\begin{aligned} \Delta\theta_j &= -\eta \frac{d\epsilon}{d\theta_j} \\ \frac{d\epsilon}{d\theta_j} &= [\bar{y}_j - A_0] \frac{1}{K} \sum_{k=1}^K \frac{dy_j^k}{d\theta_j} \end{aligned} \quad (2.60)$$

1060 Differentiating  $y_j^k$ , the expression in (2.3), with respect to  $\theta_j$  yields

$$\frac{dy_j^k}{d\theta_j} = H(y_j^k) [-C_\theta] \quad (2.61)$$

1062 Plugging (2.61) in (2.60) gives,

$$\begin{aligned} \frac{d\epsilon}{d\theta_j} &= [\bar{y}_j - A_0] \frac{1}{K} \sum_{k=1}^K H(y_j^k) [-C_\theta] \\ &= -[\bar{y}_j - A_0] C_\theta \langle H(y_j^k) \rangle_k \end{aligned} \quad (2.62)$$

1064 Therefore,

$$\Delta\theta_j = \eta [\bar{y}_j - A_0] C_\theta \langle H(y_j^k) \rangle_k \quad (2.63)$$

1066 Similar to (2.45), we assume that spiking thresholds are updated for silent KCs as well:

$$\Delta\theta_j = \eta [\bar{y}_j - A_0] C_\theta \langle (1 - H(y_j^k)) \rangle_k \quad (2.64)$$

1068 Adding (2.63) and (2.64) we obtain the spiking thresholds plasticity rule allowing KCs to achieve equal average  
 1069 activity:

$$\Delta\theta_j = \eta C_\theta [\bar{y}_j - A_0] \quad (2.65)$$

### 1071 **Tuning spiking thresholds to equalize KCs response probabilities**

1072 We tested an alternative strategy to tune  $\theta$  suggested in (Kennedy 2019): to equalize not  $\bar{y}_j$  but rather the average  
 1073 response probability of each KC across  $K$  odours without inhibition,  $P_j$ , i.e.:

$$P_j = \frac{1}{K} \sum_{k=1}^K H(y_j^k |_{\alpha=0}) \quad (2.66)$$

1075 As in (2.5), we set this target response probability,  $P_j^{target}|_{\alpha_j=0}$ , to 0.2 to match experimental findings that  
 1076 blocking inhibition approximately doubles response probability (Lin et al. 2014). We minimized the error function:

$$\epsilon = \frac{1}{2} [P_j - P_j^{target}|_{\alpha_j=0}]^2 \quad (2.67)$$

1078 by adjusting  $\theta_j$  by gradient optimization:

$$1079 \quad \Delta\theta_j = -\eta \frac{d\epsilon}{d\theta_j} \quad (2.68)$$

$$\frac{d\epsilon}{d\theta_j} = [P_j - P_j^{target}|_{\alpha_j=0}] \frac{dP_j}{d\theta_j}$$

1080 To differentiate  $P_j$ , as in (2.13), we approximated the discontinuous Heaviside function with a sigmoid:

$$1081 \quad \frac{dP_j}{d\theta_j} = \frac{1}{K} \sum_{k=1}^K \frac{dS(y_j^k |_{\alpha=0})}{d\theta_j} \quad (2.69)$$

$$\frac{dS(y_j^k |_{\alpha=0})}{d\theta_j} = S'(y_j^k |_{\alpha=0}) \frac{dy_j^k |_{\alpha=0}}{d\theta_j}$$

1082 Recalling the formula of  $y_j^k$  in (2.3), it follows

$$1083 \quad \frac{dy_j^k |_{\alpha=0}}{d\theta_j} = -C_\theta H(y_j^k) \quad (2.70)$$

1084 Combining (2.70) with (2.69), and plugging in (2.68),

$$1085 \quad \frac{d\epsilon}{d\theta_j} = - [P_j - P_j^{target}|_{\alpha=0}] C_\theta \langle S'(y_j^k |_{\alpha=0}) H(y_j^k |_{\alpha=0}) \rangle_K \quad (2.71)$$

1086 Thus,  $\theta_j$  values are updated by,

$$1087 \quad \Delta\theta_j = \eta C_\theta [P_j - P_j^{target}|_{\alpha=0}] \langle S'(y_j^k |_{\alpha=0}) H(y_j^k |_{\alpha=0}) \rangle_K \quad (2.72)$$

1088 As in (2.45), (2.64) and (2.53), we can write a symmetric rule for silent KCs:

$$1089 \quad \Delta\theta_j = \eta C_\theta [P_j - P_j^{target}|_{\alpha=0}] \langle S'(y_j^k |_{\alpha=0}) (1 - H(y_j^k |_{\alpha=0})) \rangle_K \quad (2.73)$$

1090 Adding (2.73) and (2.72) leads to an activity-dependent update rule for  $\theta_j$ , given all the incoming input odours:

$$1091 \quad \Delta\theta_j = \eta C_\theta [P_j - P_j^{target}|_{\alpha=0}] \langle S'(y_j^k |_{\alpha=0}) \rangle_K \quad (2.74)$$

1092 In this model, the sparsity constraint  $CL_{target|\alpha=0} = 0.2$  is satisfied by  $P_j^{target}|_{\alpha_j=0} = 0.2$ , because coding level  
1093 equals the average of response probabilities across KCs:

$$1094 \quad \begin{aligned} CL &= \frac{1}{K} \sum_{k=1}^K \left( \frac{1}{M} \sum_{j=1}^M H(y_j^k) \right) \\ &= \frac{1}{M} \sum_{j=1}^M \left( \frac{1}{K} \sum_{k=1}^K H(y_j^k) \right) \\ &= \langle P_j \rangle_j. \end{aligned} \quad (2.75)$$

## 1095 Optimization of the multiple objective functions

1096 As noted above, homeostatic tuning of  $w_{ji}$ ,  $\theta_j$ , or  $\alpha_j$  needs to happen while maintaining the sparsity constraints,  
1097 (2.5) and (2.6). (It is important to note that the homeostatic update rules are meant to represent a biological process  
1098 while the sparsity constraints merely fit our model to experimental data and stand in for unknown processes that  
1099 lead to a coding level of 0.1.) Since these activity-equalizing tunings both depend on and change the network's  
1100 sparsity level, we used a sequential optimization approach to optimize each objective function,  $O_i$ , at a time. For  
1101 each  $i$ , we find the optimal parameters  $\{P_i\}$  minimizing an objective  $O_i$ , using the current estimates of the other

1102 parameters  $\{P_j\}$  from all the other objectives,  $\{O_j\}$  where  $j \neq i$ . The algorithm iterates for all  $i$  to minimise each  
 1103 of the objective functions, until it reaches a minimum where the errors from all of the objective functions fall below  
 1104 a certain tolerance,  $\tau_O$ .

1105 Given an initial estimate for  $C_\theta$ ,  $\alpha$ ,  $\theta_j$  and  $w_{ji}$ , the algorithm goes as follows:

---

**Algorithm 1** Tuning of KCs parameters to equalize activity while constraining coding level

---

**Result:**  $C_\theta$ ,  $\alpha$ , parameters to be tuned for activity equalization [ $w_{ji}$  or  $\theta_j$ ]

0: Initialize  $C_\theta=1$ ,  $\alpha=0$ ,  $\epsilon_1 = \epsilon_2=1$ ,  $\epsilon_3 = \bar{1}$ ,  $\tau_1=0.2$   $\tau_2=0.01$ ,  $\tau_3 = 0.06\mathbf{A}_0$

0: Initialize tuned parameter for activity equalization  $\{w_{ji}$  or  $\theta_j\} \in U[0,1]$  **while** any in  $[\epsilon_1, \epsilon_2, \epsilon_3] > [\tau_1, \tau_2, \tau_3]$  **do**  
 0:

1. Using the current values for  $\theta_j$  and  $w_{ji}$ , update  $C_\theta$  using (2.15)

0: 2. Using the value of  $C_\theta$  from step (1) and current values for  $w_{ji}$ , and  $\theta_j$ , update  $\alpha$  using (2.19)

0: 3. Using  $C_\theta$  and  $\alpha$  from (1) and (2) respectively, update  $w_{ji}$  using (2.44) or  $\theta_j$  using (2.65)

0: 4. Re-calculate the errors for the three objectives, (2.5), (2.6) and (2.39):

$$\{\epsilon_1 = | \frac{CL|_{\alpha=0}}{CL} - 2 |, \epsilon_2 = | CL - 0.1 |, \epsilon_3 = | \bar{y}_j - \mathbf{A}_0 |\}$$

0:

=0

---

1106 In our implementation we initialize the parameters to be tuned for activity equalization ( $w_{ji}$ ,  $\theta_j$  or  $\alpha_j$ ) from a  
 1107 uniform random distribution  $U = [0, 1]$  (the non-tuned parameters follow the distributions in Fig. 2). In addition,  
 1108 we set the error for the first and second sparsity constraint, (2.5) and (2.6), to be  $\tau_1 = | \frac{CL|_{\alpha=0}}{CL} - 2 | = 0.2$ , while  $\tau_2 =$   
 1109  $| CL - 0.1 | = 0.01$  respectively. This means allowing the coding level without and with the APL feedback to fall  
 1110 within  $[1.8CL \leq CL|_{\alpha=0} \leq 2.2CL]$ , and  $[0.09 \leq CL \leq 0.11]$  respectively. For the activity equalization objective,  
 1111 the error  $\epsilon_3$  is a column vector of size  $M$ , of the differences between the target average activity value  $A_0$ , and the  
 1112 current average activity for each KC,  $\bar{y}_j$ . This objective function is satisfied when all the values in the vector  $\epsilon_3$   
 1113 are less than 6% of the target activity.

1114 Note that in the inhibition-tuning model, we tune the same parameter,  $\alpha_j$  (a vector of  $M$  values instead of a  
 1115 constant), to jointly satisfy both the sparsity and the activity-equalization objectives. In this case, step (3) above  
 1116 is removed and step (2) updates  $\alpha_j$  using (2.59).

1117 In the model where we tune  $\theta_j$  to equalize response probability rather than average activity (Fig. 2.9), equalizing  
 1118 response probability without inhibition to 0.2 also solves the coding level constraint ((2.75)). Thus, in this case, the  
 1119 algorithm iterates between 2 steps: (1) update  $\theta_j$  according to (2.74), (2) use these values to update  $\alpha$  according  
 1120 to (2.19), as follows,

1121

---

**Algorithm 2** Tuning of KCs spiking thresholds to equalize response probabilities
 

---

**Result:**  $C_\theta, \alpha, [\theta_j]$  to be tuned for equalizing KCs response probabilities
 

---

```

0: Initialize:  $[C_\theta=1, \alpha=0, \epsilon_1 = \epsilon_2=1, \tau_1=0.2, \tau_2=0.01]$ 
0: Initialize  $[\theta_j] \in U[0,1]$ 
  while any in  $[\epsilon_1, \epsilon_2] > [\tau_1, \tau_2]$  do
0: 1. update  $\theta_j$  using (2.74)
0: 2. Using these new values of  $\theta_j$  in step (1), update  $\alpha$  using (2.19)
0: 3. Re-calculate the errors for the two objectives, (2.67) and (2.6):
       $\{\epsilon_1 = |P_j - P_j^{target}|_{\alpha_j=0}, \epsilon_2 = |CL - 0.1|\}$ 
0: =0
  
```

---

1122 In our optimization pipeline, there is a potential problem in the models where KC activity is equalized by tuning  
 1123  $\alpha_j$  or  $\theta_j$ . In these models  $w_{ji}$  is not tuned, so for values of  $A_0$  that are too high relative to values of  $w_{ji}$ , excitation  
 1124 will be too low to reach the high targets given the constraints  $C_\theta\theta_j > 0$ ,  $CL = 0.1$  and  $CL|_{\alpha=0} = 0.2$ , meaning the  
 1125 algorithm does not converge. (This is not a problem when tuning  $w_{ji}$  because  $w_{ji}$  can go arbitrarily high, whereas  
 1126 thresholds cannot go below zero.) Therefore,  $w_{ji}$  values must be chosen in a sensible range relative to  $A_0$  (keeping  
 1127 in mind that the value of  $A_0$  is arbitrary: see below). Rather than further complicating the objective cost functions  
 1128 by introducing a tunable scaling factor for  $w_{ji}$ , we found that in practice the algorithm converged if  $w_{ji}$  values  
 1129 (starting from a log-normal distribution with  $\mu = -0.0507$ ,  $\sigma = 0.3527$ ) were multiplied by  $\frac{A_0}{CL}$  (where  $CL = 0.1$ ).  
 1130 The target activity  $A_0$  is arbitrary because if parameters can be found to satisfy our model constraints ( $\bar{y}_j = A_0$ ,  
 1131  $CL = 0.1$  and  $CL|_{\alpha=0} = 0.2$ ) for a particular  $A_0 > 0$ , then a solution also exists for  $\bar{y}_j = cA_0$  for any  $c > 0$ , because:

$$\begin{aligned}
 cy_j^k &= c \operatorname{Relu}\left(\sum_{i=1}^{24} w_{ji}x_i^k - \alpha_j \sum_{j=1}^M \sum_{i=1}^{24} w_{ji}x_i^k - C_\theta\theta_j\right) \\
 &= \operatorname{Relu}\left(\sum_{i=1}^{24} (cw_{ji})x_i^k - \alpha_j \sum_{j=1}^M \sum_{i=1}^{24} (cw_{ji})x_i^k - cC_\theta\theta_j\right)
 \end{aligned} \tag{2.76}$$

1133 That is, to scale  $\bar{y}_j$  by a factor  $c$ , one need only scale the parameters  $w_{ji}$  and  $C_\theta$  by  $c$ . In other words, only the  
 1134 relative magnitudes of  $A_0$ ,  $w_{ji}$  and  $C_\theta$ , not the absolute magnitudes, are meaningful. Thus, when comparing the  
 1135 distributions of  $w_{ji}$  and  $\theta$  to their experimental equivalents in Fig. 4, 2.8 and 2.9, we uniformly scaled all  $w_{ji}$  and  
 1136  $\theta$  values to make their mean match the experimental mean, and present  $\alpha$  values in arbitrary units.

### 1137 2.2.5 Robustness analysis

1138 Of the 110 odours tested in (Hallem and Carlson 2006), we took the four chemical classes with the most odours  
 1139 (acids, terpenes, alcohols and esters), so that tuning parameters on a single class would provide a reasonable number  
 1140 of odours (at least 15). Because each class had different numbers of odours, and the memory task is more difficult  
 1141 when more odours need to be classified, we equalized the number of odours in each task by randomly sampling 15  
 1142 odours from those classes that had more than 15 members (terpenes, 16; alcohols, 18; esters, 24), with a different

1143 random sampling for each model instantiation. Because of the small number of odours used for tuning, it was not  
1144 always possible to equalize the activity of every single KC, so we allowed a maximum of 5 KCs to fall outside a  
1145  $\pm 7\%$  bound on average activity.

## 1146 2.2.6 Connectome analysis

1147 KC neurite skeletons and connectivity were downloaded from the hemibrain connectome v. 1.1 (Scheffer et al.  
1148 2020). KCs (excluding those that receive significant non-olfactory input) were selected as neurons whose ‘type’  
1149 field was KCG-m, KCab-c, KCab-m, KCab-s, KCa’b’-ap2 or KCa’b’-m. PN inputs for a KC were identified as  
1150 neurons whose ‘type’ field included adPN, lPN or vPN (NB: some of these, e.g., vPNs, do not project to the  
1151 mushroom body and so were never counted) and that formed more than 2 synapses with the KC (see Fig. 6B).  
1152 KCs with truncated skeletons lacking the dendritic tree were excluded. The posterior boundary of the peduncle  
1153 was the most posterior node in a skeleton annotated as being in the ‘PED(R)’ region of interest (annotations at  
1154 <https://storage.cloud.google.com/hemibrain/v1.1/hemibrain-v1.1-primary-roi-segmentation.tar.gz>). The boundary  
1155 between the calyx and peduncle regions in the hemibrain was defined by innervation by PNs (or lack thereof). The  
1156 distance from this point to each PN-KC synapse along the KC’s neurite skeleton (i.e., not the Euclidean distance)  
1157 was measured as described in (Amin et al. 2020).

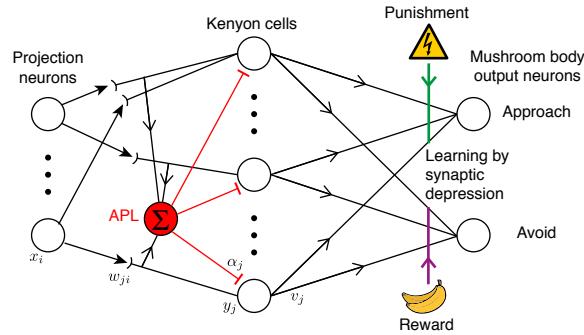
## 1158 2.2.7 Code availability

1159 Modeling and connectome analysis were carried out using custom code written in MATLAB, which is available at  
1160 <https://github.com/aclinlab/CompensatoryVariability>.

## 2.3 Results

### Realistic inter-KC variability impairs memory performance under sparse coding

To study how variability between KCs might affect the fly’s olfactory memory performance, we modelled the mushroom body as a rate-coding neural network (Fig. 2.1).



**Figure 2.1:** Schematic for the mushroom body network model. Projection neurons in the input layer relay the odour responses,  $x_i$ , downstream to the Kenyon cells ( $y_j$ ). Kenyon cells connect randomly to the projection neurons with synaptic weights  $w_{ji}$  and receive global inhibition from the APL neuron with weight  $\alpha_j$ . Learning occurs when dopaminergic neurons (DANs) carrying punishment (reward) signals from the environment depress the synapses ( $v_j$ ) between the active Kenyon cells and the mushroom body output neurons (MBONs) that lead to approach (avoidance) behavior.

To simulate the input activity from PNs, we modeled their activity as a saturating non-linear function of activity of the first-order olfactory receptor neurons (ORNs) (see Eq.(2.3) in the Methods section above; (Olsen, Bhandawat, and Wilson 2010)). We applied this function to the recorded odour responses of 24 different olfactory receptors (Hallem and Carlson 2006) to yield simulated PN activity, as in previous computational studies of fly olfaction (Luo, Axel, and Abbott 2010; Parnas et al. 2013; Krishnamurthy, Hermundstad, and Mora 2017; Kennedy 2019). To simulate variability in PN activity across different encounters with the same odour, we created several ‘trials’ of each odour and added Gaussian noise to PN activity, following the coefficients of variation reported in (Bhandawat et al. 2007). To increase the number of stimuli beyond the 110 recorded odours in (Hallem and Carlson 2006), we generated odour responses in which the activity of each PN was randomly sampled from that PN’s activity across the 110 odours used in (Hallem and Carlson 2006) (results were similar with the ‘real’ 110 odours; see below).

Each of the 2000 KCs in our model received excitatory input from a randomly selected set of  $N$  PNs, each with strength  $w$ . A KC’s response to each odour was the sum of excitatory inputs minus inhibition, minus a spiking threshold  $\theta$ ; if net excitation was below the threshold, the activity was set to zero.

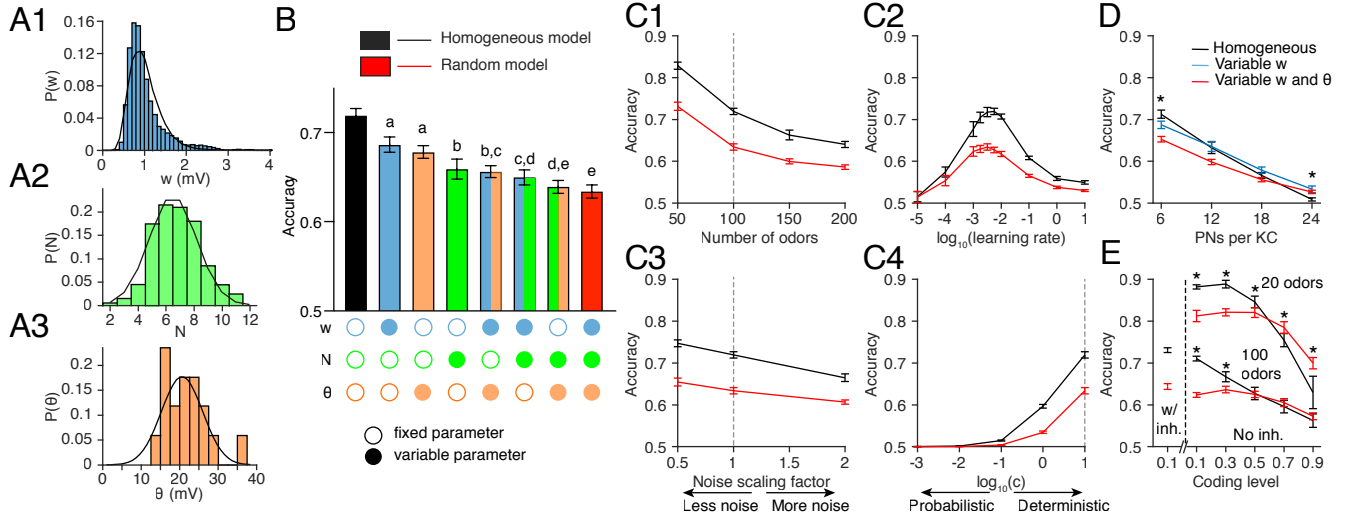
Inhibition came from the feedback interneuron APL (‘Anterior Paired Lateral’), which is excited by and inhibits all KCs (Lin et al. 2014). To avoid simulating the network in time, we simplified the feedback inhibition into pseudo-feedforward inhibition, in which APL’s activity was the sum of all post-synaptic excitation of all KCs (without the KCs’ threshold applied); we based this simplification on the fact that KCs and APL form reciprocal synapses with each other on KC dendrites (i.e., before the KCs’ spike initiation zone), and APL activity is somewhat spatially restricted between KC axons and dendrites (Amin et al. 2020). Thresholds and inhibition were scaled so that on

1184 average 10% of KCs were active for each odour ('coding level' = 0.1).

1185 Learning in flies occurs when KCs (responding to odour) are active at the same time as dopaminergic neurons  
 1186 (DANs, responding to 'reward' or 'punishment'); the coincident activity modifies the output synapse from KCs  
 1187 onto mushroom body output neurons (MBONs) that lead to behavior (e.g., approaching or avoiding an odour).  
 1188 Typically, the output to the 'wrong' behavior is depressed: for example, pairing an odour with electric shock weakens  
 1189 the output synapses from KCs activated by that odour onto MBONs that lead to 'approach' behavior (Aso et al.  
 1190 2014a; Hige et al. 2015; Cohn, Morantte, and Ruta 2015; Handler et al. 2019) (reviewed in (Amin and Lin 2019)).  
 1191 We simulated this plasticity using a simplified architecture with only two MBONs, 'approach' and 'avoid'. The  
 1192 input odours were randomly divided: half were paired with punishment and half with reward. During training, KCs  
 1193 activated by rewarded odours weakened their synapses onto the 'avoid' MBON, while KCs activated by punished  
 1194 odours weakened their synapses onto the 'approach' MBON (depression by exponential decay; see Eq.(2.20)). The  
 1195 fly's behavior then depended probabilistically (via a softmax function; see Eq.(2.21)) on whether the 'avoid' or  
 1196 'approach' MBON's activity was greater, and the model's accuracy in learning was scored as the fraction of correct  
 1197 decisions for unseen noisy variants of the trained odours (i.e., avoiding punished odours and approaching rewarded  
 1198 odours).

1199 To test the effect of realistic inter-KC variability on this model, we introduced variability step-by-step. We first  
 1200 tested the performance of the model holding constant across all KCs the 3 parameters  $N$  (number of PN inputs per  
 1201 KC),  $w$  (strength of each PN-KC connection) and  $\theta$  (KC spiking threshold). Then we added inter-KC variability  
 1202 step-by-step: first varying only one out of 3 parameters, then 2 out of 3, then all 3 parameters (thus 8 possible  
 1203 models). Inter-KC variability in  $N$ ,  $w$  and  $\theta$  followed experimentally measured distributions (Fig. 2.2A1-3) (Caron  
 1204 et al. 2013; Turner, Bazhenov, and Laurent 2008). Increasing inter-KC variability systematically degraded the  
 1205 model's performance when tested on 100 input odours: the more variable parameters there were, the worse the  
 1206 performance (Fig. 2.2B). This performance trend was the same when these 8 models were trained and tested on  
 1207 the real input odours responses from (Hallem and Carlson 2006) (Fig.2.6A). To test whether this effect is robust to  
 1208 different learning and testing conditions, we tested the two extreme cases while varying the numbers of input odours  
 1209 to be classified, the amount of noise in PN activity, the learning rate at the KC-MBON synapse (the two models  
 1210 might have different optimal learning rates:  $\eta$  in Eq.(2.20)), or the indeterminacy of the fly's decision making ( $c$  in  
 1211 Eq.(2.21)). In every case, the model with all parameters fixed (which we call the 'homogeneous' model) consistently  
 1212 outperformed the model with all parameters variable (which we call the 'random' model) (Fig. 2.2C1-4). These  
 1213 results indicate that biologically realistic variability in KC network parameters impairs the network's ability to  
 1214 classify odours as rewarded vs. punished.

1215 Our conclusion contrasts with earlier results that inter-neuronal variability between mitral cells increases infor-  
 1216 mation content (Padmanabhan and Urban 2010; Padmanabhan and Urban 2014; Tripathy et al. 2013), i.e., that  
 1217 variability is helpful, not harmful. This apparent contradiction can be resolved by noting two differences between  
 1218 our approaches. First, the mitral cell studies provided the same input to every neuron, whereas here, every KC



**Figure 2.2:** Inter-KC variability in  $w$ ,  $N$  and  $\theta$  degrades the model fly’s memory performance. **(A)** Histograms of the experimentally measured distributions for: **(A1)**  $w$  (amplitude of spontaneous excitatory postsynaptic potentials in KCs, mV; data from (Turner, Bazhenov, and Laurent 2008)), **(A2)**  $N$  (number of PN inputs per KC, measured as the number of dendritic ‘claws’; data from (Caron et al. 2013)), **(A3)**  $\theta$  (spiking threshold minus resting potential, mV; data from (Turner, Bazhenov, and Laurent 2008)). The overlaid black curves show log-normal ( $w$ ) and Gaussian ( $N$ ,  $\theta$ ) fits to the data.

**(B)** The model fly’s memory performance (given 100 input odours), varying the parameters step by step. Fixed and variable parameters are shown by empty and filled circles, respectively. The homogeneous model (all parameters fixed,  $N = 6$ ; black) performs the best and the random model (all parameters variable; red) performs the worst. All bars are significantly different from each other unless they share the same letter annotations (a, b, etc.),  $p < 0.05$  by Wilcoxon signed-rank test (for matched models with the same PN-KC connectivity) or Mann-Whitney test (for unmatched models with different PN-KC connectivity, i.e., fixed vs. variable  $N$ ), with Holm-Bonferroni correction for multiple comparisons.  $n = 30$  model instances with different random PN-KC connectivity.

**(C)** The performance trend is consistent over a range of different conditions: **(C1)** number of input odours, **(C2)** the learning rate used to update KC-MBON weights, **(C3)** amount of noise in PN activity (half, the same, or double the noise measured in (Bhandawat et al. 2007)), **(C4)** the indeterminacy in the decision making, quantified by  $\log(c)$ , where  $c$  is the constant in the soft-max function (SI Appendix, Eq. 21). The vertical dotted lines indicate the conditions used in panel B (each condition used the best learning rate).

**(D)** As KCs receive more inputs (thus more similar inputs), inter-KC variability becomes helpful, not harmful, to memory performance, especially when all KCs receive the same inputs ( $N = 24$ ). Blue: KCs vary in excitatory weights ( $w$ ); red: KCs vary in both  $w$  and thresholds ( $\theta$ ). Data for  $N = 6$  equivalent to panel B.  $n=30$ .

**(E)** Inter-KC variability improves performance in dense coding regimes (coding levels 0.7 - 0.9) at classifying 100 odours (a hard task) or 20 odours (easy task). Left of dashed line: equivalent to panel B, for comparison. Right of dashed line: increasing coding levels, in each case without inhibition (because inhibition is constrained to decrease coding level by half, which is impossible if coding level  $> 0.5$ ).  $n=50$ . \*  $p < 0.05$ , Wilcoxon signed-rank test (D) or Mann-Whitney test (E) with Holm-Bonferroni correction for multiple comparisons. Error bars show 95% confidence intervals.



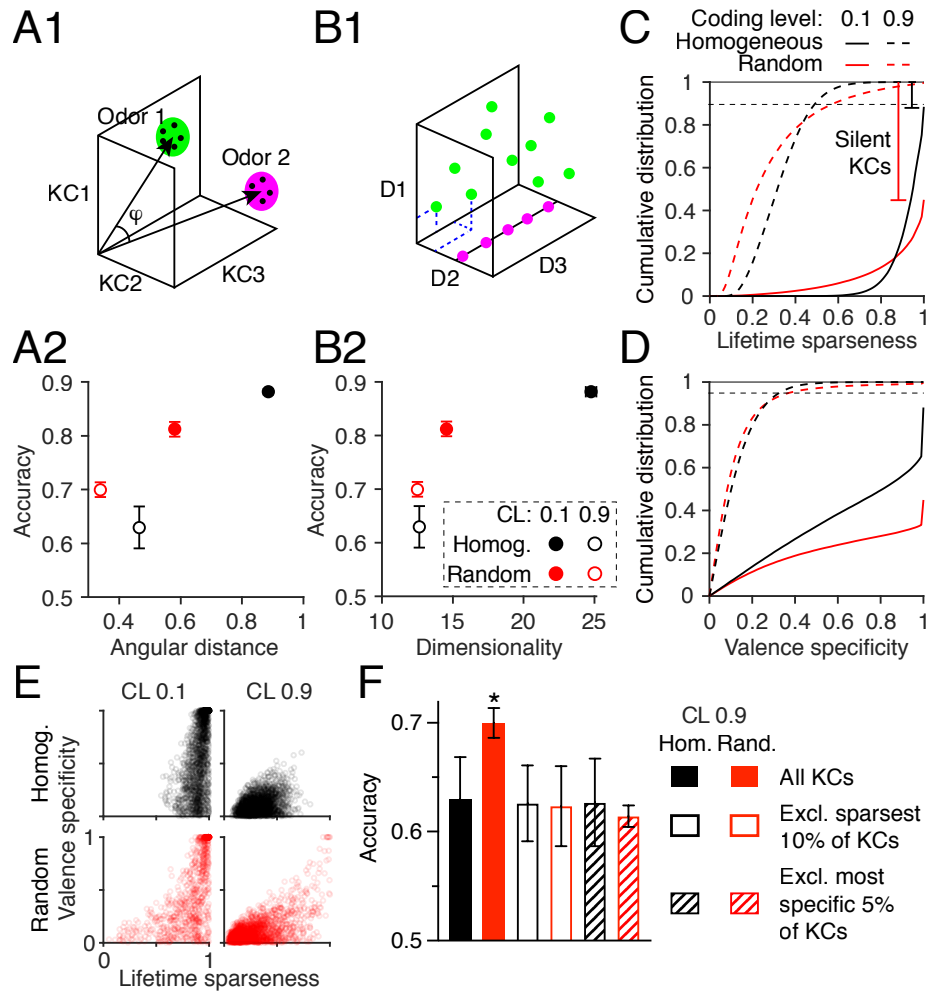
receives different inputs thanks to random PN-KC connectivity. Indeed, when we forced every KC to receive input from the same PNs ( $N = 24$ , i.e. every KC receives input from every PN; Fig. 2.2D), variability between KCs in input weights actually improved performance compared to the homogeneous model (although both models unsurprisingly performed much worse compared to the more realistic  $N = 6$ ). In other words, when all KCs receive the same input, only inter-KC variability allows them to have different odour response profiles from each other (Litwin-Kumar et al. 2017), which is required for distinct olfactory memories to be formed at KC output synapses.

Second, unlike in our model, the mitral cell studies did not enforce sparse coding where only a small fraction of cells should respond at any given time. Indeed, under dense coding (coding level = 0.9), while all models unsurprisingly performed worse than under sparse coding (coding level = 0.1), the random model out-performed the homogeneous model. While this difference was only marginal when discriminating 100 odours (possibly due to a floor effect), it was more apparent on an easier task where the network learned to classify 20 odours instead of 100 (Fig. 2.2E). Thus, while sparse coding and diverse PN inputs for each KC greatly improve learned odour classification, these features require homogeneous KCs to fully exploit their advantages, thus making inter-KC variability harmful rather than helpful under sparse coding.

### 2.3.1 Performance depends on KC lifetime sparseness

We next asked what features of KC population odour representations might account for the worse performance of the random model compared to the homogeneous model under sparse coding, but the reverse under dense coding. Learning KC-MBON weights to correctly classify rewarded versus punished odours is equivalent to finding a hyper-plane (in 2000-dimensional space) to separate KC responses to rewarded odours from those to punished odours. Finding a separating hyper-plane might be easier if (a) odours are far apart from each other in KC coding space (measured by angular distance, a scale-insensitive distance metric, Fig. 2.3A1, used in, e.g., (Turner, Bazhenov, and Laurent 2008)), or (b) odour responses occupy more independent dimensions (measured by a metric for dimensionality developed by (Litwin-Kumar et al. 2017); Fig. 2.3B1). Indeed, under sparse coding (coding level = 0.1), the random model had smaller angular distances and lower dimensionality than the homogeneous model (Fig. 2.3A,B, Fig.2.4B,C). However, surprisingly, the same was true at coding level = 0.9, even though in this condition, the random model out-performed the homogeneous model (Fig. 2.2E), suggesting that separation and dimensionality of KC odour responses alone do not explain inter-KC variability’s effect on performance, at least with the learning rule used here (i.e., depression of KC outputs to ‘wrong’ actions by exponential decay).

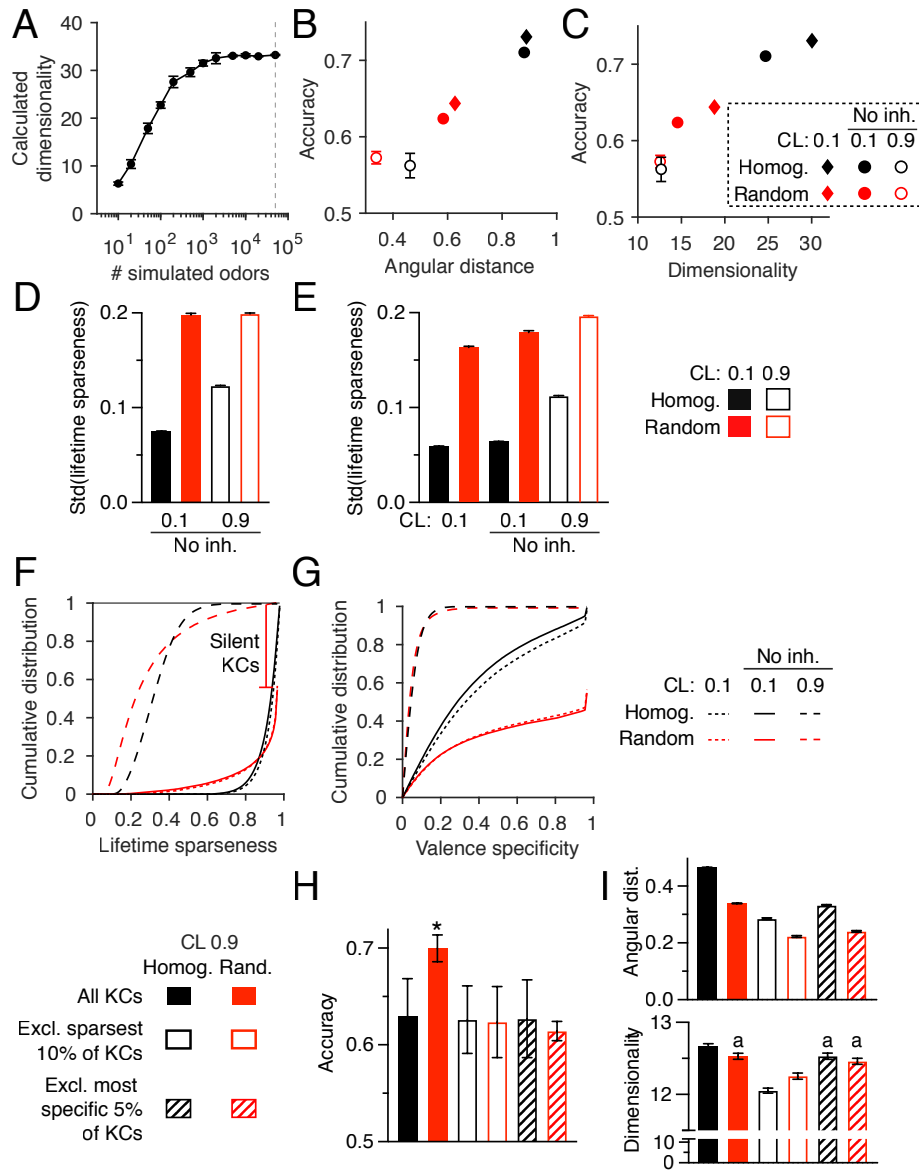
Instead, we hypothesized that inter-KC variability impairs performance under sparse coding because it makes some KCs indiscriminately active but leaves others completely silent, meaning fewer KCs provide useful odour identity information. Sparse coding requires sparseness in two dimensions: population sparseness (each stimulus activates few neurons) and lifetime sparseness (each neuron responds to few stimuli) (Willmore and Tolhurst 2001). While our models enforced population sparseness (coding level = 0.1), they did not enforce any particular lifetime sparseness. In an extreme case, a model could have very consistent population sparseness with a coding level of



**Figure 2.3:** Performance depends on KC lifetime sparseness. **(A1,B1)** Diagrams of angular distance between odours (i.e., between centroids of clusters of noisy trials; A1) and dimensionality of a system with 3 variables (B1). The system with its states scattered throughout 3D space (green) has dimensionality 3 while the system with all states on a single line (magenta) has dimensionality 1. **(A2,B2)** The homogeneous model has higher angular distance and dimensionality than the random model ( $p < 0.05$ , Mann-Whitney test), matching the performance difference when coding level = 0.1, but the opposite trend to performance when coding level = 0.9. **(C-D)** Cumulative distribution function (cdf) of the lifetime sparseness (C) or valence specificity (D) of KCs in the homogeneous (black) and random (red) models, across 50 model instantiations. The gap between 1.0 and the top of the cdf represents silent KCs (lifetime sparseness and specificity undefined). At coding level 0.1, the random model has many more silent KCs, non-sparse KCs, and non-specific KCs than the homogeneous model, but at coding level 0.9, the random model has more KCs with high lifetime sparseness and more KCs with high valence specificity. **(E)** High lifetime sparseness enables high valence specificity, although many sparse KCs have low valence specificity because of random valence assignments (data here from single model instances). **(F)** Removing the sparsest or most valence-specific KCs (corresponding to the dashed horizontal lines in C,D) removes the performance advantage of the random model under dense coding.  $n=50$  network instantiations; \*  $p < 0.05$ , Mann-Whitney test; error bars, 95% confidence interval (horizontal error bars in A2,B2 are smaller than the symbols). These results are from the 20-odour task in Fig. 2.2E.

1253 0.1 for all odours, simply by having the same 10% of cells responding equally to every odour and the other 90%  
 1254 being completely silent. In this case, no cells would provide any useful information about odour identity. We asked  
 1255 whether a less extreme version of this problem could explain the relative performance of our models.

1256 To test this, we quantified the specificity of KCs, both across all odours and for rewarded vs. punished odours.  
 1257 To quantify specificity across odours, we used lifetime sparseness, a metric that is 1 when a cell fires to one stimulus  
 1258 and no other stimuli, vs. 0 when it fires equally to all stimuli. A cell that fires to no stimuli has an undefined



**Figure 2.4:** Additional metrics supporting Figure 3. **(A)** Dimensionality can be estimated numerically using (2.23) given sufficient simulated inputs (dashed line = 50,000, the number used here). Calculations here on the homogeneous model, coding level = 0.1, with inhibition. **(B-C)** As in Fig. 3A,B, except models trained to discriminate 100 odours instead of 20 odours. The homogeneous model has higher angular distance and dimensionality than the random model ( $p < 0.05$ , Mann-Whitney test), matching the performance difference when coding level = 0.1, but the opposite trend to performance when coding level = 0.9. **(D-E)** The random model has greater standard deviation of lifetime sparseness across KCs, compared to the homogeneous model, in all conditions tested (coding level 0.1 or 0.9; with or without inhibition) using 20 odours in (D) and 100 odours in (E). Note: Inhibition was omitted for comparing coding level 0.1 vs. 0.9 because our model was constrained to have the coding level without inhibition be double the coding level with inhibition, which is impossible when the coding level with inhibition is 0.9. The results in (D) are from the 20-odour task in Fig. 2.2E. **(F-G)** As in Fig. 3C,D, except models trained to discriminate 100 odours instead of 20 odours. Cumulative distribution function (cdf) of the lifetime sparseness (C) or valence specificity (D) of KCs in the homogeneous (black) and random (red) models, across 50 model instantiations. The gap between 1.0 and the top of the cdf represents silent KCs (lifetime sparseness undefined). At coding level 0.1, the random model has many more silent KCs, non-sparse KCs, and non-specific KCs than the homogeneous model, but at coding level 0.9, the random model has more KCs with high lifetime sparseness and more KCs with high valence specificity. **(H)** Reproduced from Fig. 3F for comparison: Removing the sparsest or most valence-specific KCs removes the performance advantage of the random model under dense coding. \*  $p < 0.05$ , Wilcoxon signed-rank test (see Table S1). **(I)** Removing the sparsest or most valence-specific KCs generally reduces angular distance and dimensionality, but not in a way that matches the effect on performance shown in (H). Conditions are significantly different (Mann-Whitney test) unless they share a letter annotation.  $n=50$  network instantiations; error bars, 95% confidence interval (where error bars cannot be seen, they are smaller than the symbols); performance data in B,C from the 100-odour task in Fig. 2E.

1259 sparseness (see Methods section Eq.(2.24) ). The homogeneous model had fairly consistent lifetime sparseness  
 1260 values, with almost 80% of KCs having a lifetime sparseness between  $\sim 0.85$  and 1. In contrast, the random model  
 1261 had KCs with much more variable lifetime sparseness, with a long tail of KCs with low sparseness (below 0.7) and  
 1262 more than 50% of KCs having undefined sparseness (i.e., completely silent). (These figures are when considering 20  
 1263 odours; when considering 100 odours, there are fewer silent KCs but the overall pattern is the same: Fig. 2.4F,G.)  
 1264 The contrasting distributions of lifetime sparseness can be seen in the cumulative distribution functions (cdfs) of  
 1265 lifetime sparseness in Fig. 2.3C and Fig. 2.4F, in how the steep curve of the homogeneous model and the shallow  
 1266 curve of the random model cross each other. This result can also be seen in the larger standard deviation of lifetime  
 1267 sparseness across KCs in the random model (Fig. 2.4 D,E). The silent KCs can be seen as the fraction of missing  
 1268 KCs needed for the cdf curves to reach 1; the random model has many more silent KCs than the homogeneous  
 1269 model.

1270 To quantify KCs' specificity for rewarded vs. punished odours, we defined 'valence specificity' for each KC as  
 1271 the absolute value of the difference between total activity for all rewarded vs. all punished odours, divided by total  
 1272 activity for all odours. Again, under sparse coding, the homogeneous model had more KCs with higher valence  
 1273 specificity than the random model (Fig. 2.3D). Given random valence assignments, high lifetime sparseness does  
 1274 not guarantee high valence specificity, but does make it more probable (the two measures are correlated: Fig. 2.3E),  
 1275 for the same reason that flipping a coin 5 times is more likely to give all heads than flipping a coin 50 times: a KC  
 1276 active for only a few odours is more likely to be active only for rewarded (or punished) odours, compared to a KC  
 1277 active for many odours.

1278 Under dense coding, KCs also have more variable lifetime sparseness in the random model (dashed lines, Fig.  
 1279 2.3C; Fig. 2.4F). However, here, the inter-KC variability is helpful rather than harmful: whereas KCs in the  
 1280 homogeneous model have uniformly low lifetime sparseness (and thus are uniformly useless for odour discrimination),  
 1281 in the random model, the inter-KC variability allows a small minority of KC to have relatively high lifetime  
 1282 sparseness and valence specificity (though still worse than under sparse coding; Fig. 2.3C-E). To test whether this  
 1283 minority of relatively specific KCs explains the better performance of the random model under dense coding, we  
 1284 removed the 10% of KCs with the highest lifetime sparseness or the 5% of KCs with the highest valence specificity  
 1285 (fractions correspond to the approximate parts of the cdfs where the random model had higher values: dashed  
 1286 horizontal lines on 2.3C,D), and replaced them with useless KCs (either silent or responding equally to all odours,  
 1287 to preserve the 0.9 coding level). Indeed, in these cases, the random model no longer outperformed the homogeneous  
 1288 model (Fig. 2.3F). However, these changes did not correspond to the effects of removing the sparsest or most specific  
 1289 KCs on angular distance or dimensionality (Fig. 2.4I), again indicating that angular distance and dimensionality  
 1290 do not always correspond to performance in our model.

1291 Together, these results indicate that under sparse (but not dense) coding, introducing realistic inter-KC variabil-  
 1292 ity in  $w$ ,  $N$ , and  $\theta$  worsens the performance of the network by making KCs' odour response profiles less consistently  
 1293 sparse and thus less specific to rewarded/punished odours. Because the real mushroom body uses sparse coding,

1294 we focus the rest of our analysis on the sparse coding condition (coding level = 0.1).

### 1295 **2.3.2 Compensatory tuning of KC parameters rescues memory performance**

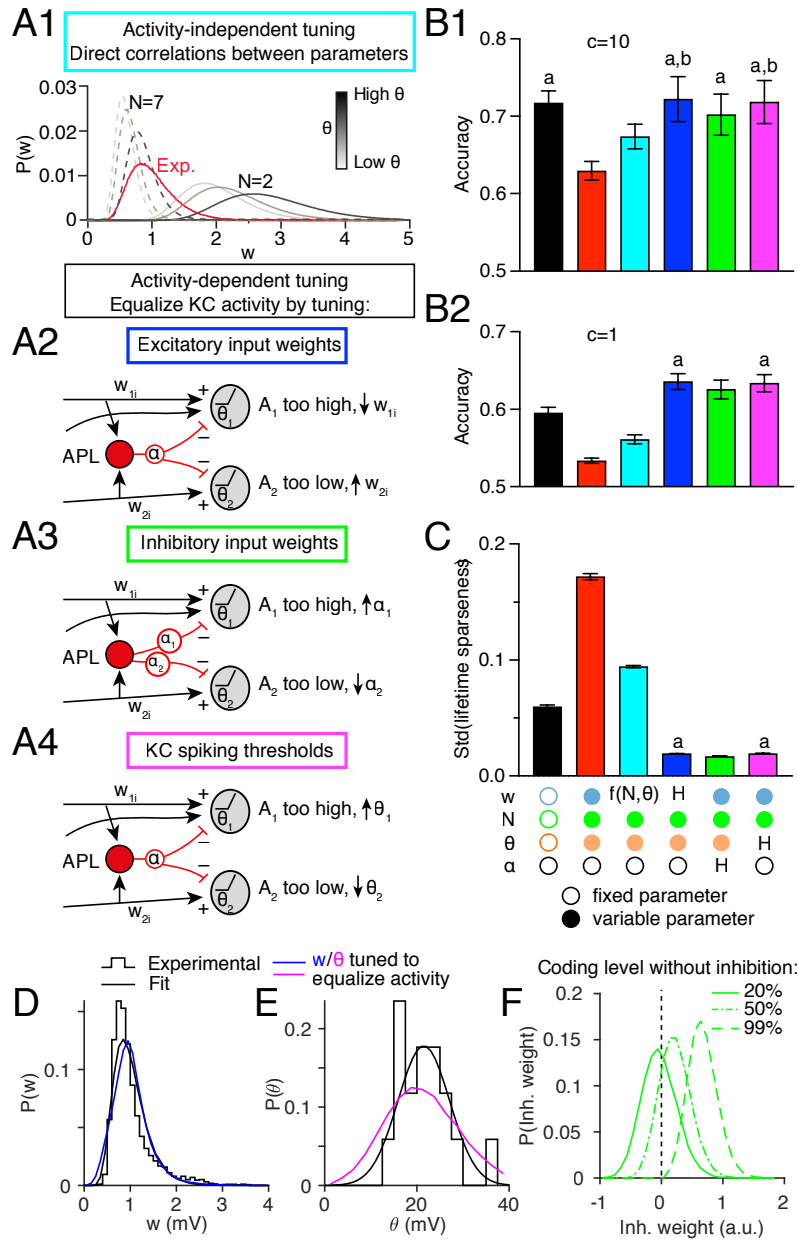
1296 Because the central problem for memory performance in the random model was inter-KC variability in activity, we  
1297 hypothesized that performance could be rescued in models where KCs could achieve roughly equal activity across  
1298 the population, while still preserving experimentally realistic variability in spiking thresholds and number/strength  
1299 of excitatory inputs.

#### 1300 **Activity-independent tuning of excitatory input weights**

1301 First, we tested a model that equalizes KC activity indirectly, by making parameters compensate for each other in  
1302 an activity-independent way. In particular, we modeled KCs as adjusting input synaptic weights ( $w$ ) to compensate  
1303 for variability in spiking threshold ( $\theta$ ) and number of PN inputs ( $N$ ). Thus, an individual KC with low  $\theta$  or high  
1304  $N$  would have low  $w$ , while a KC with high  $\theta$  or low  $N$  would have high  $w$ . We simulated these correlations  
1305 ( $w \propto \sqrt{\theta}$ ;  $w \propto 1/\sqrt{N}$ ) constrained by experimental data. To do this, we sampled  $N$  and  $\theta$  from the distributions in  
1306 Fig. 2.2A, and sampled  $w$  from a posterior compensatory distribution,  $P(w | N, \theta)$ , whose overall shape across all  
1307 KCs was constrained to be the same as the experimental  $P(w)$  in Fig. 2.2A1 but which was composed of multiple  
1308 distributions of  $P(w)$  for different values of  $N$  and  $\theta$ . For example, a KC with a relatively high  $N = 7$  would  
1309 sample its weights from a  $P(w)$  shifted to the left (lower  $w$ ) (Fig. 2.5A1, dashed lines), while a KC with a relatively  
1310 low  $N = 2$  would sample its weights from a  $P(w)$  shifted to the right (higher  $w$ ) (Fig. 2.5A1, solid lines). The  
1311 same would be true for different values of  $\theta$  (Fig. 2.5A1, different shadings). We fitted these component  $P(w)$   
1312 curves so that with experimentally observed distributions of  $N$  and  $\theta$ , the sum of the components would produce  
1313 the experimentally observed distribution of  $w$  across all KCs (see Methods section). (Note that this algorithm  
1314 is not meant to describe an actual biological mechanism, merely to create correlations between  $w$  vs.  $N$  and  $\theta$   
1315 while constraining the parameters to experimentally realistic distributions. Biologically, such correlations could  
1316 arise through several mechanisms; see Discussion.) This compensatory mechanism rescued the fly's performance,  
1317 producing significantly higher accuracy at classifying odours than the random model (Fig. 2.5B, Fig. 2.6B, cyan  
1318 bars), likely resulting from the reduced variability in KC lifetime sparseness (Fig. 2.5C). (Note however that this  
1319 model did not perform quite as well as the homogeneous model.)

#### 1320 **Activity-dependent tuning of KC parameters**

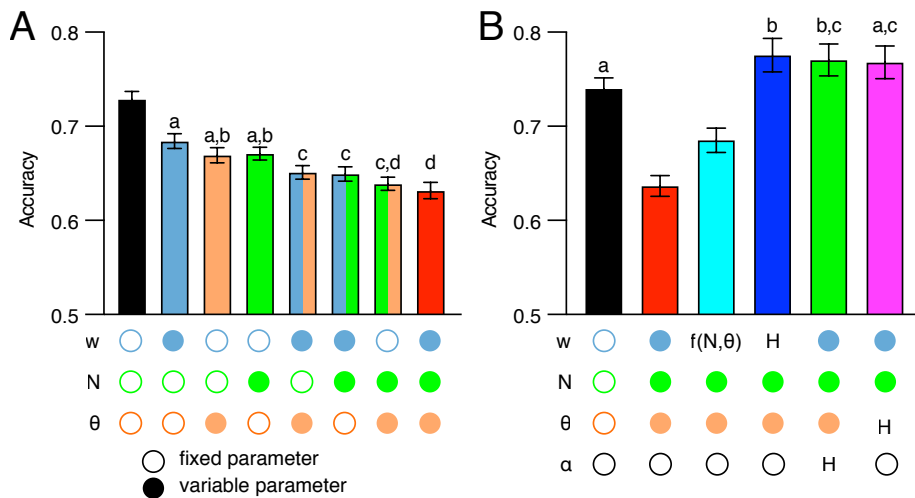
1321 We next tested compensatory mechanisms based on activity rather than explicit correlations between network  
1322 parameters. Here, each KC has the same desired average activity level across all odours,  $A_0$  (with a tolerance of  
1323  $\pm 6\%$ ). We tested three models, each of which equalized average KC activity  $A_0$  by tuning a different parameter:  
1324 input excitatory weights ( $w$ ), inhibitory weights ( $\alpha$ ), or spiking thresholds ( $\theta$ ). The non-tuned parameters followed  
1325 the distributions in Fig. 2.2A (inhibitory weights were constant when non-tuned), while individual KCs adjusted the



**Figure 2.5:** Compensation in network parameters rescues memory performance. **(A)** Schematics of different compensation methods. **(A1)** Activity-independent compensation. Lognormal fit of experimental distribution of the synaptic weights (Exp., red), and its component distributions for different  $N$  and  $\theta$ , for high  $N = 7$  (dashed) or low  $N = 2$  (solid). Shadings of gray indicate different values of  $\theta$ . **(A2-4)** Mechanisms for activity-dependent homeostatic compensation. Overly active KCs weaken excitatory input weights ( $w_{ji}$ , A2), strengthen inhibitory input weights ( $\alpha_j$ , A3), or raise spiking thresholds ( $\theta_j$ , A4). Inactive KCs do the reverse. **(B1)** Compensation rescues performance, alleviating the defect caused by inter-KC variability in the random model (red) compared to the homogeneous model (black), whether compensation occurs by setting  $w$  according to  $N$  and  $\theta$  (cyan; A1), using activity-dependent homeostatic compensation to adjust excitatory weights (blue; A2), inhibitory weights (green; A3) or spiking thresholds (magenta; A4). **(B2)** Differences between models are more apparent when the task is more stochastic decision-making ( $c = 1$  instead of  $c = 10$  in the softmax function). **(C)** Compensation reduces variability in KC lifetime sparseness.  $n = 20$  model instances with different random PN-KC connectivity; error bars, 95% confidence interval. All bars are significantly different from each other unless they share the same letter annotations,  $p < 0.05$ , by Wilcoxon signed-rank test (for matched models with the same PN-KC connectivity) or Mann-Whitney test (for unmatched models with different PN-KC connectivity, i.e., fixed vs. variable  $N$ ), with Holm-Bonferroni correction for multiple comparisons. Annotations below bars indicate whether parameters were fixed (empty circle), variable (filled circle), or variable following a compensation rule ('H' for homeostatic tuning,  $f(N, \theta)$  for activity-independent tuning). Results here are for 100 synthetic odours; see Fig. 2.6B for similar results with odours from (Hallem and Carlson 2006). **(D)** KC excitatory input synaptic weights ( $w$ ) after tuning to equalize average activity (blue) follow a similar distribution to experimental data (black, from Fig. 2.2A1) **(E)** KC spiking thresholds ( $\theta$ ) after tuning to equalize average activity (magenta) have wider variability than the experimental distribution (black, from Fig. 2.2A3). **(F)** Tuning KC inhibitory weights ( $\alpha$ ) to equalize average activity requires many inhibitory weights to be negative, unless the coding level without inhibition is as high as 99%.

1326 tuned parameter according to whether their activity was too high or too low. For example, a relatively highly active  
 1327 KC (whether because it has high  $w$  or  $N$ , low  $\theta$ , or simply receives input from highly active PNs) would scale down  
 1328 its excitatory weights (Fig. 2.5A2), scale up its inhibitory weights (Fig. 2.5A3), or scale up its spiking threshold  
 1329 (Fig. 2.5A4). Likewise, a relatively inactive (or indeed totally silent) KC would do the reverse (see Methods for  
 1330 details of the update rules underlying the homeostatic tuning and discussion of variant update rules in Fig. 2.8,  
 1331 2.9).

1332 All three homeostatic models performed as well as the homogeneous model (Fig. 2.5B1, Fig. 2.6B, blue, green,  
 1333 magenta bars), and indeed even out-performed the homogeneous model when decision-making was more stochastic  
 1334 (lower value of  $c$  in the softmax function; Fig. 2.5B2). The more stochastic decision-making makes the task more  
 1335 difficult and thus brings out the enhanced coding by the homeostatic models. Indeed, the variability in KC lifetime  
 1336 sparseness was even lower in the homeostatic models was even higher than in the homogeneous model (Fig. 2.5C).  
 1337 (As average activity and lifetime sparseness are not the same thing, it is notable that tuning to equalize average  
 1338 activity also tended to equalize lifetime sparseness.)



**Figure 2.6:** Similar analyses to Fig. 2 and 4, using the original 110 odour responses from (Hallem and Carlson 2006). **(A)** Inter-KC variability degrades the memory performance. **(B)** Compensation as in Fig. 4 improves memory performance.  $n = 30$  (A) or 20 (B) model instances with different random PN-KC connectivity; error bars, 95% confidence interval. The indeterminacy constant  $c$  from the softmax equation was set to 10. Bars within a graph that do not share the same letter annotation are significantly different,  $p < 0.05$ , Mann-Whitney or Wilcoxon test as in Fig.2.2,2.5.

1339 What distributions of excitatory weights, inhibitory weights, or spiking thresholds emerge after activity-dependent  
 1340 tuning to equalize KC activity? Do they match experimentally observed distributions? Tuning excitatory weights  
 1341 led to a distribution fairly similar to the approximately log-normal experimentally observed distribution of EPSP  
 1342 amplitudes (Fig. 2.5D). Tuning spiking thresholds led to a distribution with greater variance than the experimental  
 1343 distribution, although with a qualitatively similar Gaussian shape (Fig. 2.5E). This larger variance of thresholds  
 1344 suggests that natural variation of  $\theta$  is too small, on its own, to equalize KC activity given the variation in the  
 1345 number/strength of excitatory inputs.

1346 The tuned distribution of inhibitory weights differed even more strongly from experimental results. While

1347 there are no experimental measurements of inhibitory weights, equalizing KC activity by tuning inhibitory weights  
1348 required many of them to be negative (Fig. 2.5G), which is unrealistic, because negative inhibition is actually  
1349 excitation, and there are no reports of GABAergic excitation of KCs (Inada, Tsuchimoto, and Kazama 2017). Our  
1350 model required negative inhibition because of the constraint that inhibition is only strong enough to reduce the  
1351 fraction of active KCs by half (from 20% to 10%, based on results from (Lin et al. 2014)). In other words, 80%  
1352 of the time, KCs are silent even without inhibition, thanks to high thresholds; such responses cannot be increased  
1353 by reducing inhibition unless inhibition becomes negative (i.e., excitatory). Indeed, if we relax the constraint that  
1354 coding level be 0.2 without inhibition, such that sparseness is enforced by inhibition alone (not thresholds), then  
1355 variable inhibition can equalize KC activity without becoming negative (Fig. 2.5F). However, in this case, the  
1356 coding level without inhibition was 99%, which is not observed experimentally (Lin et al. 2014). Even allowing a  
1357 coding level without inhibition of 50%, equalizing KC activity still requires some APL-KC inputs to be negative  
1358 (Fig. 2.5F). Interestingly, these unrealistic models, where sparseness is mainly driven by inhibition rather than high  
1359 thresholds, perform better than the three models shown here (Fig. 2.9A) suggesting that biological constraints may  
1360 limit network performance. Overall, these results suggest that tuning inhibitory weights cannot compensate on its  
1361 own for variability in other KC parameters. More likely, the system optimizes multiple parameters at once (see Fig.  
1362 2.10 and Discussion).

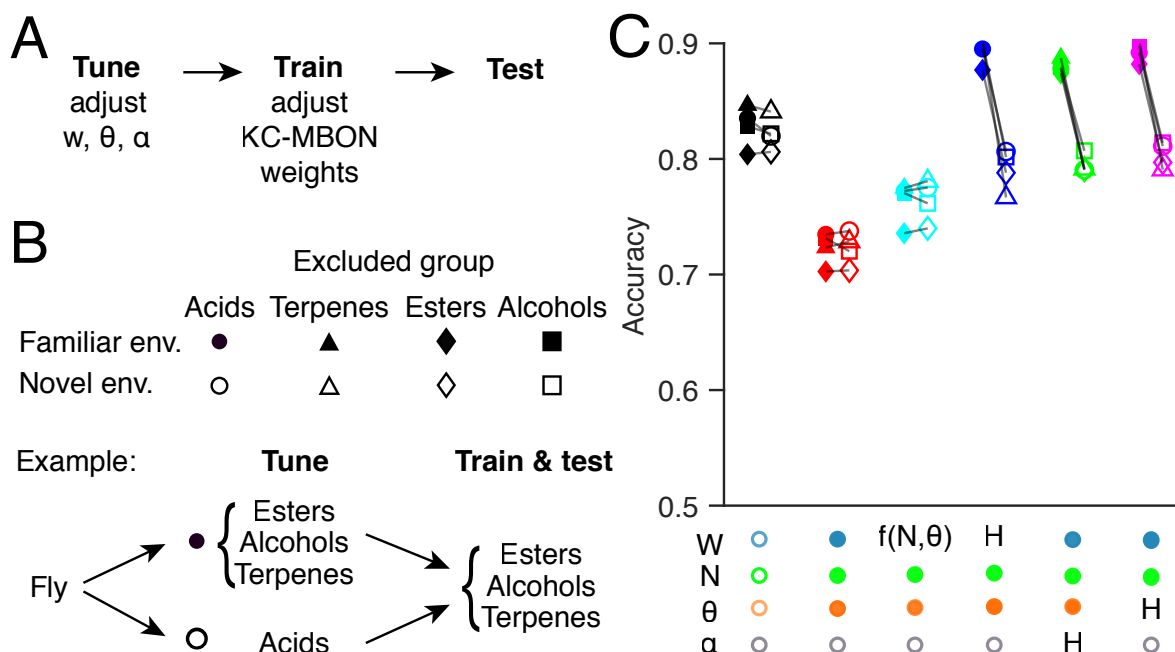
1363 We also tested whether memory performance can be rescued by equalizing not KC average activity, but rather  
1364 KC response probability (equivalent to average activity if KC activity is binarized, i.e., 0 or 1). Equalizing response  
1365 probability (as opposed to average activity) by tuning KC spiking thresholds has been shown to improve separation  
1366 of KC odour representations in a different computational model (Kennedy 2019). However, in our model, this tech-  
1367 nique (tuning thresholds to equalize KC response probability) produced somewhat worse classification performance  
1368 compared to tuning thresholds to equalize KC average activity (Fig. 2.9B,C), though still better than the random  
1369 model (compare Fig. 2.9 to Fig. 2.5).

## 1370 **Robustness of pre-tuned compensations in new environments with novel** 1371 **odours**

1372 Any activity-dependent tuning depends on the model’s context. If a fly tunes its network parameters based on  
1373 experience in one odour context (e.g., smelling only odours of one chemical family), will it still perform well at  
1374 classifying odours in a novel environment with different odours (e.g., odours of a different chemical family)? We  
1375 hypothesized that performance would depend more on tuning context with the activity-dependent compensation  
1376 mechanisms than the activity-independent mechanism.

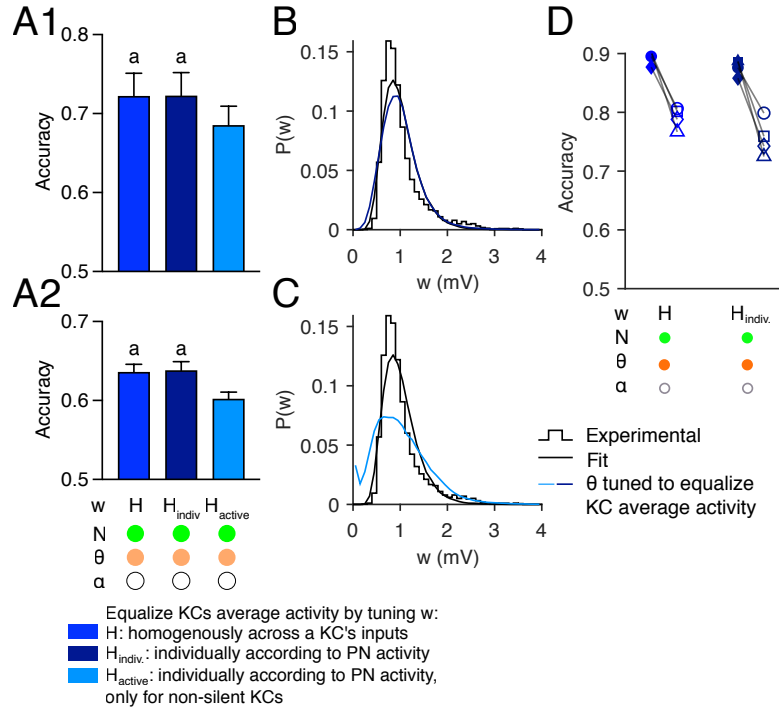
1377 To test this, we tuned the parameters in our models using only a subset of odours from (Hallem and Carlson  
1378 2006), grouped by chemical class, and then trained and tested the models on odour-reward/punishment associations  
1379 using the other odours. We took the four chemical classes that had the most odours in the dataset: acids, terpenes,



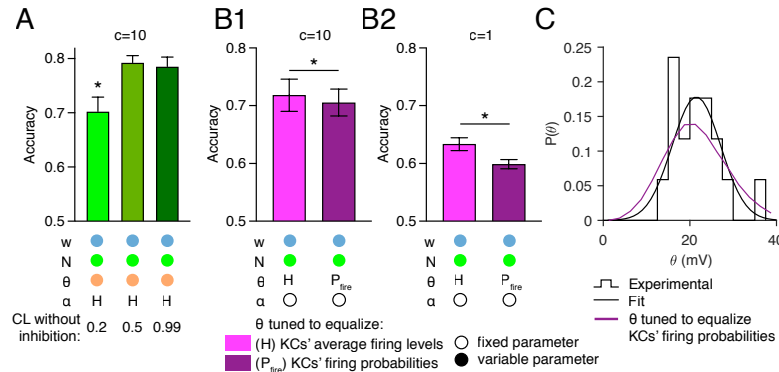


**Figure 2.7:** Robustness of pre-tuned compensations with novel odours. **(A)** For each model fly, network parameters are tuned as in Fig. 2.5, on a subset of odours. At this stage, no rewards or punishments are given, and KC output weights are not modified. Then, the model is trained to classify rewarded and punished odours that are the same as or different from the odours used for tuning. Finally, the model is tested on new noisy variants of the odours used for training. **(B)** Empty symbols (‘novel’ environment): models were tuned on odours from one chemical group ( $G_i$ : acids - circles, terpenes - triangles, esters - diamonds, or alcohols - squares), then trained and tested on odours from the other three groups ( $G_{i \neq j}$ ). Each empty symbol is paired with a matched control (filled symbols) showing how that model would have fared in a ‘familiar’ environment: a model tuned, trained, and tested all on the same three groups of odours as the matched ‘novel’ model was trained and tested on ( $G_{i \neq j}$ ). **(C)** Models with activity-dependent compensation (blue, magenta, green) performed significantly worse in novel environment than familiar environments (matching indicated by connecting lines) ( $p < 0.05$ , Wilcoxon signed-rank test with Holm-Bonferroni correction). In contrast, models with no compensation (black, red), or activity-independent compensation (cyan), performed similarly in novel and familiar environments ( $p > 0.05$  except for homogeneous (black), acids, and random (red), terpenes). Mean of 20 model instantiations, where each instantiation received a different permutation of odours (see SI Appendix). Annotations below graph indicate whether parameters were fixed (empty circle), variable (filled circle), or variable following a compensation rule (‘H’ for homeostatic tuning,  $f(N, \theta)$  for activity-independent tuning).

1380 alcohols and esters. For each class, we tuned the model’s parameters on that class and then trained the model  
 1381 to classify odours in the other 3 classes (‘novel’ environment). For matched controls, we trained models that  
 1382 had been tuned on the same 3 classes used for training/testing (‘familiar’ environment). As expected, the three  
 1383 activity-dependent models performed worse in novel environments than familiar environments, while the activity-  
 1384 independent model performed consistently regardless of tuning environment (blue, green and magenta vs. cyan in  
 1385 Fig. 2.7C). However, in general, tuning odours on one class but training/testing on different classes does not fatally  
 1386 damage the activity-dependent compensation strategies: although performance is worse in novel environments, it  
 1387 remains better than the random model. Thus, activity-dependent compensation is still a good strategy to overcome  
 1388 the pernicious effects of inter-KC variation, even if the compensation environment differs from the classification  
 1389 environment (at least within the range of the odours in (Hallem and Carlson 2006)).



**Figure 2.8:** Alternative update rules for tuning KCs' input excitatory weights. **(A)** Performance of different models at different indeterminacy constants (A1:  $c = 10$ ; A2:  $c = 1$ ). Blue, left: the method in the main figures, (2.47), where a given KC's input weights are all adjusted equally ('H'); dark blue, middle: (2.46), where a given KC's input weights are adjusted individually according to the average activity of the PN ('H<sub>indiv</sub>'); light blue, right: (2.44), where only non-silent KCs adjust their input weights ('H<sub>active</sub>').  $n = 20$  model instances with different random PN-KC connectivity. Error bars show 95% confidence interval. Bars with the same letter annotations are not significantly different from each other; all other comparisons are significant  $p < 0.05$ , by Wilcoxon signed-rank test with Holm-Bonferroni correction for multiple comparisons. **(B,C)** Probability distribution of the tuned excitatory weights (compare to Fig. 4E). **(D)** The 'H<sub>indiv</sub>' model performs worse than the 'H' model in novel environments (see legend of Fig. 5; the drop in performance from familiar to novel environments is significantly greater for the 'H<sub>indiv</sub>' model,  $p < 0.05$  by Wilcoxon signed-rank test).



**Figure 2.9:** Variants of activity-dependent compensation models.

**(A)** Tuning inhibitory weights to equalize KC average activity improves performance more when we remove the constraint that the coding level without inhibition be double (0.2) the coding level with inhibition (0.1). Coding level without inhibition was 0.2 (left, light green), 0.5 (middle, medium green) or 0.99 (right, dark green).

**(B)** Better performance when spiking thresholds are tuned to equalize KC average activity (magenta) rather than KC response probability (dark magenta), under both more ( $c = 10$ , B1) and less ( $c = 1$ , B2) deterministic decision-making.

**(D)** Probability distribution of spiking thresholds ( $\theta$ ) after tuning them to equalize KCs' response probabilities (compare to Fig. 4E).

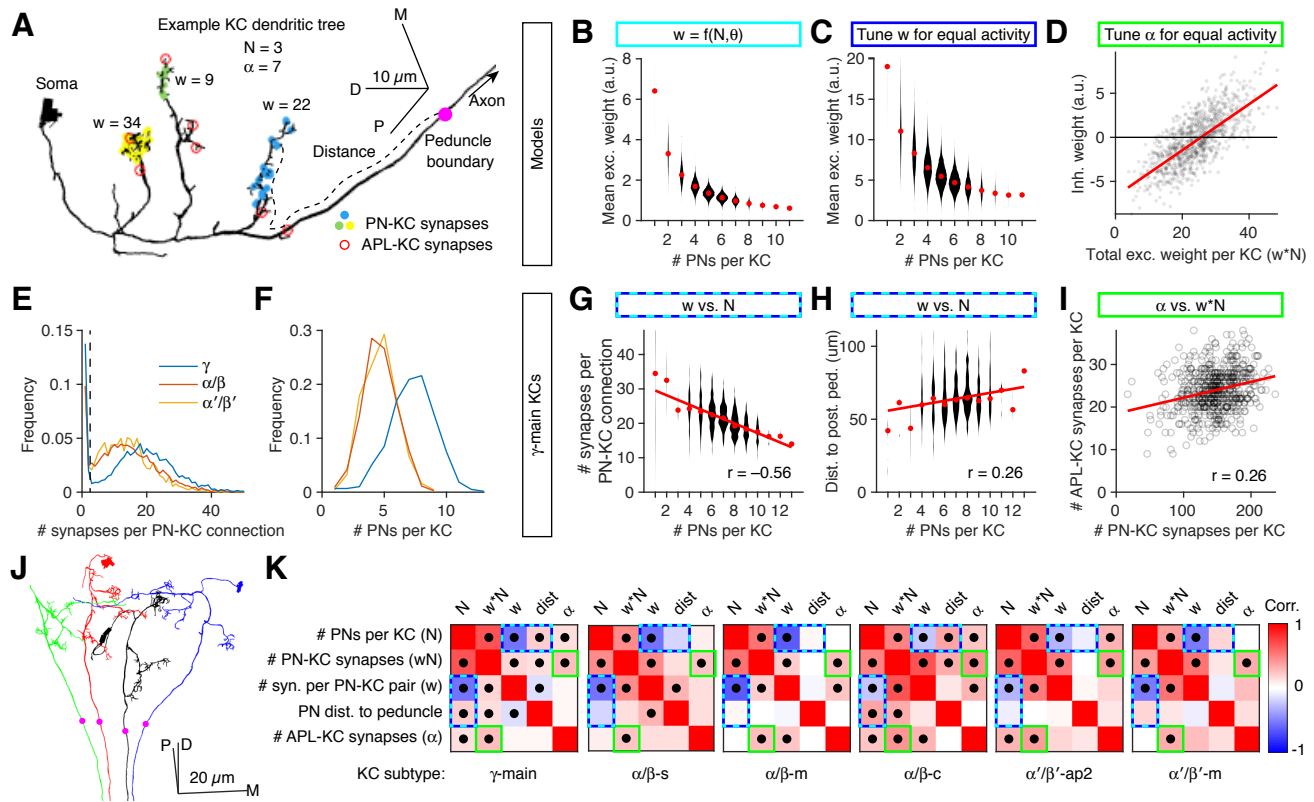
$n = 20$  model instances with different random PN-KC connectivity. Error bars show 95% confidence interval. \*  $p < 0.05$ , by Mann-Whitney test with Holm-Bonferroni correction (A) or Wilcoxon signed-rank test (B).

## 2.4 Connectome reveals compensatory variation of input strength and numbers

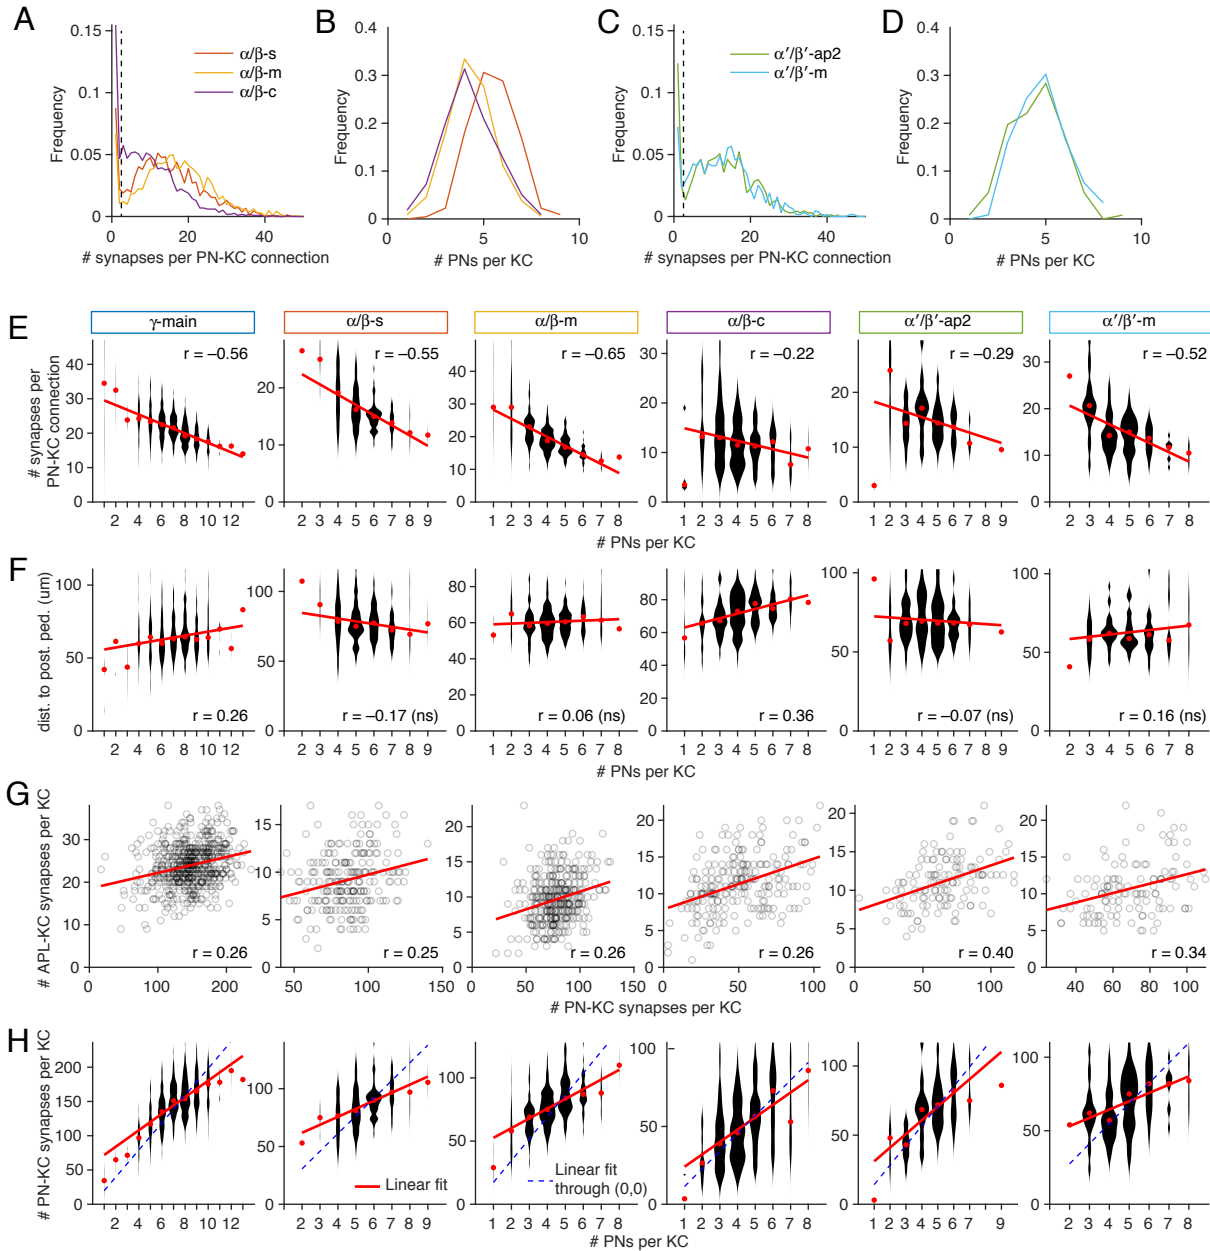
Our proposed compensatory mechanisms predict correlations between the key model parameters. Excitatory weights ( $w$ ) should be inversely correlated to number of PNs per KC ( $N$ ) where  $w$  is tuned to compensate for variable  $N$  and  $\theta$  (Fig. 2.10B) or where  $w$  is tuned to equalize KC activity (Fig. 2.10C). Meanwhile, inhibitory weights ( $\alpha$ ) should be positively correlated to the sum of excitatory weights ( $\sum w$ , or  $\bar{w}N$ , where  $\bar{w}$  is the mean  $w$  per KC) where inhibitory weights are tuned to equalize KC activity (Fig. 2.10D). Such correlations have been observed in larvae (Eichler et al. 2017), but they have not yet been analyzed in the adult mushroom body.

To test these predictions, we analyzed the recently published hemibrain connectome (Scheffer et al. 2020; Li et al. 2020), which annotates all synapses between PNs and KCs in the right mushroom body of one fly. The connectome reveals three of our parameters: the number of PN inputs per KC ( $N$ ), the strength of each PN-KC connection ( $w$ ), and the strength of inhibitory inputs ( $\alpha$ ). Although the anatomy does not directly reveal  $w$  and  $\alpha$  (which can only be measured electrophysiologically), we used an indirect proxy for synaptic strength: the number of synapses per connection (i.e., number of sites between two neurons where neuron 1 has a T-bar and neuron 2 has a postsynaptic density, counted by machine vision; Fig. 2.10A). It seems reasonable to presume that, all else being equal, connections with more synapses are stronger. Indeed, in the *Drosophila* antennal lobe, when comparing connections from ORNs to ipsilateral PNs vs. contralateral PNs, ipsilateral connections are both stronger (Gaudry et al. 2013) and have more synapses per connection (Tobin, Wilson, and Lee 2017). Moreover, synaptic counts approximate synaptic contact area throughout the larval *Drosophila* nervous system (Barnes, Bonnerly, and Cardona 2020) and synaptic area approximates EPSP amplitude in mammalian cortex (Holler et al. 2021).

Therefore, to test if mean  $w$  and  $N$  are inversely correlated across KCs, we asked if the number of PN inputs per KC was inversely correlated to the number of synapses per PN-KC connection. We ignored PN-KC connections with 2 or fewer synapses, because the number of synapses per PN-KC connection formed a bimodal distribution with a trough around 3-4 (Fig. 2.10E); we presumed that connections with only 1-2 synapses represent annotation errors. We divided KCs into their different subtypes as annotated in the hemibrain (Li et al. 2020), because different subtypes have different numbers of PN inputs per KC and different numbers of synapses per PN-KC connection ((Caron et al. 2013); Fig. 2.10E,F, Fig. 2.11A,B,C,D). We excluded KCs that receive significant non-olfactory input ( $\gamma$ -d,  $\gamma$ -t,  $\alpha\beta$ -p,  $\alpha'\beta'$ -ap1). In all analyzed subtypes of KCs ( $\gamma$ -main,  $\alpha\beta$ -s, -m and -c;  $\alpha'\beta'$ -ap2 and -m), the number of PN inputs per KC ( $N$ ) was inversely correlated to the mean number of synapses per PN-KC connection, averaged across the PN inputs onto a KC (proxy for  $\bar{w}$ ) (Fig. 2.10G,K, Fig. 2.11E). Linear regression showed that on average, there were  $\approx 6-15\%$  fewer input synapses per PN-KC connection ( $\bar{w}$ ), for each additional PN per KC ( $N$ ) (compare to the equivalent slopes for the linear fits to the activity-independent (-22%) and activity-dependent (-18%) model parameters in Fig. 2.10B,C). This negative correlation meant that the number of total PN-KC synapses per KC increased only sublinearly relative to the number of PN inputs per KC (Fig. 2.11H).



**Figure 2.10:** Connectome analysis reveals compensatory variation in excitatory and inhibitory input strengths. **(A)** Example  $\alpha\beta$ -c KC (bodyId 5901207528) with inputs from 3 PNs (yellow/green/blue dots) and 7 dendritic APL-KC synapses (red circles). The magenta circle shows the posterior boundary of the peduncle. Line widths not to scale. **(B,C)** Mean synaptic weight ( $w$ ) per PN-KC connection is inversely related to the number of input PNs in models that tune input weights given  $N$  and  $\theta$  **(B)**, or that tune input weights to equalize average activity levels across KCs **(C)**. **(D)** In the model that tunes input inhibitory synaptic weights ( $\alpha$ ) to equalize average activity levels across KCs, inhibitory weights are directly related to the sum of excitatory weights per KC (i.e.,  $wN$ ). Note the negative values of  $\alpha$  (discussed in text). **(E,F)** Probability distributions of the number of synapses per PN-KC connection **(E)** and the number of input PNs per KC **(F)** in the different KCs subtypes ( $\alpha\beta$ ,  $\gamma$ ,  $\alpha'\beta'$ ). Dashed line in **(E)** shows our threshold for counting connections as genuine. **(G)** Mean number of input synapses per PN-KC connection (averaged across PNs for each KC) is inversely related to the number of input PNs per KC, in  $\gamma$ -main KCs (see SI Appendix, Fig. S5 for other KC types). **(H)** Mean distance of PN-KC synapses to the posterior boundary of the peduncle (presumed spike initiation zone) is directly related to the number of input PNs per KC. **(I)** The number of APL-KC synapses per KC is directly related to the total number of PN-KC synapses per KC. **(J)** Four  $\alpha\beta$ -c KCs, one from each neuroblast clone. The posterior boundary of the peduncle (magenta circles) lies where the KC axons begin to converge. **(K)** Grids show Pearson correlation coefficients ( $r$ ) between various KC parameters for all KC subtypes tested (red: positive; blue: negative). Dots indicate  $p < 0.05$  (Holm-Bonferroni corrected). Coloured outlines indicate predictions of models (cyan/blue: models tuning  $w$  **(G,H)**; green: model tuning  $\alpha$  **(I)**). Number of KCs for each subtype, left to right: 588, 222, 350, 220, 127, 119. In **(B,C,G,H)**, red dots are medians and the widths of the violin plots represent the number of KCs in each bin. Trend lines in **(D,G,H,I)** show linear fits to the data. Scale bars in **(A,J)**: D, dorsal, P, posterior, M, medial.



**Figure 2.11:** Connectome analysis on all KC subtypes ( $\gamma$ -main,  $\alpha\beta$ -s, -m and -c;  $\alpha'\beta'$ -ap2 and -m). **(A-D)** Probability distributions of the number of synapses per PN-KC connection (A,C) and the number of input PNs per KC (B,D) in  $\alpha\beta$  and  $\alpha'\beta'$  KCs separated out by subtype (compare to Fig. 6E,F). **(E)** Mean number of input synapses per PN-KC connection is inversely related to the number of input PNs per KC. **(F)** Mean distance of PN-KC synapses to the posterior boundary of the peduncle (presumed spike initiation zone) is directly related to the number of input PNs per KC in  $\gamma$  and  $\alpha\beta$ -c KCs. **(G)** The number of APL-KC synapses per KC is directly related to the total number of PN-KC synapses per KC. **(H)** The number of PN-KC synapses per KCs grows sublinearly with the number of PN inputs per KC. Red dots: medians. Red lines: linear fits. Blue dashed lines: linear fits through the origin (if every PN-KC connection had the same number of synapses). Note that the red dots follow a concave function relative to both linear fits.

1424 We also tested another anatomical proxy of excitatory synaptic strength. Because KCs sum up synaptic inputs  
 1425 linearly or sublinearly, their dendrites likely lack voltage-gated currents that would amplify inputs, so synaptic  
 1426 input currents likely propagate passively (Gruntman and Turner 2013). Therefore, an excitatory input would make  
 1427 a smaller contribution to a KC's decision to spike the farther away it is from the spike initiation zone (Williams and  
 1428 Stuart 2003). While the spike initiation zone cannot be directly observed in the connectome, the voltage-gated  $\text{Na}^+$   
 1429 channel *para* and other markers of the axon initial segment (also called the 'distal axonal segment') are concentrated  
 1430 at the posterior end of the peduncle, near where axons from KCs derived from the four neuroblast clones converge  
 1431 (Ravenscroft et al. 2020; Trunova, Baek, and Giniger 2011). This location can be approximated in the connectome  
 1432 as the posterior boundary of the 'PED(R)' region of interest (ROI) (magenta dots, Fig. 2.10A,J). From this point, we  
 1433 measured the distance along each KC's neurite skeleton (i.e., not the Euclidean distance) to each PN-KC synapse.  
 1434 In the  $\alpha\beta$ -c and  $\gamma$ -main KCs (but not other KCs), this distance was positively correlated with the number of PNs  
 1435 per KC (Fig. 2.10H,K, Fig. 2.11F). That is, the more PN inputs a KC has, the farther away the input synapses  
 1436 are from the putative spike initiation zone (and thus the weaker they are likely to be). Intriguingly, of all the KC  
 1437 subtypes,  $\alpha\beta$ -c KCs show the strongest correlation between number of PN inputs and PN-peduncle distance, but  
 1438 the weakest correlation between number of PN inputs and number of synapses per PN-KC connection (Fig. 2.10K),  
 1439 suggesting that different types of KCs might use different mechanisms to achieve the same compensatory end.

1440 To test if inhibitory and excitatory input are positively correlated across KCs (as predicted in Fig. 2.10D), we  
 1441 approximated  $\alpha$  by counting the number of synapses from the APL neuron to every KC in the calyx (the 'CA(R)'  
 1442 ROI). In all types of KCs, the more total PN-KC synapses there were per KC, the more calyx APL-KC synapses  
 1443 there were (Fig. 2.10I,K, Fig. 2.11G), indicating that indeed, inhibitory and excitatory synaptic input are correlated.

1444 These results confirm the predictions of our compensatory models. That correlations exist for both excitation  
 1445 and inhibition suggests that the mushroom body tunes more than one parameter simultaneously (thresholds may be  
 1446 tuned as well, but cannot be measured in the connectome). Such multi-parameter optimization likely explains (1)  
 1447 why the correlations in the connectome are not as steep as when only a single parameter is tuned in our models (Fig.  
 1448 2.10D-F), and (2) why natural compensatory variation of tuned parameters need not be as wide as the variation of  
 1449 tuned parameters in our models (Fig. 2.5F).

## 1450 2.5 Discussion

1451 Here we studied under what conditions inter-neuronal variability would improve vs. impair associative memory.  
 1452 Using a computational model of the fly mushroom body, we showed that under sparse coding conditions, associative  
 1453 memory performance is reduced by experimentally realistic variability among Kenyon cells in parameters that  
 1454 control neuronal excitability (spiking threshold and the number/strength of excitatory inputs). These deficits arise  
 1455 from unequal average activity levels among Kenyon cells. However, memory performance can be rescued by using  
 1456 variability along one parameter to compensate for variability along other parameters, thereby equalizing average  
 1457 activity among KCs. These compensatory models predicted that certain KC features would be correlated with

each other, and these predictions were borne out in the hemibrain connectome. In short, we showed (1) the computational benefits of compensatory variation, (2) multiple mechanisms by which such compensation can occur, and (3) anatomical evidence that such compensation does, in fact, occur.

Note that when we say ‘equalizing KC activity’, we do not mean that all KCs should respond the same to a given odour. Rather, in each responding uniquely to different odours (due to their unique combinations of inputs from different PNs), they should keep their *average* activity levels the same. That is, while KCs’ odour responses should be heterogeneous, their average activity should be homogeneous.

It will be interesting to use empirically recorded KCs activities to directly test for the existence of compensatory variability mechanisms. First, one could examine the distribution of the KCs average activities in an empirically recorded dataset. That is, one could test if the KCs recorded responses averaged across the input odours (it can be using the odours in the Hallem-Carlson data) are almost equal or not. In addition, since the high variability in the KCs lifetime sparsity values in the random model accounted the most to its impaired memory performance compared to the homogeneous and compensatory variability models (Fig. 2.3 and 2.5C), we can test for the existence of the compensatory variability mechanisms by computing the distribution of the KCs lifetime sparsity levels from their recorded activities and quantifying its fit to the ones obtained from our compensatory models.

How robust are our connectome analyses? We found correlations between anatomical proxies for the physiological properties predicted to be correlated in our models (i.e., KCs receiving excitation from more PNs should have weaker excitatory inputs, while KCs receiving more overall excitation should also receive more inhibition). In particular, we measured the number of synapses per connection as a proxy for the strength of a connection. As described above, this proxy seems valid based on matching anatomical and electrophysiological data (Tobin, Wilson, and Lee 2017; Barnes, Bonnery, and Cardona 2020; Holler et al. 2021). However, other factors affecting synaptic strength (receptor expression, post-translational modification of receptors, pre-synaptic vesicle release, input resistance, etc.) would not be visible in the connectome. Of course, such factors could further enable compensatory variability (see below). It is also worth noting that the connectome data is from only one individual.

We also used the distance between PN-KC synapses and the peduncle as a proxy for the passive decay of synaptic currents as they travel to the spike initiation zone. In the absence of detailed compartmental models of KCs, it is hard to predict exactly how much increased distance would reduce the effective strength of synaptic inputs, but it is plausible to assume that signals decay monotonically with distance. Note that calcium signals are often entirely restricted to one dendritic claw (Gruntman and Turner 2013; Li et al. 2013). Another caveat is that the posterior boundary of the peduncle is only an estimate (though a plausible one: (Ravenscroft et al. 2020; Trunova, Baek, and Giniger 2011)) of the location of the spike initiation zone. However, inaccurate locations should only produce fictitious correlations for Fig. 2.10J and 2.11F if the error is correlated with the number of PN-KC synapses per KC (and only in  $\alpha\beta$ -c and  $\gamma$ -main KCs, not other KCs), which seems unlikely.

Our work is consistent with prior work, both theoretical and experimental, showing that compensatory variability can maintain consistent network behavior (Golowasch et al. 2002; Achard and DeSchutter 2006; Tobin and Calabrese

1493 2006; Taylor, Goillard, and Marder 2009; Marder and Goillard 2006; Schulz, Goillard, and Marder 2006; Schulz,  
1494 Goillard, and Marder 2007; MacLean et al. 2003; MacLean et al. 2005; O’Leary and Marder 2016; Parrish et al.  
1495 2014; Grashow, Brookings, and Marder 2010; Kazama and Wilson 2008). However, to our knowledge, we are the  
1496 first to analyze the computational benefits of equalizing activity levels across neurons in a population (as opposed  
1497 to across individual animals or over time). A recent pre-print showed that equalizing response probabilities among  
1498 KCs reduces memory generalization (Kennedy 2019), but we showed that equalizing average activity outperforms  
1499 equalizing response probabilities (Fig. 2.9). Another model of the mushroom body used compensatory inhibition,  
1500 in which the strength of inhibition onto each KC was proportional to its average excitation (Luo, Axel, and Abbott  
1501 2010), similar to our inhibitory plasticity model (Fig. 2.5A2). However, the previous work did not analyze the specific  
1502 benefits from the compensatory variation; it also set the inhibition strong enough that average net excitation was  
1503 zero, whereas we show that when inhibition is constrained to be only strong enough to reduce KC activity by  $\approx$   
1504 half (consistent with experimental data: (Lin et al. 2014)), inhibition alone cannot realistically equalize KC activity  
1505 (Fig. 2.5G). In addition, there is experimental support for our models’ predictions that KCs with more PN inputs  
1506 would have weaker excitatory inputs: when predicting whether calcium influxes in individual claws would add up  
1507 to cause a supra-threshold response in the whole KC, the most accurate prediction came from dividing the sum of  
1508 claw responses by the log of the number of claws (Li et al. 2013). However, the functional benefits of this result  
1509 only become clear with our computational models. Finally, the larval mushroom body shows a similar relationship  
1510 between number and strength of PN-KC connections: the more PN inputs a KC has, the fewer synapses per PN-KC  
1511 connection (Eichler et al. 2017); however, again, the larval work did not analyze the computational benefits of this  
1512 correlation.

1513 We modeled two forms of compensation: direct correlations between neuronal parameters (Fig. 2.5A1) and  
1514 activity-dependent homeostasis (Fig. 2.5A2-4). Both forms improve performance and predict observed correlations  
1515 in the connectome. Certainly, activity-dependent mechanisms are plausible, as KCs regulate their own activity  
1516 homeostatically in response to perturbations in activity (Apostolopoulou and Lin 2020). Indeed, different KC sub-  
1517 types use different combinations of mechanisms for homeostatic plasticity (Apostolopoulou and Lin 2020), consistent  
1518 with the different correlations observed in the connectome for different KC subtypes. Our activity-dependent mod-  
1519 els lend themselves to straightforward biological interpretations. Excitatory or inhibitory synaptic weights could be  
1520 tuned by activity-dependent regulation of number of synapses per connection or expression/localization of receptors  
1521 or other post-synaptic machinery. Spiking thresholds could be tuned by altering voltage-gated ion conductances  
1522 or moving/resizing the spike initiation zone (Grubb and Burrone 2010; Trunova, Baek, and Giniger 2011). Such  
1523 homeostatic plasticity would be akin to the sensory gain control implemented by feedback inhibition, but on a  
1524 slower timescale.

1525 On the other hand, KCs are not infinitely flexible in homeostatic regulation; for example, complete blockade  
1526 of inhibition causes the same increase in KC activity regardless of whether the blockade is acute (16 - 24 h) or  
1527 constitutive (throughout life) (Apostolopoulou and Lin 2020). This apparent lack of activity-dependent down-



1528 regulation of excitation suggests that activity-independent mechanisms might contribute to compensatory variation  
1529 in KCs, as occurs for ion conductances in lobster stomatogastric ganglion neurons (MacLean et al. 2003; MacLean  
1530 et al. 2005). For example, the inverse correlation of  $w$  and  $N$  arises from the fact that the number of PN-KC  
1531 synapses per KC increases only sublinearly with increasing numbers of claws (i.e., PN inputs) (Fig. 2.11H). Perhaps  
1532 a metabolic or gene regulatory constraint prevents claws from recruiting postsynaptic machinery in linear proportion  
1533 to their number. (Interestingly, this suppression is stronger in larvae, where the number of PN-KC synapses per KC is  
1534 actually constant relative to the number of claws: (Eichler et al. 2017).) Meanwhile, the correlation between number  
1535 of inhibitory synapses and number of excitatory synapses might be explained if excitatory and inhibitory synapses  
1536 share bottleneck synaptogenesis regulators on the post-synaptic side. Although activity-dependent compensation  
1537 produced superior performance in our model compared to activity-independent compensation thanks to its more  
1538 effective equalization of KC average activity (Fig. 2.5) (most likely because it takes into account the unequal  
1539 activity of different PNs), activity-dependent mechanisms suffered when the model network switched to a novel  
1540 odour environment (Fig. 2.7). Given that it is desirable for even a newly enclosed fly to learn well, and for flies to  
1541 learn to discriminate arbitrary novel odours, activity-independent mechanisms for compensatory variation may be  
1542 more effective in nature.

1543 Compensatory variability to equalize activity across neurons could also occur in other systems. The vertebrate  
1544 cerebellum has an analogous architecture to the insect mushroom body; cerebellar granule cells are strikingly  
1545 similar to Kenyon cells in their circuit anatomy, proposed role in ‘expansion recoding’ for improved memory, and  
1546 even signaling pathways for synaptic plasticity (Modi, Shuai, and Turner 2020; Farris 2011; Litwin-Kumar et al.  
1547 2017; Marr 1969; Handler et al. 2019; Aso et al. 2019). Whereas cortical neurons’ average spontaneous firing rates  
1548 vary over several orders of magnitude (Buzsáki and Mizuseki 2014), granule cells are, like Kenyon cells, mostly silent  
1549 at rest, and it is plausible that their average activity levels might be similar (while maintaining distinct responses  
1550 to different stimuli) (Powell et al. 2015). Granule cell input synapses undergo homeostatic plasticity (Delvendahl,  
1551 Kita, and Müller 2019), while compartmental models suggest that differences in granule cells’ dendritic morphology  
1552 would affect their activity levels, an effect attenuated by inhibition (Houston et al. 2017), raising the possibility  
1553 that granule cells may also modulate inter-neuronal variability through activity-dependent mechanisms. Future  
1554 experiments may test whether compensatory variability occurs in, and improves the function of, the cerebellum or  
1555 other brain circuits. Finally, activity-dependent compensation may provide useful techniques for machine learning.  
1556 For example, we found that performance of a reservoir computing network could be improved if thresholds of  
1557 individual neurons are initialized to achieve a particular activity probability given the distribution of input activities  
1558 (Manneschi, Lin, and Vasilaki in press).

## 1559 Note on originality

1560 This chapter was previously published as **Abdelrahman NY, Vasilaki E, Lin AC (2021). Compensatory**  
1561 **variability in network parameters enhances memory performance in the Drosophila mushroom body.**

1562 **PNAS 118:e2102158118. <https://doi.org/10.1073/pnas.2102158118>.**

1563 The only edits compared to the published version are to include the supplemental methods and data in the main  
1564 text. The publication was also edited by my supervisor Andrew C Lin.

1565 Some of the plots and analyses were made by my supervisor Andrew C Lin. In particular, the KC schematic in  
1566 panel (A) of Fig. 2.10, and the analysis of the number of PNs per KCs versus the distance to the posterior peduncle  
1567 in Fig. 2.10 panel (H). Besides, my supervisor Andrew C Lin has used a custom code I created initially to analyse  
1568 the KCs sub types  $\gamma$ ,  $\alpha/\beta$  and  $\alpha'/\beta'$  and to identify inhibitory and excitatory synapses, and he extended it to  
1569 analyse KCs sub types in a more detailed level which is  $\alpha/\beta$ -s,  $\alpha/\beta$ -m,  $\alpha/\beta$ -c,  $\alpha'/\beta'$ -ap2 and  $\alpha'/\beta'$ -m. He analysed  
1570 and plotted the results for the more detailed KCs sub types in panel (K) Fig. 2.10. Last but not least, he also did  
1571 the analysis and plots in Fig. 2.11.

## 1572 Chapter 3

# 1573 Computational benefits of learning by 1574 depression emerge in the model fruit fly 1575 under a divisive normalisation decision 1576 making policy

### 1577 3.1 Introduction

1578 In fruit flies, the repetitive coupling of an input odour (conditioned stimulus) with a reward or punishment (uncon-  
1579 ditioned stimulus) induces long term plasticity in the KCs' output synapses onto the MBONs. In particular, this  
1580 plasticity happens via long term depression (LTD), where the output synapses from the KCs are weakened onto the  
1581 MBONs that encode a behaviour (approach or avoidance) opposite to the valence of the unconditioned stimulus  
1582 (reward or punishment); in contrast to learning by potentiating the right behaviour.

1583 There exists a great body of literature which has firmly established the biological mechanisms responsible for  
1584 inducing long term potentiation or depression (LTP and LTD, respectively) in mammals (Malenka 1991; Markram  
1585 et al. 1997; Dan and Poo 2006) and other species. Nonetheless it remains elusive if the direction of plasticity in  
1586 these neural circuits has evolved in a non-random way to optimize some aspects of learning.

1587 The idea of optimality in neural circuits (and natural systems) has provided some explanations for long left  
1588 unanswered questions in neuroscience. One among many was to try to explain the observed variability in the  
1589 plasticity loci, presynaptically or postsynaptically, (Bolshakov and Siegelbaum 1995; Zakharenko, Zablow, and  
1590 Siegelbaum 2001; Lisman and Raghavachari 2006). Indeed, it was only recently suggested that this variability  
1591 could be a result of an optimization protocol of the postsynaptic neuron statistics (Costa et al. 2017). In addition,  
1592 other computational studies showed that the network structure in some circuits has evolved to optimise for different

1593 aspects like, energy constraints or learning performance. For example, the degree of neuronal expansion from the  
1594 PNs to KCs in the fruit fly MB or in the mossy fibres to granule cells in the human cerebellum optimises the learning  
1595 performance and sensory encoding (Litwin-Kumar et al. 2017). Another work showed that the number of neurons  
1596 in the third layer of the olfactory circuits, granule cells in mammals and KCs in flies and other insects, scales as the  
1597 number of the input neurons in a manner that optimises for the learning performance and depends on the animal's  
1598 lifetime (Hiratani and Latham 2022). Along the same line of thought here I ask: Can the learning rule observed in  
1599 the real KCs-MBONs synapses also be a result of an optimization framework?

1600 What are the cases where learning by depression would offer better memory performance and data encoding  
1601 than in potentiation?

1602 I approached this question by modeling an associative memory task in the mushroom body using a rate coding  
1603 network. I modelled two functions for the decision making strategy in this network: a soft-max (see the next section  
1604 and Chapter 4 for details) and *divisive* normalisation function, which normalizes the difference between outputs of 2  
1605 MBONs encoding opposing behaviours with their sums (see details in the next section). Interestingly, I found that  
1606 the memory performance was indeed optimal when learning happens by depression compared to potentiation but  
1607 only when the decision making strategy was like a divisive normalisation. I also suggest that a Winner-Takes-All  
1608 circuit (WTA) architecture between the MBONs outputs can serve as a bio-plausible implementation for divisive  
1609 normalisation. The analysis of this WTA circuit model reveals that depression outperforms potentiation only in  
1610 the presence of multiplicative noise in the MBONs responses. This steers attention to an attractive avenue where  
1611 these theoretical predictions can be tested by quantifying the type of the MBONs' noise in-vivo.

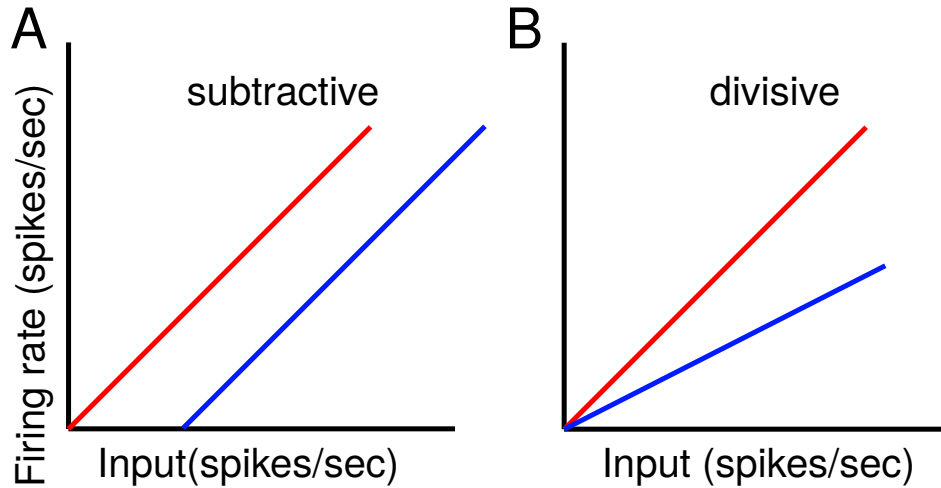
## 1612 3.2 Normalisation as a canonical operation in neural circuits

1613 To understand the predictions drawn from my model, one should first get familiar with the concept of normalisation,  
1614 its potential benefits and existence in neural circuits across different species.

1615 Normalisation has been identified as one of the canonical operations in neural computations. Generally, nor-  
1616 malisation can be seen as the way for the brain to implement gain control on its inputs. Gain control is essentially  
1617 the ability of a neuron to control its output responses to stay within a certain dynamic range, irrespective of the  
1618 absolute values of its inputs (mean energy of the inputs).

1619 It can have different forms: (a) subtractive (or additive) which is usually implemented via a shunting inhibition-  
1620 like mechanism (Holt and Koch 1997). This type of normalisation doesn't change the shape of the neuron's input-  
1621 output curve, that is also referred to as firing rate - Input current (f-I) curve as in (Fig3.1A). In contrast, there  
1622 is another type of normalisation: (b) divisive normalisation, this can change the slope of the neural input-output  
1623 behaviour Fig(3.1B). In this case the normalisation factor is a constant that scales the slope of the neural linear  
1624 output function. In contrast, I will use a non-linear neural output function throughout this chapter as depicted in  
1625 Fig(3.2) (using Eq.(3.1) and (3.2)), where the neural output function becomes non-linear beyond a certain input  
1626 value and saturates at a maximum level.

1627 In an elegant review by (Carandini and Heeger 2012) divisive normalisation has explained numerous non-linear  
 1628 phenomena in neural responses across various sensory circuits. For instance, it has successfully accounted for the  
 1629 surround suppression phenomenon in the primary visual cortex (Heeger, Landy, and Movshon 1991; Heeger 1992;  
 1630 Albrecht and Geisler 1991; Carandini, Heeger, and Movshon 1997), neural responses in olfactory systems (Olsen,  
 1631 Bhandawat, and Wilson 2010), and in the brain areas responsible for context-based decision making (Louie, Grattan,  
 1632 and Glimcher 2011).



**Figure 3.1:** Types of normalisation. The red curve corresponds to the original neuron’s (f-I) curve before gain control. (A) the blue curve is the result of modulating the red curve with a subtractive normalisation. (B) the blue curve is the result of modulating the red curve with a divisive normalisation.

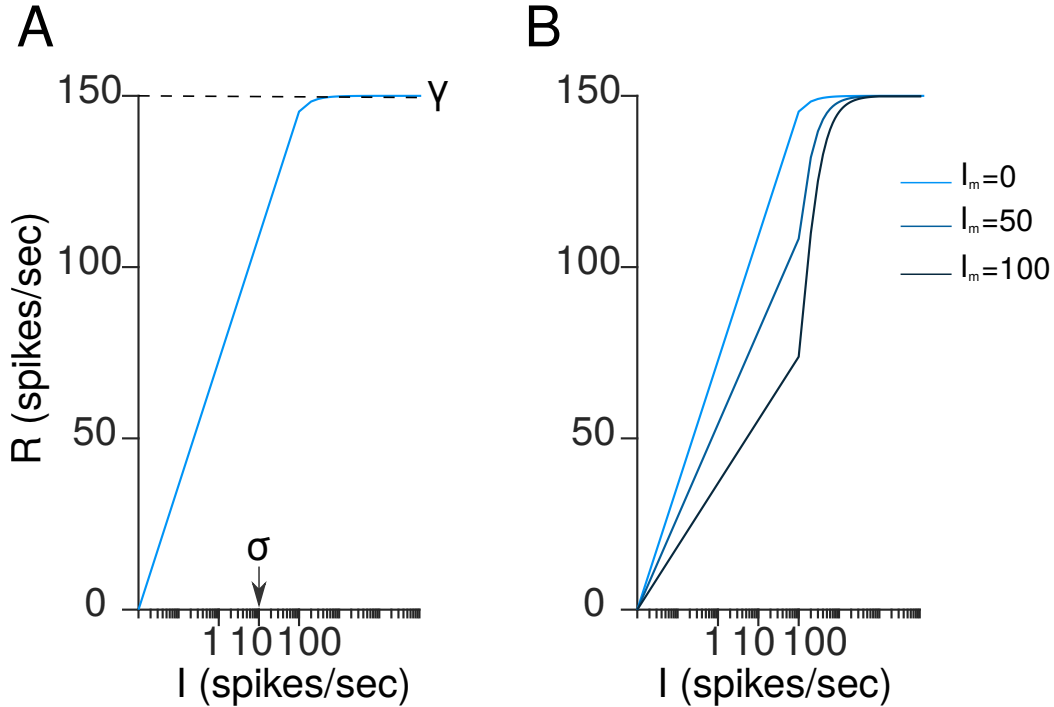
1633 In fruit flies, the outputs in the antennal lobe neurons was explained using a divisive normalisation transfer  
 1634 function. A neuron response will saturate after the input odour strength increases past a certain value, as in Eq.  
 1635 (3.1); similar to the results from Fig. 2 in (Olsen, Bhandawat, and Wilson 2010). The input value after which the  
 1636 neuron response saturates as well as the value at which it will saturate both depend on the variables  $\gamma$ ,  $\sigma$  and  $n$  in  
 1637 Eq.(3.1).

$$R = \gamma \frac{I^n}{\sigma^n + I^n} \quad (3.1)$$

1639 In addition, an antennal lobe neuron’s response was found to be suppressed in the presence of another odour,  
 1640 referred to as ”mask” odour, which would not normally evoke a response in this neuron by its own. This suppression  
 1641 happens by means of divisive normalisation as in Eq.(3.2), as shown in (Fig.3.2B), such that the activity in a neuron  
 1642 will be normalised by dividing its response  $I$ , with the pool of activity from the other neurons responding to the  
 1643 mask odour,  $I_m$ . This explains the shift in the antennal neural responses to the right as the concentration of the  
 1644 mask odour increases. Note that the equation below is similar to the equation I used in Chapter 2 (see Methods,

Eq.2.1) to simulate the PNs inputs to KCs.

$$R = \gamma \frac{I^n}{\sigma^n + I_m^n + I^n} \quad (3.2)$$



**Figure 3.2:** Normalisation in the olfactory neurons in the fruit fly. (A) The response of an antennal lobe neuron to input odour concentration as in Eq.(3.1).  $\gamma$  and  $\sigma$  values are labeled on the graphs.

Divisive normalisation helps to maximize the sensitivity of neural responses by shifting the steepest regions (most sensitive to input changes) in their responses curves towards the mean energy in the input stimuli. This was found in the antennal neurons in the fruit fly, as shown in (Fig.3.2B), but was also evident in the retinal neurons, which allowed them to adapt to the wide range of light levels present in the same scene. Through divisive normalisation they can utilise their response ranges maximally to encode relevant features irrespective of the background illumination (Boynton and Whitten 1970; Normann and Perlman 1979; Burkhardt 1994). This results in scale-invariant neural responses, where the retinal neurons will encode the contrast levels in a visual scene the same way if its background illumination was scaled up or down.

In addition, divisive normalisation successfully accounts for the cross-orientation suppression phenomenon (among other non-linear properties) observed in the primary visual cortex neurons (V1) (Carandini, Heeger, and Movshon 1997; Bauman and Bonds 1991; Freeman et al. 2002). V1 neurons selectively respond to test gratings which have specific orientations and lie spatially within the neuron's preferred field of response. However, a neuron's response to a preferred stimulus is reduced if a different, non-preferred stimulus (e.g., another grating moving in a different direction), is superimposed. It was unclear what could cause this non-linear behaviour, until some

1661 experimental work expressly designed to test the normalisation model has strongly suggested its good quantitative  
 1662 fits with the experimental data (Carandini, Heeger, and Movshon 1997).

1663 Not only does the divisive normalisation offers scale invariant population responses, but it also helps to increase  
 1664 the neural coding efficiency and reduce its redundancy. Neural responses are considered inefficient if they are highly  
 1665 correlated for different inputs. Divisive normalisation was found to increase the statistical independence among  
 1666 the primary visual cortex neurons responses(V1) (Schwartz and Simoncelli 2001a; Lyu and Simoncelli 2009). As  
 1667 well, it is responsible for eliminating the dependency between the projection neurons compared to the upstream  
 1668 (pre-normalised) olfactory receptor neurons responses (Luo, Axel, and Abbott 2010; Olsen, Bhandawat, and Wilson  
 1669 2010).

### 1670 **3.3 Mathematical analysis: Learning by depression is optimal under** 1671 **a divisive normalisation decision making policy.**

1672 Not only divisive normalisation was observed as a canonical operation across many sensory neural circuits, but it was  
 1673 also evident in cortical areas involved in decision making. It was found that divisive normalisation accounts for the  
 1674 values modulation and relative evaluation of rewarded options presented to the participants in (Louie, Grattan, and  
 1675 Glimcher 2011). In particular, they studied the neurons in the lateral intraparietal cortex (LIP) area in monkeys,  
 1676 these neurons are unique since they are at the nexus of decision making and sensory stimuli representation. LIP  
 1677 neurons are activated by visual stimuli in a certain visual field, as well they fire when the monkey decides to pick  
 1678 an option from a set of visual inputs by saccading eye movements (Platt and Glimcher 1999). In Louie et al, they  
 1679 hypothesized that the LIP neurons activities are modulated via divisive normalisation with respect to the pooled  
 1680 activity of the other LIP neurons, which encode other saccades directions associated with higher or lower rewards.

1681 MBONs in the fruit fly are similar to the LIP neurons, to some extent. They receive olfactory responses from  
 1682 the KCs as their inputs and are upstream to behaviour-guiding neurons. It is plausible that MBONs' outputs may  
 1683 be combined by divisive normalisation. The question that naturally follows, what will be the benefits of divisive  
 1684 normalisation if it exists in the MBONs circuitry? The answer is not too straightforward. It will require formal  
 1685 analyses to reveal it. To this end, I will introduce in this section a mathematical treatment which suggests that  
 1686 the learning rule between the KCs-MBONs weights could have evolved by depression to optimise the memory  
 1687 performance compared to the other alternative rule by potentiation. The computational benefits of learning by  
 1688 depression will follow if and only if the bias towards one decision (difference between the opposite decisions values  
 1689 which are encoded by the MBONs firing rates) is normalised by dividing it with the sum of both decisions values  
 1690 (sum of the MBONs firing rates), i.e. context-based value learning.

1691 Consider the learned bias,  $R$ , to avoid an odour, ranging from -1 to 1, which is related to the probability of  
 1692 avoiding the odour (ranging from 0 to 1) by,

$$1693 \quad R = 2[probability]-1 \quad (3.3)$$

1694 Under softmax decision policy, the learned bias to avoid an odour can be written as,

$$\begin{aligned}
 R &= \frac{e^{nAvoid} - e^{nApproach}}{e^{nAvoid} + e^{nApproach}} \\
 &= 2 \left[ \frac{e^{nAvoid}}{e^{nApproach} + e^{nAvoid}} \right] - 1 \\
 &= 2 \left[ \frac{1}{e^{-n(Avoid-Approach)} + 1} \right] - 1
 \end{aligned}
 \tag{3.4}$$

1696 Where *Avoid* and *Approach* are the responses of the avoidance and approach MBONs, respectively, to an odour  
 1697 after memory formation. The multiplicative factor  $n$  in Eq.(3.4) controls the steepness of the Sigmoid function,  
 1698 the speed at which the probability shifts given a difference between the *Avoid* and *Approach* MBONs readouts. A  
 1699 decision policy with divisive normalisation will use a similar form to Eq. (3.4), but will be:

$$R = g \left[ \frac{Avoid^n - Approach^n}{k^n + Avoid^n + Approach^n} \right]
 \tag{3.5}$$

1701 Again, the exponent  $n$  expresses how steeply the bias changes with the ratio of Avoid and Approach.  $g$  is the  
 1702 gain and  $k$  is a constant to prevent division by zero that ends up controlling how ‘divisive’ the normalisation is.

1703 Consider a case with one Avoidance MBON and one Approach MBON. The activity of each MBON is the sum  
 1704 of active KCs, weighted by the KC-MBON synaptic strength, e.g., if  $x_i$  is the activity of the  $i$ th KC and  $W_{av,i}$  is  
 1705 the weight of synapse from the  $i$ th KC to the Avoid MBON, then  $Avoid = \sum_i^n W_{av,i} x_i$ . Consider the symmetrical  
 1706 cases: depressing KC-MBON synaptic weights toward incorrect actions from 1 to  $z$  ( $0 \leq z \leq 1$ ), Vs. potentiating  
 1707 them towards the correct actions, from  $z$  to 1. Suppose odour A activates a set of  $a$  unique KCs, shown by red  
 1708 colour in Fig.3.3, odour B activates a set of  $b$  unique KCs, shown by green in Fig.3.3, and both A and B activate  
 1709 the same set of  $c$  overlapping KCs, assuming one KC for simplicity shown by blue in Fig.3.3. Suppose further that  
 1710 KCs have binary activity of 0 or 1, and the fly experiences a sequential learning of [A+punishment], followed by  
 1711 [B+ reward].

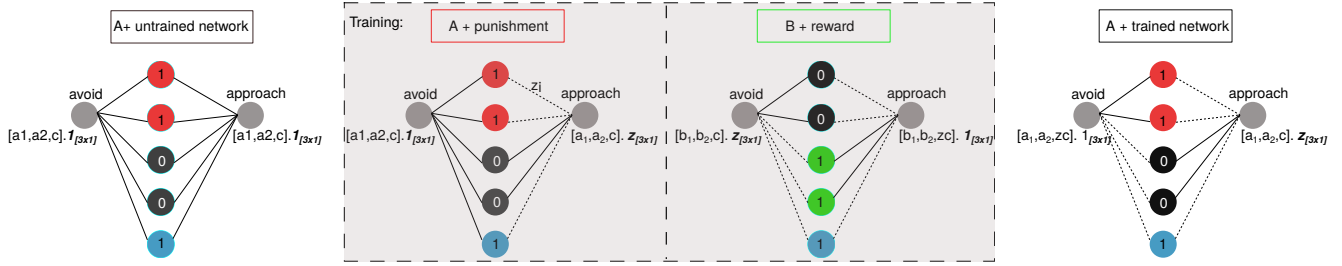




**Figure 3.3:** Each of odour A and B activates unique KCs shown in red and green respectively. Both odours activate an overlapping KC shown in blue.

1712 In the case of learning by depression, the Approach MBON response to odour A is  $(a + c)$  before training (the  
 1713 left panel in Fig.3.4) and  $z(a + c)$  after training (the mid panel in Fig.3.4) because all of the  $(a + c)$  KCs have  
 1714 their output synapses depressed from 1 to  $z$ . The Avoidance output will stay the same at this stage of training,  
 1715 i.e  $(a + c)$ . Yet, after pairing [B+reward] the Avoidance MBON response to odour B will decrease from  $(b + c)$  to  
 1716  $z(b + c)$ . Importantly, after this coupling the Avoidance MBON output response to odour A will not stay the same,  
 1717 instead it will decrease from  $(a + c)$  to  $(a + zc)$ , because only the  $c$  overlapping KCs were depressed by the pairing  
 1718 of odour [B+reward], as shown by the right panel in Fig.3.4.

1719 On the other hand, the case of potentiation, both MBONs' responses to odour A start out at  $z(a + c)$ ; the  
 1720 Avoidance MBON's response increases to  $(a + c)$  (all  $(a + c)$  KCs have their output synapses potentiated from  $z$   
 1721 to 1), while the Approach MBON's responses increases from  $(z(a + c))$  before training to  $(za + c)$  after pairing  
 1722 [B+reward] (the  $a$  KCs unique to odour A remain unchanged while the  $c$  overlapping ones have their output synapses  
 1723 potentiated from  $z$  to 1).



**Figure 3.4:** Changes in the output weights from KCs to MBONs after learning by depression. Learning by penalising the wrong action.

1724 Under a softmax policy and learning by depression, the  $P_{avoid}$  will look like:

$$\begin{aligned}
 P_{avoid} &= \frac{1}{1 + e^{-n(Avoid - Approach)}} \\
 &= \frac{1}{1 + e^{-n(a + zc - z(a + c))}} \\
 &= \frac{1}{1 + e^{-n(1-z)a}}
 \end{aligned} \tag{3.6}$$

1726 Likewise, the  $P_{avoid}$  will be the same when learning by potentiation, because the  $c$  terms in the exponent in the  
 1727 denominator will also cancel out. That is, depression and potentiation are equivalent under softmax.

1728 On the contrary, under divisive normalisation and learning by depression, the learned bias ( $R_d$ ) and probability  
 1729 of avoiding odour A ( $P_{avoid|d}$ ) will be,

$$\begin{aligned}
 R_d &= g \left( \frac{Avoid^n - Approach^n}{k^n + Avoid^n + Approach^n} \right) \\
 &= g \left( \frac{(a + zc)^n - (z(a + c))^n}{k^n + (a + zc)^n + (z(a + c))^n} \right)
 \end{aligned} \tag{3.7}$$

1731 And,

$$\begin{aligned}
 P_{avoid|d} &= 0.5(R_d + 1) \\
 &= 0.5 \left( g \frac{(a + zc)^n - (z(a + c))^n}{k^n + (a + zc)^n + (z(a + c))^n} + 1 \right) \\
 &= 0.5 \left( g \frac{2(a + zc)^n + k^n}{k^n + (a + zc)^n + (z(a + c))^n} \right)
 \end{aligned} \tag{3.8}$$

1733 Whilst under potentiation the learned bias ( $R_p$ ) and probability of avoiding odour A ( $P_{avoid|p}$ ) will be,

$$\begin{aligned}
 R_p &= g \left( \frac{Avoid^n - Approach^n}{k^n + Avoid^n + Approach^n} \right) \\
 &= g \left( \frac{(a + c)^n - (za + c)^n}{k^n + (a + c)^n + (za + c)^n} \right)
 \end{aligned} \tag{3.9}$$

1735 And,

$$\begin{aligned}
 P_{avoid|p} &= 0.5(R_p + 1) \\
 &= 0.5 \left( g \frac{(a + c)^n - (za + c)^n}{k^n + (a + c)^n + (za + c)^n} + 1 \right) \\
 &= 0.5 \left( g \frac{2(a + c)^n + k^n}{k^n + (a + c)^n + (za + c)^n} \right)
 \end{aligned} \tag{3.10}$$

1737 if  $n=1$ , then Eq.(3.8) and (3.10) become,

$$1738 \quad P_{avoid|d} = 0.5g \left( \frac{2a + 2zc + k}{k + (1+z)a + 2zc} \right) \quad (3.11)$$

1739 And,

$$1740 \quad P_{avoid|p} = 0.5g \left( \frac{2a + 2zc + k}{k + (1+z)a + 2c} \right) \quad (3.12)$$

1741 Comparing Eq. (3.11) and (3.12) and since  $z < 1$ , the probability to avoid the punished odour  $A$  is lower in  
 1742 potentiation than in depression. If  $c=0$ , i.e. no overlapping KCs between the responses for odours  $A$  and  $B$ , then  
 1743 learning by depression and potentiation are equivalent.

1744 We can find the conditions under which the learned bias under potentiation is less than that under depression for  
 1745 the more general cases where  $n > 0$ . For the sake of simplifying the expression later on, let's define, the following  
 1746 quantities:  $[v = a + zc]$ ,  $[p = z(a + c)]$  and  $[x = (1 - z)c]$ . Hence, we can rewrite both  $R_p$  (Eq.(3.9)) and  $R_d$   
 1747 (Eq.(3.7)) in terms of these quantities  $v$ ,  $p$  and  $x$  as follows,

$$1748 \quad R_d = g \left( \frac{(a + zc)^n - (z(a + c))^n}{k^n + (a + zc)^n + (z(a + c))^n} \right) \quad (3.13)$$

$$= g \left( \frac{v^n - p^n}{k^n + v^n + p^n} \right)$$

$$1749 \quad R_p = g \left( \frac{(a + c)^n - (za + c)^n}{k^n + (a + c)^n + (za + c)^n} \right) \quad (3.14)$$

$$= g \left( \frac{(v + x)^n - (p + x)^n}{k^n + (v + x)^n + (p + x)^n} \right)$$

1750 When there is no overlap between the KCs responses for odours  $A$  and  $B$ , i.e.  $c=0$  and  $x = 0$ , again, the learned  
 1751 bias by potentiation and depression are the same, Eq.(3.13) and (3.14) are the same. However, as  $c$  increases, the  
 1752 learned bias under potentiation will always be less than that under depression in the cases where  $n < 1$ . To see  
 1753 this, first we need to compute the derivative of the learned bias under potentiation with respect to  $x$ , hence:

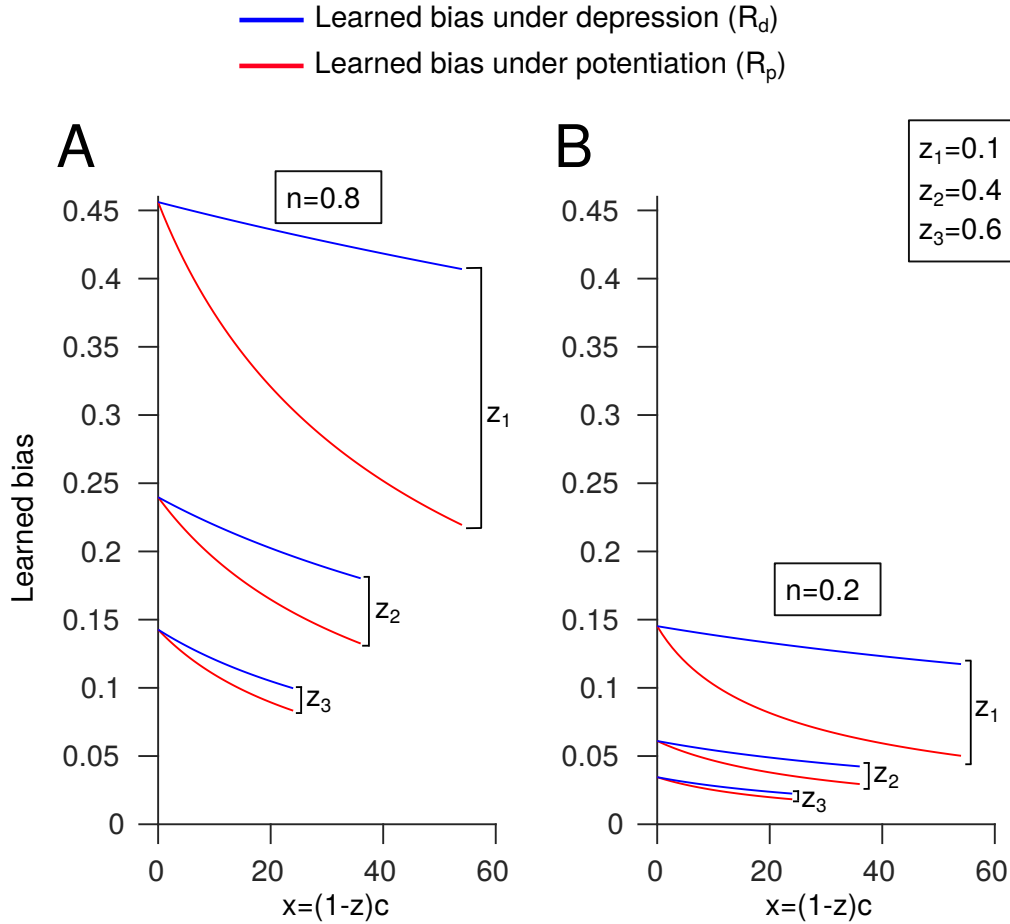
$$1754 \quad \frac{\delta R_p}{\delta x} = \frac{\delta}{\delta x} g \left( \frac{(v + x)^n - (p + x)^n}{k^n + (v + x)^n + (p + x)^n} \right) \quad (3.15)$$

$$= n \frac{k^n [(v + x)^{n-1} - (p + x)^{n-1}] + 2(p - v)(p + x)^{n-1}(v + x)^{n-1}}{(k^n + (v + x)^n + (p + x)^n)^2}$$

1755 if  $n < 1$ , and  $p < v$ , then the numerator in the expression above in Eq.(3.15) will always be negative for any  $x > 0$ ,  
 1756 i.e.  $\frac{\delta R_p}{\delta x}$  will be always negative. Thanks to the rearrangement of the variables  $(a, b, c)$  into  $(v, p, x)$ , it becomes  
 1757 more obvious the different effects that  $c$  will have on the learned biases under depression compared to potentiation.

1758 In particular when  $n < 1$  the increase in  $c$  will cause more decrease in  $R_p$  than in  $R_d$ , captured in the quantity  
 1759  $x$ , which equals 0 in  $R_d$ , ( $R_d = R_p|_{x=0}$ ).  $R_d$  will also drop as  $c$  grows, but it will always stay higher than  $R_p$ .  
 1760 This is depicted in (Fig3.5) under different values of  $n$  and  $z$ , while holding the other parameters constant:  $k = 10$ ,  
 1761  $a = b = 10$ . As  $c$  is varying in an arbitrary range from 15 to 60, so is  $x$  where  $x = (1 - z)c$ . In both cases of  $n$   
 1762 values shown below in (Fig.3.5), as  $z$  increases in the values from  $[0.1, 0.4, 0.6]$  the range of  $x$  decreases, with the

1763 blue and red curves at value of  $z = 0.6$  being shorter than those at  $z = 0.1$ , for the same  $\{n, a, b, k\}$ , (Fig.3.5 A  
 1764 and B). Also, as  $z$  increases (or  $x$  decreases) the learned bias under depression will start to be similar to that under  
 1765 potentiation, bringing the blue and red curves with the same value of  $z$  closer to each other, (Fig.3.5 A and B).



**Figure 3.5:** Learned bias in depression is always higher than potentiation for any  $n < 1$ : (A)  $n = 0.8$ , (B)  $n = 0.2$ . Learned bias by depression remains higher than potentiation for different conditions of  $z$ : results are shown for  $z=0.1, 0.4$  and  $0.6$ . The gap between the learned biases under depression and potentiation closes as the magnitude of the synaptic plasticity,  $z$ , increases.

1766 For the other region of the  $n$  space, when  $n > 1$ , there will be cases where  $R_p$  is less than  $R_d$  and others where  
 1767 the reverse will happen. That depends on the relative values of  $k$  and  $a$ .  $R_p$  will be less than  $R_d$  for small values  
 1768 of  $k$ , in particular when  $k^n$  is less than the quantity below,

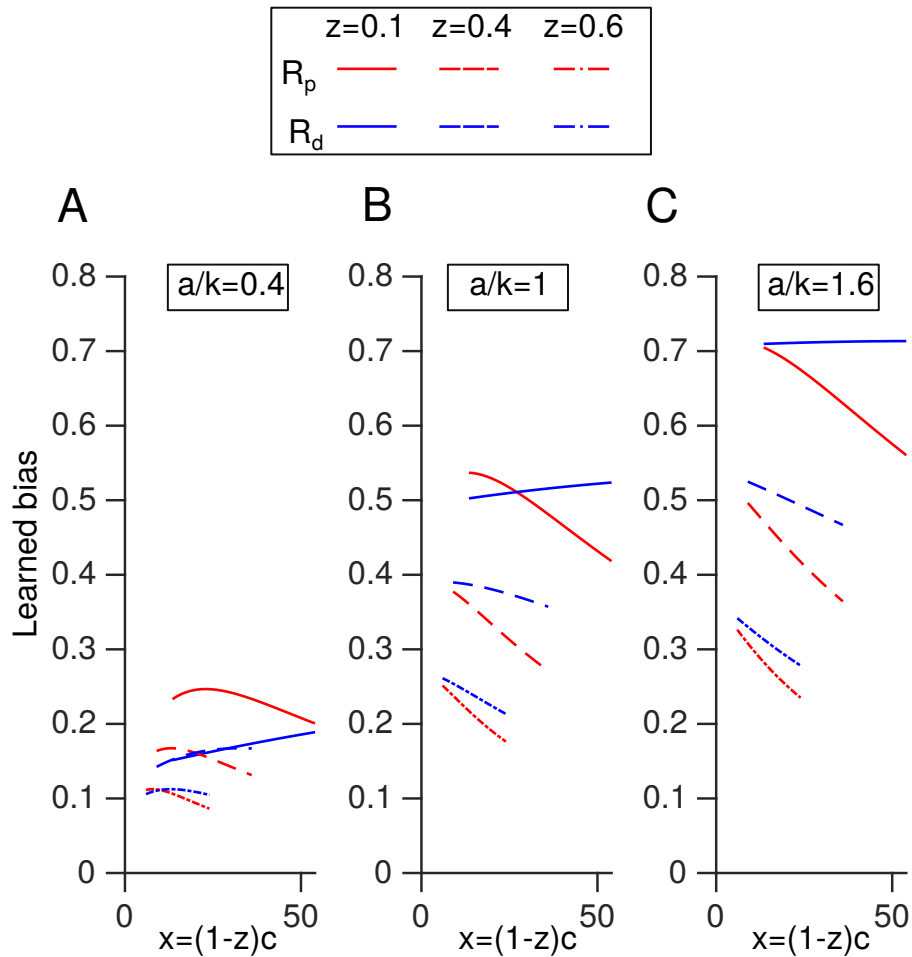
$$k^n < \frac{2(v^n(p+x)^n - p^n(v+x)^n)}{((v+x)^n - v^n - (p+x)^n + p^n)} \quad (3.16)$$

$$k^n < \frac{2((a+zc)^n(za+c)^n - z^n(a+c)^{2n})}{((1+z^n)(a+c)^n - (a+zc)^n - (za+c)^n)}$$

1770 rewriting Eq.(3.16) using  $y = \frac{c}{a}$ , i.e., ratio of overlapping cells to unique cells

$$k^n < 2a^n \frac{(1+zy)^n(z+y)^n - z^n(1+y)^{2n}}{(1+z^n)(1+y)^n - (1+zy)^n - (z+y)^n} \quad (3.17)$$

1772 This reformulation shows that learning by depression will outperform potentiation for low values of  $k$ , in particular  
 1773 when it is small enough relative to  $a$  (the number of unique active KCs). Intuitively if an odour elicit sparser  
 1774 responses, then the chance of it having overlapping KCs with other odour responses will be slimmer, which means  
 1775 that  $k$  will dominate the denominator in Eq. (3.13) and (3.14) and the learned bias will behave more like a softmax  
 1776 policy. Figure 3.6 illustrates the effect of the relative magnitudes of  $k$  and  $a$ , at different values of  $z$  and  $c$ , on  $R_p$   
 1777 and  $R_d$ .



**Figure 3.6:** Learned bias by depression is better than potentiation for high values of  $\frac{a}{k}$ , when  $n > 1$ . **A** learning by potentiation outperforms depression for low values of the ratio  $\frac{a}{k}$ . The order switches as  $z$  increases. As  $z$  increases, learning by depression crosses the learning curve of potentiation at lower values of  $c$ . **B** learning by depression outperforms potentiation under more conditions of  $z$  and  $c$  when  $\frac{a}{k} = 1$ . **C** Depression outperforms potentiation under all conditions of  $z$  when  $k$  is less than  $a$ ,  $\frac{a}{k} > 1$ , equals 1.6. Learned biases for different values of  $z$  are shown by the solid, dashed and dotted curves respectively.

1778 (Please see notes of originality for more details on my partial contribution to the this subsection)

### 3.4 Simulation results: Learning by depression enhances the memory performance in the fruit fly model under a divisive normalisation decision policy

In the last section, I introduced a mathematical reasoning for why learning by depression can be better than potentiation when opposing MBONs are integrated by divisive normalisation.

In this subsection I sought to, first visualize the effect of divisive normalisation on the probability of making the right decision (irrespective of the learning rule) in a different view, with more exploration for the parameter space for  $n$ ,  $k$  and  $Approach(Avoidance)$ . Second, I test the hypothesis of optimality via learning by depression using a MB model with realistic input PN responses. I simulated the learning task (detailed in Chapter 3) using the tuned flies instantiations and inputs in the 6 models (random, homogeneous, activity dependent and independent models) from the previous simulation in Chapter 3 (see Fig.2.5). Then I trained these networks once by depression and another time by potentiation, and calculated their performances in each learning rule using 2 decision making policies: **(a)** softmax policy as in Eq.(3.6) **(b)** divisive normalisation policy similar to equations (3.8) and (3.10).

For the first aim of this subsection, I used arbitrary values to simulate the MBONs firing rates that encodes opposite actions. The MBONs firing rates are varied arbitrarily in a range from 1 to 10, in steps of 0.5. Then, I measured the probability of choosing the  $i$ th action ( $Z_i$ ) using the equation below, which is identical to Eq.(3.8) and (3.10), but using the variables  $Z_i$  and  $Z_j$  in place of  $Approach$  and  $Avoidance$ .

$$\begin{aligned}
 P_i &= 0.5(R_i + 1) \\
 &= 0.5g \left( \frac{(Z_i)^n - (Z_j)^n}{k^n + (Z_i)^n + (Z_j)^n} + 1 \right)
 \end{aligned}
 \tag{3.18}$$

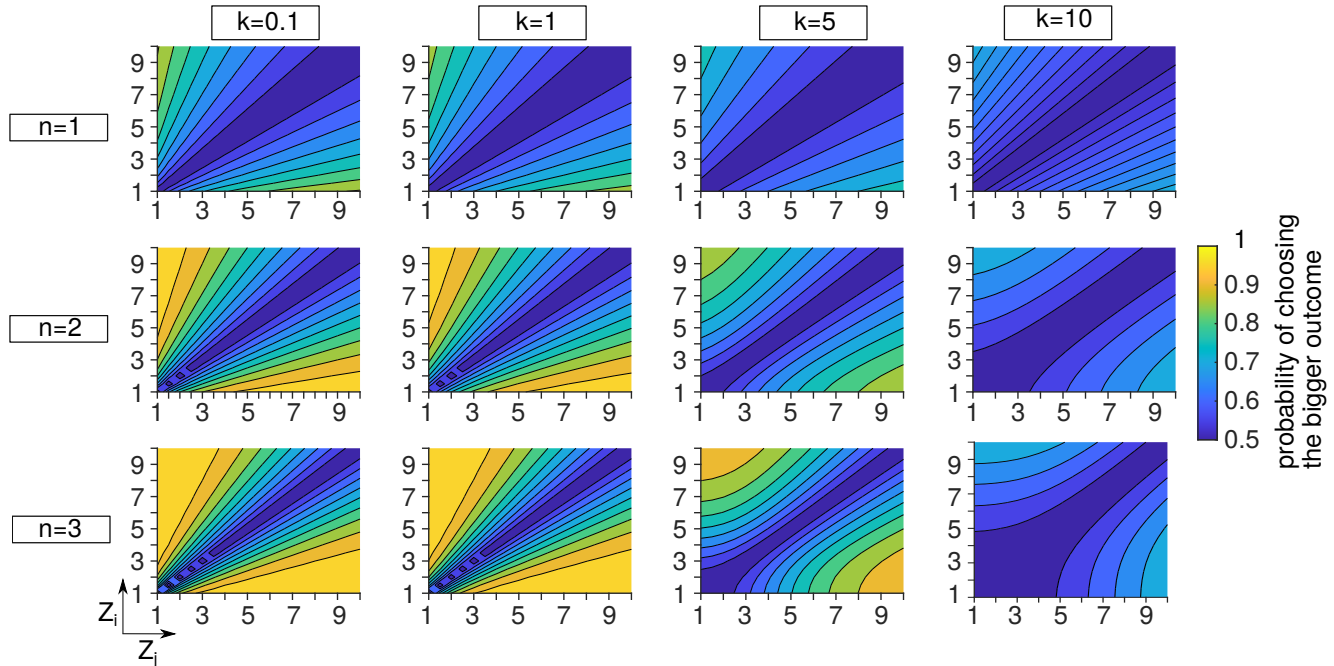
For each pair of action values  $\{Z_i \text{ and } Z_j\}$ , I calculated the probability of making the right action by picking the option with the bigger value  $p$  as follows,

$$p = \begin{cases} P_i, & \text{if } \{Z_i > Z_j\} \\ P_j, & \text{if } \{Z_j > Z_i\} \end{cases}
 \tag{3.19}$$

I explored the accuracy values under different combinations for  $n$  and  $k$ . First,  $n$  varied from 1 to 3 in steps of 1. And  $k$  arbitrarily changes from 0.1 to 10 by factors of 10, however an extra scale of  $k$  was explored when  $k > 1$ , that is  $k = 5$ .

This result in the grid of probability maps in Fig.3.7. An important observation in the figure below, that the ratio of  $\frac{k}{Z_i(\text{or } Z_j)}$  determines how much the decision making policy will behave either as a divisive normalisation function or a soft-max. In the probability maps where  $k = 0.1$ , and given the same gap between both the action values  $\Delta$ , the probability of making the right action (picking the larger outcome) is more if the two action values were small than if they were both large. For instance consider the cases with the same gap  $\Delta = 1$ , the probability

1808 of picking  $Z_i$  over  $Z_j$ , is bigger when their values are 2 and 1 respectively, than if they were 10 and 9. This is  
 1809 true for all  $n$ . This means that for low values of  $k$ , relative to the actions values, the decision making will be more  
 1810 like a divisive normalisation, where the same gap between action values is treated more reliably if their absolute  
 1811 values are low than if they were high. In contrast, if  $k$  is high or comparable to the action values, like the heat  
 1812 maps for  $k = 5$  and  $k = 10$ ,  $k$  dominates the denominator of Eq.(3.18), so the probability of picking the right  
 1813 action is mostly a function of the gap between both values, more like a softmax function. The normalisation effect  
 1814 can be visualised in these probability maps by examining the shape of what is called ‘isolines’. Isolines represent  
 1815 equal values of probabilities, shown by black curves in Figure 3.7. For low values of  $k$  where decision looks more  
 1816 like a divisive normalisation function, these curves diverge in a fan shaped beam from the bottom left corner in a  
 1817 probability map, as shown in the probability maps of low  $k$ ,  $k=0.1$  and 1. In contrast they will run parallel for high  
 1818 values of  $k$ , where the probability is the same for a given gap between action values irrespective of their order of  
 1819 magnitudes, like in the probability maps with  $k = 5$  or 10.



**Figure 3.7:** Toy example to show the effect of the ratio between  $k$  and an option value  $\{Z_i \text{ or } Z_j\}$  on shifting the probability of picking the bigger outcome from a softmax function to a divisive normalisation one. Each pixel in every heat map in this 3x4 grid is the probability of picking the bigger outcome between  $Z_i$  and  $Z_j$ . The heat maps with higher  $k$  and/or  $n$  reveals a decision function that depends only on the gap between the given options values  $[\Delta = Z_i - Z_j]$ . For lower  $k$ ,  $k = 0.1$  or  $k = 1$ , and for the same gap between both options the probability of picking the bigger option grows as the values of  $Z_i$  and  $Z_j$  drops. This is depicted by the isolines divergence in a fan shaped beam from the left bottom corner in these maps. Whereas the isolines in the heat maps of higher  $k$  run in parallel.

1820 **3.4.1 Learning by depression enhances the memory performance under a divisive**  
 1821 **normalisation decision policy and using real odour input responses**

1822 For the second aim of this subsection, I will test the hypothesis that learning by depression is better than potenti-  
 1823 ation using realistic odour input responses. I simulated the mushroom body network with real odour responses in  
 1824 (Hallem and Carlson 2006) and calculated its performance in the memory task using two decision making policies:  
 1825 (a) softmax policy (b) divisive normalisation policy. To implement the divisive normalisation policy I chose a value  
 1826 for  $k$  to be significantly lower than the order of magnitude of the MBONs firing rates,  $k = 0.001$ . For simplicity I  
 1827 set  $n = 1, g = 1$ .

1828  
 1829 For the softmax policy, learning by depression and potentiation are equivalent, as in Fig.3.8. Performances are  
 1830 shown for the different models presented in Chapter 2, the accuracy values shown below are at the peak learning  
 1831 rate in each model. These results agree with the analytical treatment following Eq.(3.6).

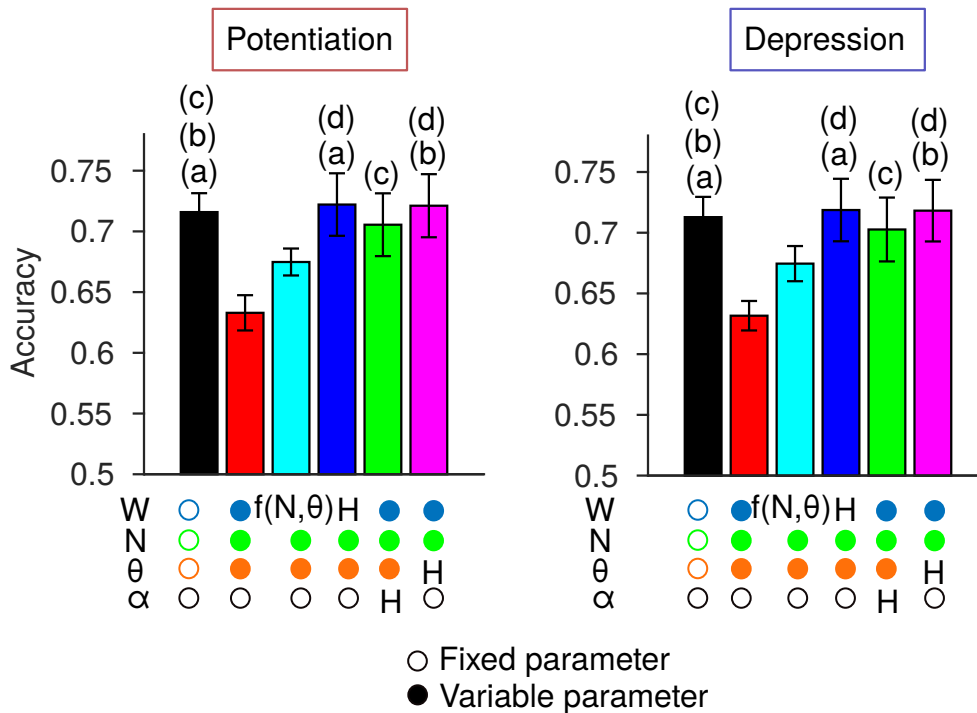
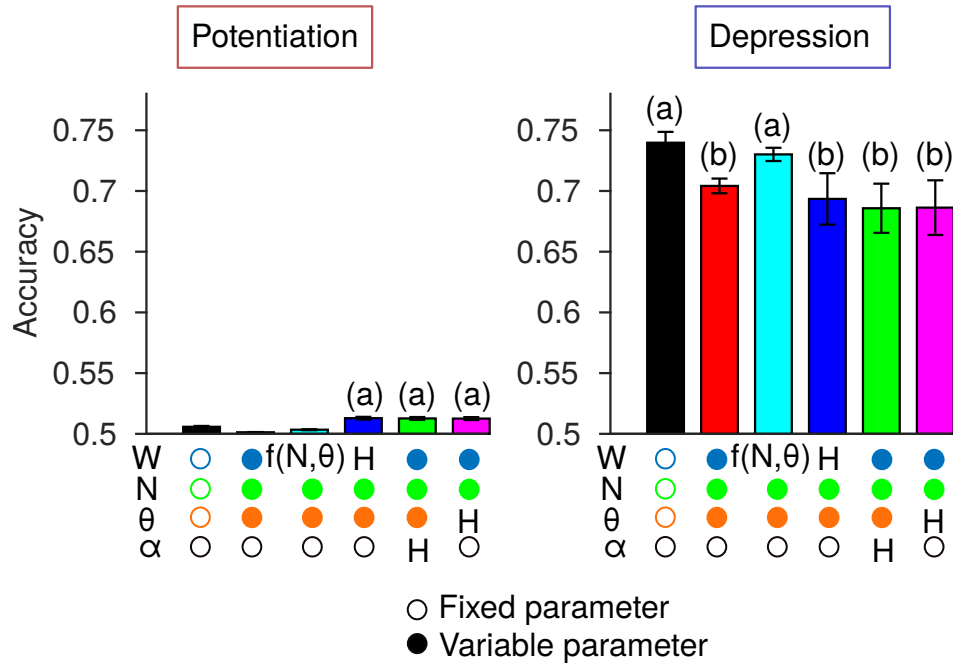


Figure 3.8: Learning by potentiation and depression are equivalent under a softmax policy.

1832 However, and as we suggested, learning by depression is "optimal" under a divisive normalisation policy, as in  
 1833 Fig.3.9.





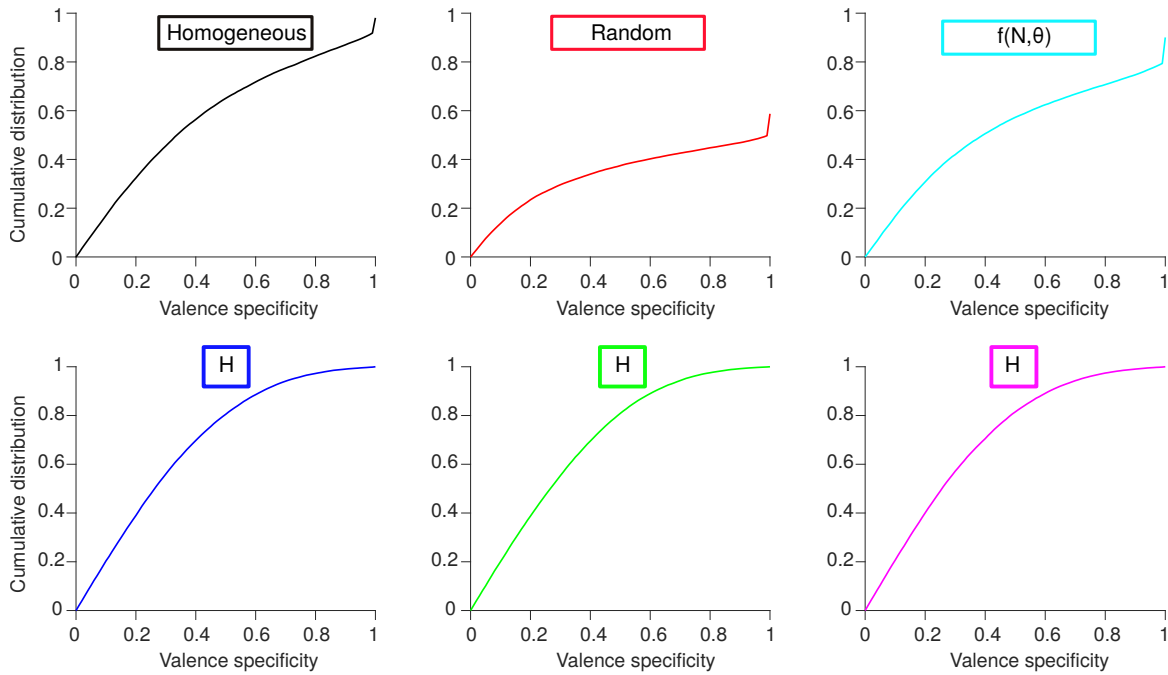
**Figure 3.9:** Learning by depression is better than potentiation when the difference from the opposing MBONs outputs is normalised by their sum.  $n=20$  random fly networks instantiations.

1834 Although the activity-dependent compensatory models have outperformed the other models under learning by  
 1835 depression and potentiation with a softmax decision policy, as well under potentiation and a divisive normalisation  
 1836 decision policy, it was interesting to see here how this pattern has changed under learning by depression and a  
 1837 divisive normalisation decision policy.

1838 Under learning by depression and divisive normalisation decision policy, the activity dependent models had  
 1839 lower memory performances than the random, homogeneous and the activity-independent equalisation models: as  
 1840 the blue, green and magenta bars are smaller than the red, black and cyan bars in Fig.3.9 (and Fig.3.16 in the next  
 1841 section using a winner take all circuit).

1842 I hypothesized that this observed change could be due to the large variance in the KCs valence specificity levels  
 1843 in the random, homogeneous and the activity-independent equalisation models. While it was non-beneficial under  
 1844 a softmax policy, the large variance in the KCs valence specificities (and sparsity levels) will become useful, in  
 1845 particular, under depression and a divisive normalisation decision policy.

1846 KCs in the activity-dependent equalisation models are more similar among each other in their valence specificity  
 1847 values. Although the KCs in these models are still highly sparse, there is less percentage of them that are perfectly  
 1848 specific compared to the random, homogeneous and activity-independent equalisation models. This is shown in  
 1849 Fig.3.10, the cumulative distribution functions (cdf) of the valence specificity in the homogeneous, random and  
 1850 activity independent equalisation models (top panels) are more shifted to the right and have shallower slopes  
 1851 (higher average valence specificity) than the activity dependent equalisation models in the bottom panels in Fig.  
 1852 3.10.



**Figure 3.10:** The cumulative distribution function of valence specificity in the different models. odour inputs and flies networks instantiations ( $n=20$ ) used here are similar to those in Chapter 2 Fig. 2.5.

The random, homogeneous and activity-independent equalisation models (top panels; black, red and cyan curves) have higher average valence specificity values than the activity dependent models in the bottom panels (blue, green, magenta curves). The homogeneous, random and cyan models have few perfectly specific KCs (specificity =1) while the activity dependent equalisation models have none. The sharp rise at the end of the black, red and cyan curves in the top panels account for the mass of the perfectly specific KCs.

1853 Under learning by depression, the highly specific KCs will maintain higher differences between their output  
 1854 weights onto the MBONs. Since the homogeneous, random and activity-independent equalisation models have more  
 1855 perfectly specific KCs than the activity-dependent equalisation models, thus we expect that as learning progresses  
 1856 the homogeneous, random and activity-independent models would have a higher number of KCs with large difference  
 1857 between their MBONs weights. To test this hypothesis, I plotted the cumulative distribution functions (cdf) of the  
 1858 absolute difference between the MBONs weights in each of the 6 different models. Since the difference between the  
 1859 MBONs weights depends on the learning rate, I plotted these cdf plots at different learning rates under each learning  
 1860 rule (depression and potentiation) in Fig. 3.11. As learning progresses, the cdf curves of the absolute difference  
 1861 between the MBONs weights will move towards the upper left corner as in Fig. 3.11. Notably I also found that the  
 1862 peak learning rate (peak performances) is different under each decision making policy, shown by the dashed and  
 1863 solid lines in Fig. 3.11; also as we will see in Fig. 3.12.

1864 Under a divisive normalisation policy and learning by depression and at the peak learning rates (peak per-  
 1865 formances) the KCs in the homogeneous, random and activity-independent equalisation models retained higher  
 1866 absolute differences between their MBONs weights than in the activity-dependent equalisation models, (bold solid  
 1867 curves in top and bottom panels, respectively, under depression in Fig. 3.11). The cdfs of the activity-dependent  
 1868 equalisation models are sharp L shaped curves with the bend near the upper left corner. Hardly any cells have  
 1869 high difference between their MBONs weights, whereas the random, homogeneous, and activity-independent models

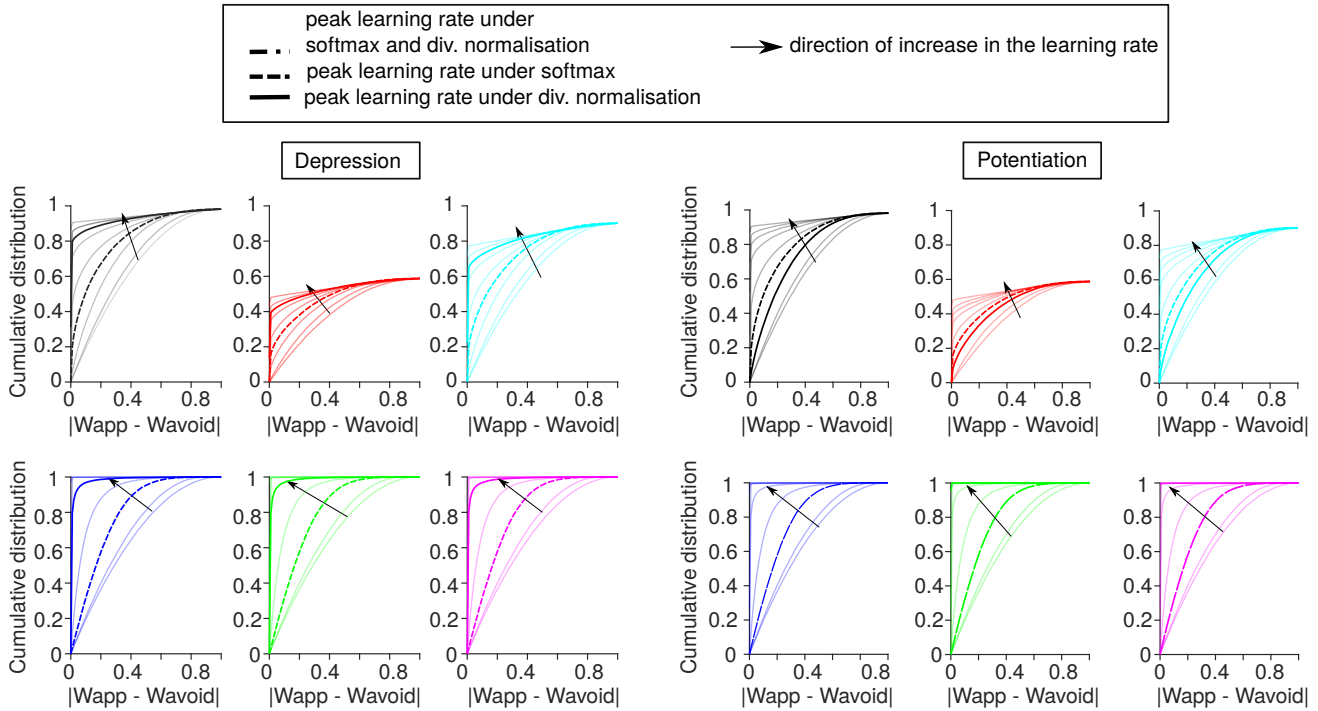
1870 slope more gradually.

1871 In contrast, under a softmax policy, each of the 6 models had the same peak learning rate in learning by  
 1872 depression and potentiation; compare the dashed curves in the homogeneous, random, and activity-independent  
 1873 equalisation models (top panels) in Fig. 3.11 under potentiation and depression, and the dashed and dotted curves  
 1874 in the activity-dependent equalisation models (bottom panels) under potentiation and depression. In fact, under  
 1875 a softmax policy and depression, or softmax policy and potentiation the cdfs of the activity-dependent models  
 1876 (dashed curves under depression and dotted curves under potentiation in bottom panels in Fig. 3.11) are actually  
 1877 a bit shifted right compared to the homogeneous, random, and activity-independent models (dashed curves under  
 1878 depression and potentiation in top panels in Fig. 3.11); that's why the activity-dependent models do better under  
 1879 these conditions. The same trend was found under learning by potentiation and a divisive normalisation policy,  
 1880 shown by the bold solid curves for the homogeneous, random, and activity-independent models (top panels) and  
 1881 dotted curves for the activity-dependent ones (bottom panels) under potentiation in Fig. 3.11.

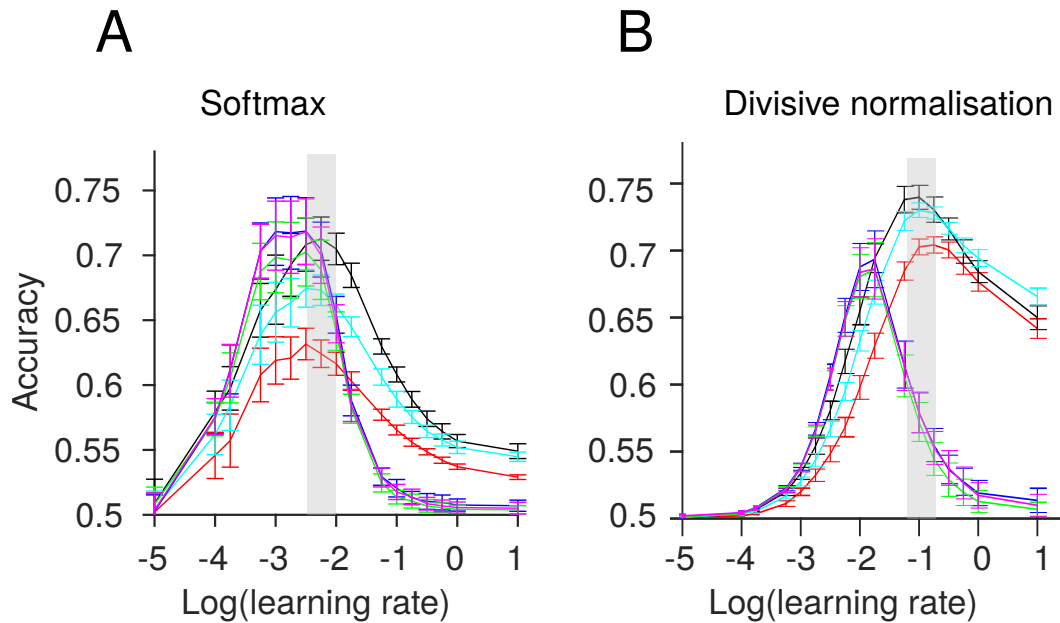
1882 Overall, these results means that the effect of the perfectly specific KCs in the homogeneous, random and  
 1883 activity-independent equalisation models would only be reflected in the models performances under a depression  
 1884 learning rule and a divisive normalisation decision policy. Under these conditions, the more the non-specific and  
 1885 non-sparse KCs are having their output weights silenced, the clearer the right action will become thanks to the  
 1886 responses from the most specific and sparsest neurons. This effect would be enlarged though under a divisive  
 1887 normalisation decision policy, since the differences between action values due to the sparsest and most specific KCs  
 1888 responses will be blown up when normalised by the minute responses from the majority of the non-sparse (or useless)  
 1889 KCs. This is opposed to the case of learning by potentiation (and especially under a divisive normalisation policy),  
 1890 where the higher differences between the opposite action values from the perfectly specific KCs will be diluted when  
 1891 normalised by the high magnitudes of action values from the majority of the non-sparse and non-specific KCs.

1892 This also means that the peak performances under learning by depression and a divisive normalisation policy  
 1893 will be happen at higher learning rates than these under learning by depression and a softmax decision policy; i.e.  
 1894 models performances will increase as the KCs-MBONs weights are updated using bigger incremental steps that is  
 1895 when the non-specific KCs will be silenced the most. Indeed as shown in Fig.3.12, for the same model flies and  
 1896 in learning by depression, the random, homogeneous and activity-independent equalisation models have their peak  
 1897 performances under a divisive normalisation policy (Fig.3.12 B) shifted more to the right, higher learning rates than  
 1898 their peak learning rates under a softmax policy (Fig.3.12 A). Same applies for the activity-dependent equalisation  
 1899 models, not highlighted on the graph.

1900 At the end of this section, we saw that learning by depression was better than potentiation on average across  
 1901 the different model types if MBONs are integrated by divisive normalisation. This motivated us to think of a  
 1902 bio-plausible mechanistic implementation for divisive normalisation in the MBONs. One appealing option is the  
 1903 Winner-take-all (WTA) circuit architecture. WTA competition between a set of  $\{N\}$  options will output the option  
 1904 with the maximum value. The probability of an option to win increases when its value is greatly higher than the



**Figure 3.11:** The cumulative distribution function (cdf) of the differences between the MBONs weights for the different models, plotted for different learning rates and under learning by depression (left half of the figure) and potentiation (the right half). Odor inputs and flies random networks instantiations ( $n=20$ ) are the same as in Chapter 2, Fig. 2.5. The black arrow shows the direction of increase in the learning rate. All models have the same cdf of the differences between their MBONs weights (at their peak learning rates) after learning by depression or potentiation under a softmax policy. This is shown by the dashed curves (black, red and cyan dashed curves) under potentiation and depression for the homogeneous, random and activity-independent equalisation models, and by the dashed and dotted curves for the activity-dependent equalisation models (blue, green, and magenta curves) under depression and potentiation, respectively. The peak learning rates in the activity-dependent equalisation models are roughly the same in learning by potentiation and under a softmax or a divisive normalisation policy (the dotted curves in the bottom panels under potentiation).



**Figure 3.12:** The random, Homogeneous and activity-independent equalisation models have their peak performances at different learning rates under each of the decision making policies: softmax and divisive normalisation. The peak scores in the random, homogeneous and activity-independent equalisation models, shown by the red, black and cyan curves respectively, are shifted more to the right under divisive normalisation (B) than in a softmax decision policy (A),  $n=20$ . The gray shading highlights the region of maximum performance in the red, black and cyan models under each decision rule.

1905 rest of the inputs. Intuitively, this behaviour is similar to the probability of choosing this option under a divisive  
 1906 normalisation function; the higher its value compared to the other options, the higher the probability of picking it  
 1907 under a divisive normalisation decision function. In the next section, I will present the general concept of WTA  
 1908 circuits, the conditions under which they can approximate a divisive normalisation policy, and I will predict the  
 1909 nature of noise in the MBONs that will be required to realize the divisive normalisation policy.

### 1910 3.5 Winner-take-all (WTA) circuit model to approximate divisive nor- 1911 malisation in different neural circuits

1912 Divisive normalisation can give rise to 2 operating regimes based on the values of the input signals. It can either act  
 1913 as an averaging kernel if all the inputs are relatively equal, or as a kernel that find the maximum value among  $N$   
 1914 choices if one of the inputs is significantly higher than the other values. The process of finding the maximum value  
 1915 among a pool of choices is also referred to as max-pooling and sometimes Winner-Take-All competition (WTA).


1916 WTA behaviour is evident in many neural systems that underlie psychophysical decision tasks and attention visual  
 1917 search (Desimone and Duncan 1995; Gold and Shadlen 2007a; Churchland and Ditterich 2012).

1918 WTA dynamics is also responsible for the sparse population responses in the fruit fly's KCs (Stevens 2015),  
 1919 also in the hippocampal place cells (Amaral and Witter 1989; Sreenivasan and Fiete 2011). It is also generally  
 1920 thought to be observed in brain areas with strong inhibition like, the basal ganglia (Bogacz and Gurney 2007; Mink  
 1921 1996; Redgrave, Prescott, and Gurney 1999). In the upcoming sections I predict that WTA competition between  
 1922 the MBONs or between neurons downstream of MBONs can replicate the divisive normalisation behaviour under  
 1923 certain conditions which are: (a) self-excitatory but mutually inhibitory MBONs (b) multiplicative source of noise  
 1924 in the MBONs outputs.

### 1925 Canonical WTA circuit model between the MBONs outputs reproduce the results 1926 from divisive normalisation

1927 I used the canonical WTA architecture from (Coultrip, Granger, and Lynch 1992; Kriener, Chaudhuri, and Fiete  
 1928 2020a; Xie, Hahnloser, and Seung 2002) to simulate the competition between the MBONs outputs. In the canonical  
 1929 WTA circuit there are noisy  $b$  options in  $\{\mathbf{B}\}$  which are sampled by a set of observing neurons,  $z_n$  in  $\{Z_n\}$ . After  
 1930 time  $T$ , the circuit will reach a steady state where the maximum option in  $\{\mathbf{B}\}$ :  $b_{max}$  will be output by its respective  
 1931 observer neuron  $z_{max}$ .

1932 At time step  $(t + \delta t)$  each observer neuron ( $z_n$ ) will sample the input value from the source  $b_n$ , receive excitation  
 1933 from its own state at the previous time step  $t$  (self-excitation), as well it will be inhibited by the rest of the observer  
 1934 neurons  $z_{m,(m \neq n)}$ . This architecture is illustrated as in the figure below,



**Figure 3.13:** Schematic of the WTA canonical circuit

1935 In the case of the mushroom body output neurons there will only be 2 competing options in the set  $\{\mathbf{B}\}$ :  
 1936 Avoidance and Approach. KCs responses for the  $j_{th}$  odour will be integrated by each MBON, we can define the  
 1937 dynamic state of a MBON,  $z_n^j$ , as follows:

$$1938 \quad \tau \frac{dz_n^j}{dt} + z_n^j = Relu(\alpha z_n^j - \beta \sum_{m,(m \neq n)} z_m^j + b_n^j + \eta_n) \quad (3.20)$$

1939 where  $b_n^j$  is the synaptic activation onto the  $n_{th}$  MBON, from the KCs responses to the  $j_{th}$  odour, as

$$1940 \quad b_n^j = \sum_{k=1}^K W_{nk} Y_k^j \quad (3.21)$$

1941 Here  $Relu(.) = max(0, .)$ , a rectifying linear unit which ensures positive firing rates (similar to that used in  
1942 Chapter 2; see Methods), whereas  $\tau$  determines the time constant of the MBONs dynamics. In the simulations,  
1943  $\tau$  was set to 1 on arbitrary time units. The change in the neural state of the  $n_{th}$  MBON,  $\frac{dz_n^j}{dt}$ , is function of the  
1944 strength of the neuron's self-excitation,  $\alpha$ , the strength of lateral inhibition from the other MBON, denoted by  $\beta$ ,  
1945 and the input synaptic activation that it receives from the KCs odour responses,  $b_n^j$ .

1946 In a non-deterministic WTA circuit the external evidences in the channels  $b_n^j$  are noisy time series, with an  
1947 overriding noise process  $\eta_n$ . In reality MBONs synaptic activations can be noisy due to random fluctuations in the  
1948 MBONs firing responses, or stochastic synaptic failures. I modelled the noise process in each MBON's synaptic  
1949 activations using private and statistically identical Ornstein–Uhlenbeck (O-U) processes (Gillespie 1996) given by,

$$1950 \quad \tau_\eta \frac{d\eta_n(t)}{dt} + \eta_n(t) = \sigma_\eta \sqrt{2\tau_\eta} \xi_n(t) \quad (3.22)$$

1951 Where  $\xi_n(t)$  is Gaussian white noise such that:  $\langle \xi_n(t) \rangle = 0$  and  $\langle \xi_n(t) \xi_m(t') \rangle = \delta_{nm} \delta(t - t')$ . For the simulation of  
1952 the O-U process, I referred to the numerical solution of Eq.(3.22) over a time grid with an increment  $\Delta t$ , given in  
1953 (Kriener, Chaudhuri, and Fiete 2020a)(Gillespie 1996), as follows

$$1954 \quad \eta_n(t + \Delta t) = \eta_n(t) e^{-\frac{\Delta t}{\tau_\eta}} + \sigma_\eta \sqrt{1 - e^{-\frac{2\Delta t}{\tau_\eta}}} \xi(t) \quad (3.23)$$

1955 Here  $\tau_\eta$  is the time constant for the O-U noise process, which I set to 0.5 on an arbitrary scale of time units. The  
1956 idea was to set  $\tau_\eta$  to be small enough relative to  $\tau$ , such that the noise process dynamics change fast enough for the  
1957 MBON states to capture that change.  $\sigma_\eta$  though represent the variance of the O-U noise, that is the stochasticity  
1958 of the MBONs' synaptic activations.

1959 I set the initial conditions  $z_n^j(0) = 0$  for all  $n \in \{Approach, Avoidance\}$ . To guarantee WTA dynamics with  
1960 a unique winner and convergence of the circuit to a steady state, I have set the parameters values of  $\alpha$  and  $\beta$  as  
1961 in (Xie, Hahnloser, and Seung 2002; Kriener, Chaudhuri, and Fiete 2020a). First, for the system to converge to a  
1962 steady state, if it exists, the strength of the self excitation has to be less than 1,  $\alpha < 1$ . Second, to guarantee a  
1963 WTA competition with a unique winner ( $\beta$ ) has to be greater than  $(1 - \alpha)$ . These constraints can be explained by  
1964 re-writing Eq.(3.20) in a vector notation for the 2 competing MBONs, in the absence of noise. Thus, I will define  
1965  $\vec{z}^j$ ,  $z^j$ ,  $W$ , and  $\vec{b}^j$  to denote the following: 2x1 vector of the MBONs states, a 2x1 vector of the gradient of MBONs  
1966 states, 2x2 coupling matrix which is function of  $\alpha$  and  $\beta$  (see the equation below), and 2x1 vector of the input  
1967 activation from KCs odour responses, respectively:

$$1968 \quad \tau \vec{z}^j + z^j = Relu(W \vec{z}^j + \vec{b}^j) \quad (3.24)$$

1969 where,  $W$  is,

$$W = \begin{bmatrix} \alpha & -\beta \\ -\beta & \alpha \end{bmatrix} \quad (3.25)$$

The coupling matrix,  $W$  has one eigenvalue  $\lambda_{com} = -1 + \alpha - \beta$  with uniform eigenvector  $\mathbf{1} = [1, 1]$ , and another eigenvalue  $\lambda_{diff} = -1 + \alpha + \beta$  with eigenvector of the difference modes between both MBONs whose entries sum to zero. For the circuit stability, the first eigenvector (uniform vector) has to be decaying as time evolves, this means that  $\lambda_{com}$  should be negative. Therefore, having  $\alpha < 1$ , and since  $\beta > 0$ , will be sufficient for the circuit's convergence to a steady state. To implement the WTA dynamics the lateral inhibition ( $\beta$ ) has to be strong enough to allow for one unique winner. In particular, the second eigenvector which represents the differential modes has to be unstable, that will require  $\lambda_{diff}$  to be positive (Xie, Hahnloser, and Seung 2002), i.e.,

$$\begin{aligned} -1 + \alpha + \beta &> 0 \\ \therefore \beta &> (1 - \alpha) \end{aligned} \quad (3.26)$$

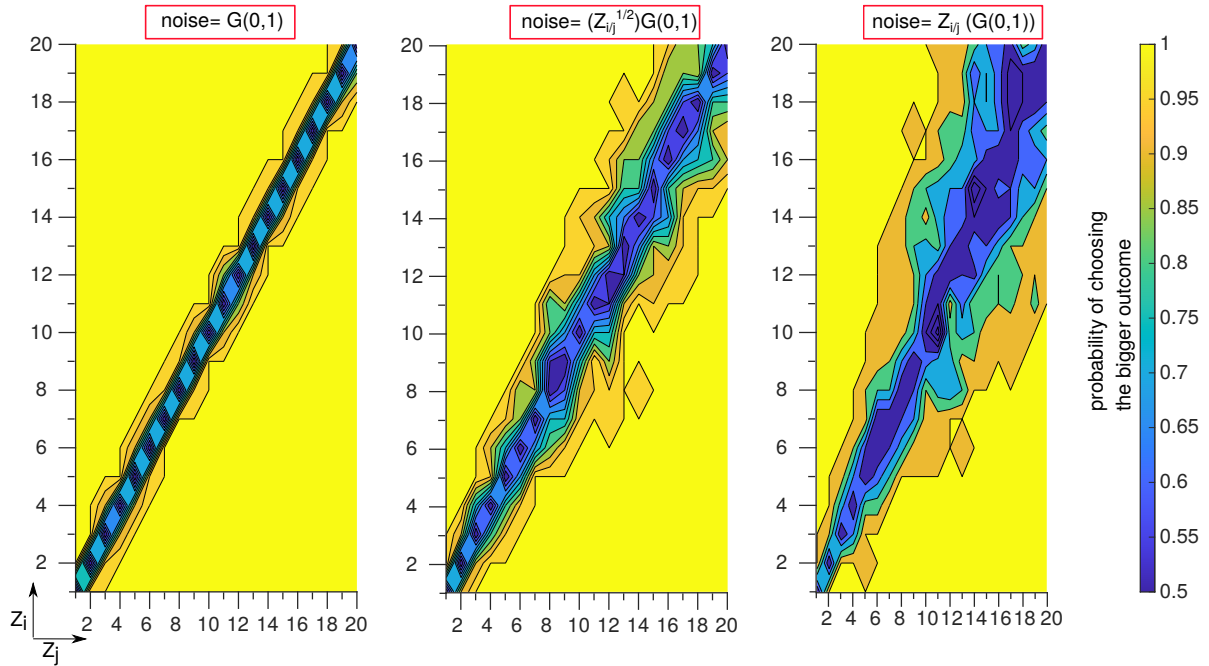
The winner neuron, or more generally the winner neurons group, should converge to an asymptotic value of  $x_{inf} = \frac{b_w}{1-\alpha}$  as was shown in (Xie, Hahnloser, and Seung 2002), where  $b_w$  is the maximum value in the input vector  $b^j$ . In our case this will be the KCs' outputs drive unto the MBON with the highest synaptic activations. Next, I simulated the non-deterministic WTA dynamics in model MBONs. I sought to explore the effect of various types of the noise on the memory performance and on their approximation to the behaviour of divisive normalisation. To do this, I have defined the variance of the O-U noise process  $\sigma_\eta$  as follows,

$$\sigma_{\eta_n}^j = (b_n^j)^\psi \quad (3.27)$$

I varied the variable  $\psi$  such that,  $\psi \in [0, 0.5, 1]$ . When  $\psi$  is 0 the noise exponent ( $\sigma_{\eta_n}^j$ ) in Eq.(3.27) will be 1, and the O-U noise process defined in Eq.(3.23) will be a Gaussian noise with zero mean and unit variance, that is an additive noise independent of the MBONs responses. However, at  $\psi=0.5$  or 1 the noise in the MBONs will be a multiplicative like noise, so the higher is the MBON's output, the more noisy it is.

I simulated a toy example to calculate the probability of choosing the bigger outcome between  $z_i$  and  $z_j$  in a WTA circuit (similar to Fig.3.7). Under each noise regime  $\psi \in \{0, 0.5, 1\}$ , the probability of choosing the bigger outcome was calculated as the average number of successful actions across ( $T=20$ ) random trials where the winner neuron has the biggest value between  $(Z_i, Z_j)$ , as follows:





**Figure 3.14:** Toy example similar to Fig. 3.7 in a WTA circuit to show the effect of the noise on shifting the probability of picking the bigger outcome from a softmax like to a divisive normalisation function. Each pixel in every heat map in this 1x3 grid is the probability of picking the bigger outcome between  $Z_i$  and  $Z_j$ . For each pair of options values, the probability of picking the bigger option is averaged over ( $n=20$ ) random trials. In each random trial (under a certain noise regime) the noise values were randomly added to the WTA circuit's neurons.

The panels show 3 different noise regimes, (left panel): additive Gaussian white noise, (middle panel): Gaussian noise with zero mean and variance equal  $\sqrt{Z_i/z_j}$ , (right panel): Gaussian noise with zero mean and variance equal  $Z_i/z_j$ . The heat map with additive noise (left) reveals a decision function that depends only on the gap between the given options [ $\Delta = Z_i - Z_j$ ]. Similar to the maps in (Fig. 3.7) with high  $k$  values. The multiplicative noise regimes (middle and right) show a behaviour similar to a divisive normalisation function. Compare to the maps in (Fig. 3.7) with low values of  $k$ .

The WTA circuit parameters in this simulation were set as:  $\alpha=0.7$ ,  $\beta=0.9$ ,  $\tau_\eta=0.5$ ,  $\Delta t=\frac{\tau_\eta}{50}$ ,  $\tau=1$ .  $Z_i$  and  $Z_j$  run from 1 to 20 in steps of 1.

$$Acc = \frac{1}{T} \sum_{t=1}^T p_t$$

$$p_t = \begin{cases} 1, & \text{if } \{Z_i > Z_j\} \& \{WTA(winner) = i\} \\ 1, & \text{if } \{Z_j > Z_i\} \& \{WTA(winner) = j\} \\ 0.5, & \text{if } \{Z_i = Z_j\} \end{cases} \quad (3.28)$$

1994

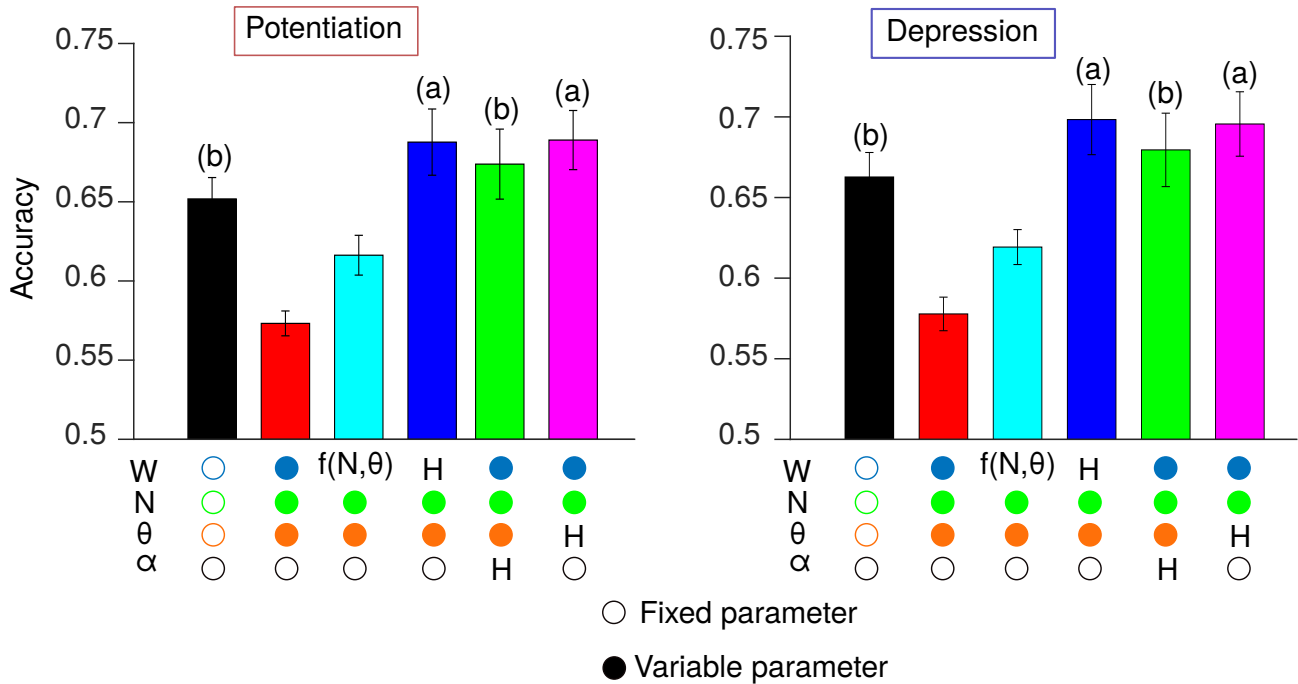
1995 I found that the results under an additive noise ( $\psi = 0$ ), shown in the left panel in Fig. 3.14, were similar  
 1996 to these in Fig.3.7 when  $k \gg a$  (e.g.  $k=5$  and  $a=2$ ) that is when the decision making policy was behaving like  
 1997 a softmax function. In addition, the isolines in the left panel of (Fig.3.14) run in parallel similar to those in the  
 1998 panels on the 3rd and 4th columns of (Fig.3.7) where  $k \gg a$ . This means in the additive noise regime the difference  
 1999 between the smaller and bigger outcome is what determines the winning neuron.

2000 Note that the off-diagonal probability values in the left panel of Fig. 3.14 are significantly higher than those in  
 2001 (Fig. 3.7) when  $k \gg a$ ; the off-diagonal values in the left panel are mostly 1. The reason is that in the WTA example

I set the probability of choosing the right action to 1 if the winner neuron has the bigger value, as in Eq.(3.28). This contrasts with the probability calculation in (Fig. 3.7) using Eq.(3.19) and (3.18), where the probability value is a real number that drops as  $k$  increases in the denominator.

On the other hand, when  $(\psi)$  is 1 or 0.5 the amount of noise in each outcome value (Eq.(3.27)) will be equal to either the outcome value itself or to its square root, respectively. Thus, the variance of the Gaussian noise in Eq.(3.23) will be modulated by a multiplicative factor that is function of the outcome value. Indeed, in this multiplicative noise regime the probability of choosing the bigger outcome will decrease as the values of the outcomes  $(z_i$  or  $z_j)$  increase. This is shown by the divergence of the isolines from the bottom left corners of the middle and right panels in Fig.3.14, a behaviour that is similar to Fig.3.7 when  $k \ll a$  when the decision making function behaves like a divisive normalisation function.

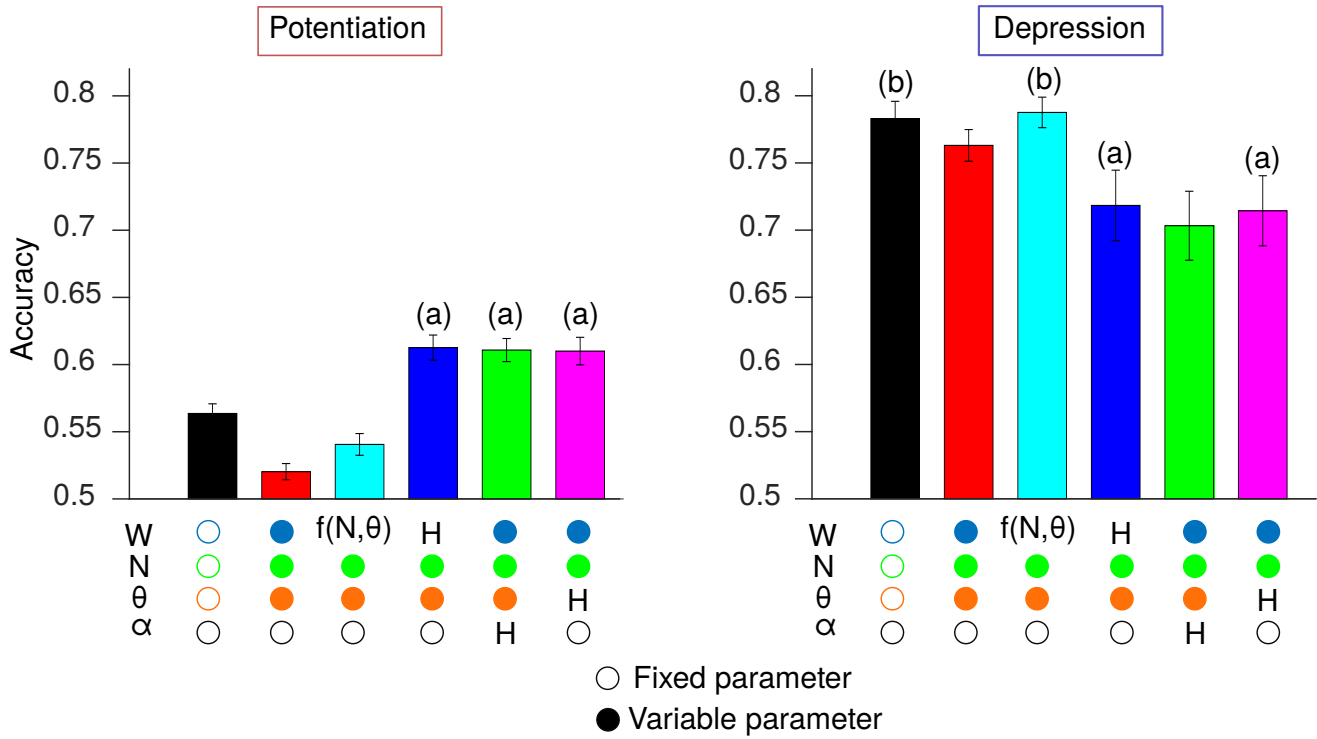
Next, I used the odour inputs and tuned flies instantiations from Chapter 3 in (Fig. 2.5) to test the models performances using a WTA model between the MBONs outputs. Under additive noise regime, i.e.  $\psi=0$ , the models performed the same when learning happened by potentiation and depression, as shown in Fig. 3.15. This agrees with the results obtained earlier under a softmax decision making strategy shown in Fig. 3.8.



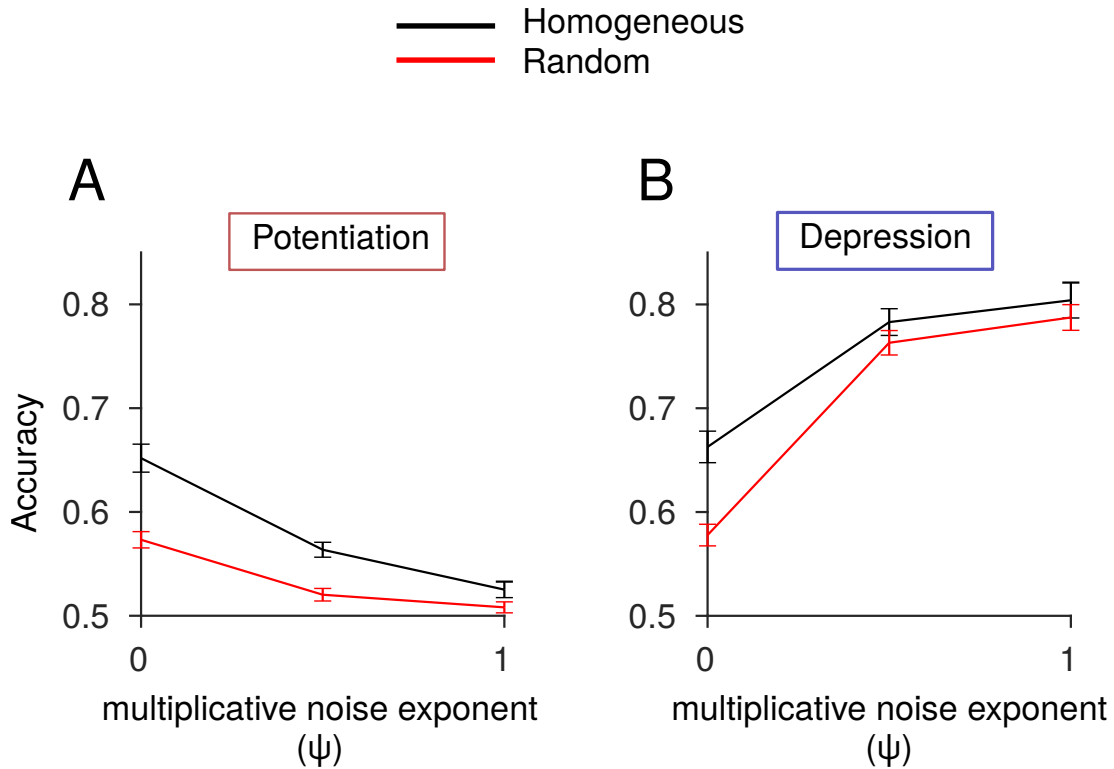
**Figure 3.15:** Learning by depression and potentiation are equivalent in a WTA circuit with additive noise.

In contrast, when the noise was implemented in a multiplicative way the performance under the depression learning rule was significantly higher than that under potentiation (Fig.3.16 for  $(\psi = 0.5)$ ) similar to the results obtained under a divisive normalisation decision function in Fig. 3.9.

Notably, this relation was also true for higher powers of the multiplicative noise as shown in Fig.3.17: like with  $(\psi = 1)$  that is when the noise in the MBON's response is equal to the MBON's synaptic activation.



**Figure 3.16:** performance in learning under depression is higher than in potentiation in a WTA circuit with multiplicative noise. Noise variance is the square root of the MBONs firing rate.



**Figure 3.17:** Performance versus the power of the multiplicative noise in the MBONs in the WTA model. Learning by depression is better than potentiation under various scales of the noise variance. Data shown for the peak performances in two models only (Homogeneous and Random) for simplicity.

## 3.6 Discussion

In this chapter, I presented an analytical and empirical account for the benefits of learning by depression over potentiation using the fruit fly model. In this model, learning by depression outperforms potentiation only under a divisive normalisation decision function; when the difference between the MBONs outputs, that encode opposite behavioural outputs, is normalised by their sum. Divisive normalisation had been observed as a canonical modality in many neural circuits, though its operating behaviour can lead to a Winner-Takes-All (WTA) like behaviour if the inputs vary significantly in their strengths. Since the WTA circuits had also been found in circuits responsible for decision making tasks, this motivated me to use it also here as a bio-plausible implementation for divisive normalisation.

In the WTA model between the MBONs outputs, I found that the noise in the MBONs will need to be of a multiplicative nature in order for the divisive normalisation function to be reproduced.

This creates an interesting avenue to probe the nature of noise in real MBONs. However, this can be challenging as one will have to retrieve the component of noise due to the MBONs only from the other noise sources: noise in the upstream neurons (PNs and KCs) and the imaging devices' noise.

This model aligns with the idea of optimality in the neural circuits development (Costa et al. 2017). Neurons which encode opposite decision values are susceptible to noise and synaptic failures. Thus, the direction of learning (here is depression) might have been a result of natural evolution to optimise the circuit memory performance in the face of this noise. In addition, a long track of behavioural analyses and studies has suggested that humans (and animals) evaluate an option relative to the context of a choice set, i.e. normalised relative to the available choices values. It was found that humans and even small animals like honeybees and gray jays do not assign a fixed value for an option but rather its attractiveness is modulated by presenting an extra choice (Tversky and Simonson 1993; Shafir, Waite, and Smith 2002). Thus, it is appealing to see if context-based value modulation also applies in the fruit fly between the MBONs outputs. It will then become be intriguing to model the motor guiding neural circuit downstream to MBONs.

### Note on originality

Some equations and derivations in this chapter were inspired and done by my supervisor Andrew C.Lin. Andrew came up with the main idea of formulating equations, Eq. 3.7, 3.8, 3.9 and 3.10 to prove that learning by depression is better than potentiation under a divisive normalisation decision policy. These are the equations that used different variables-  $a, b$  and  $c$ - to refer to the magnitudes of unique and overlapping KCs in the MB responses to two different odors.

In addition, Andrew reformulated the previous equations' variables into new parameters set  $-v, p$  and  $x$ - which made deriving equations Eq.3.15, 3.16 and 3.17 easier and more intuitive. He also has carried out the derivations for Eq. 3.15 and 3.16.

## Chapter 4

# General Discussion and Future Work

### 4.1 Compensatory variability rescues the memory performance

In this work I bridged the gap in understanding the link between inter-neuronal variability, homeostasis and its computational benefits for the brain. Variability is embedded in the nature of our brains and the characteristics of its building units, the neurons. Numerous studies have linked the role of homeostasis for the neurons stability with the inter-neuronal variability. In previous work it safeguarded the neuronal intended behaviour across individual animals and within the same animal in face of noisy gene expressions. However, there have been few studies elucidating on the computational consequences of variability among neurons of the same type (same computing nodes in a circuit) within the same neural network. Are there any benefits of inter-neuronal variability? What would happen if neurons of the same type (unrealistically) all had the same intrinsic parameters?

Using the fruit fly model to answer these questions, I found that the memory performance was the highest in an 'unrealistic' model where all KCs had the same intrinsic parameters. The inter-neuronal homogeneity was rather desired over a random variability. I also showed that this aspect of homogeneity can be realised realistically by equalising the average activity levels among KCs whilst maintaining the inherent variability among the KCs parameters.

In this work I suggest that activity independent and homeostatic-like (activity dependent) compensatory models restore the fly's memory performance to the levels of the 'unrealistic' homogeneous model and significantly higher than the random model (Fig.2.5B1,B2). The reason the random variability was undesirable in this memory task is that under sparse coding regime only few percentage of the neurons can be active. In the random model, there were few neurons in the circuit (which had more and/or stronger inputs and low spiking thresholds) that were highly active and fired for any input stimulus, whereas the rest were silent. This defies the very objective of the brain to have sparsely encoding neurons. These neurons are supposed to disentangle the broadly tuned responses from a lower number of upstream neurons (PNs) and encode the odours identities. The utility of such biological normalisation mechanisms was also shown to improve the learning capacity and performance of artificial neural

2079 networks (Shen, Wang, and Navlakha 2021). In (Shen, Wang, and Navlakha 2021), they showed empirical results of  
2080 how the typical normalisation techniques used in artificial neural networks (like the batch normalisation and drop  
2081 out) would enhance the memory capacity and the accuracy of classification as good as the bio-inspired normalisation  
2082 techniques (like synaptic scaling).


2083 In my computational model I used realistic input odour responses and neuronal parameters distributions to build  
2084 a realistic view for the extent to which random variability degrades the performance. Although I explored different  
2085 difficulty levels of the learning task, for e.g., number of input odours or levels of noise imposed on the inputs,  
2086 these conditions cannot exhaust the entire tasks space where random variability could have potentially been less  
2087 undesirable. As an example, the number of silent KCs and variance among the KCs sparsity levels were significantly  
2088 higher in the random model than these in the activity independent and dependent compensatory models (which  
2089 had no idle KCs) (Fig.2.5C). These idle neurons might be beneficial in another scenario like when the distribution of  
2090 rewards/punishments associated with inputs is non-uniform, i.e. some odours are more 'good' ('bad') than others.  
2091 In this instance the idle neurons could be used to encode these extremely rewarding (or punishing) odour inputs.  
2092 In contrast, in the activity dependent models there are no silent neurons, which means there will be no available  
2093 nodes dedicated specially to encode these relatively more important odours.

2094 The connectome data has indeed confirmed correlations among the KCs parameters similar to what I predicted  
2095 in my compensatory models (Fig.2.10 B-I). The correlations observed in the connectome data are inconsistent  
2096 with random sampling, however the orders of the correlations functions were lower than these in my models. In  
2097 particular, the order of correlation between the number of input PNs and the average number of synapses per PN-  
2098 KC connection in the connectome data was of a first order, while the correlation between the number of input PNs  
2099 and the average synaptic weight per PN-KC connection in both the cyan and blue models were of a second order  
2100 (quadratic) (Fig. 2.10 B,C,G,H). This suggests that more than one of these compensatory modes could potentially  
2101 coexist in real KCs.

2102 Although these findings in the connectome strikingly support my models, they remain an approximation because:  
2103 (a) the connectome data is from a single fly ( $n=1$ ) (b) I used the number of PN-KC synapses as an anatomical  
2104 proxy for the KCs input synaptic weights. It will be interesting to see if these correlations (Fig. 2.10B-I) though  
2105 also hold true when EPSPs are recorded in KCs in-vivo.

#### 2106 **4.1.1 Extending compensatory variability models to the cerebellum and other MB-** 2107 **like neural circuits**

2108 The MB circuit structure is quite reminiscent of the cerebellum in mammals, see (Fig.4.1). Cerebellum is the brain  
2109 center responsible of motor learning and movement control. Inside the cerebellum, there are around 200 million  
2110 mossy fiber cells which carry sensory input to the granule cells (more than 50 billion), which is an expansion ratio  
2111 around (200:1) in humans. This is analogous to 150 PNs: 2000 KCs, in flies i.e. expansion ratio (13:1).



**Figure 4.1:** Similarity between the circuit structure of the MB and the cerebellum. Image courtesy of Modi et al.,2020

2112 Granule cells respond to the integrated mossy fibers inputs and relay their outputs to the Purkinje cells. Synaptic  
 2113 plasticity is induced at the outputs of granule cells onto the Purkinje cells when the climbing fibre cell is activated  
 2114 (Albus 1971; Ito 1989; Modi, Shuai, and Turner 2020). This is similar to the memory circuitry in the MB, where  
 2115 KCs integrate inputs from the PNs and relay their output to the MBONs. Similar to the plasticity in the cerebellum,  
 2116 learning in the MB happens by long term depression between the KCs outputs onto the MBONs when the DANs  
 2117 are activated by reward or punishment (US) (Albus 1971; Ito 1989; Aso et al. 2014b). Memory traces formed at  
 2118 the input synapses of the Purkinje cells or MBONs will correct the animal's motor control or bias its behaviour  
 2119 (approach or avoidance) in the future in the presence of the CS, respectively, without the need for the activity in  
 2120 the climbing fibres or the DANs (US) (Aso et al. 2014a; Aso et al. 2014b; Albus 1971).

2121 Given this similarity between the memory circuits in the MB and the cerebellum, could the cerebellar granule  
 2122 cells also show compensatory variability mechanisms similar to the KCs? Cerebellar granule cells are regarded  
 2123 simple and small neurons (Eccles 1967), yet it remained technically challenging to quantify their morphological  
 2124 features due to their small size. A quantitative account of the granule cells morphology and its consequences on  
 2125 their functionality has only been fulfilled recently in (Houston et al. 2017). This study found that granule cells  
 2126 varied among each other in: the distance between the claw-like endings in their dendrites, their axons displacement;  
 2127 which means that in some cases the axon originates from a dendrite rather than from the soma, and what they  
 2128 called the dendrites complexity, which they found to be positively correlated with the distance between the dendrite  
 2129 and axon.

2130 In this study they also measured the effect of these different morphological features on the granule cells intrinsic  
 2131 excitability and hence its functionality. For example, they found that the dendrites complexity, that is defined as  
 2132 the ratio of the total dendrite length (dendrite length to claw + dendrite length within the claw) to the dendrites  
 2133 surface areas, decreases as the dendrites are distant from the axon. This can lead the granule cells to preferentially  
 2134 select for the mossy fibers information arriving close to the axons, because the dendrites would be longer and have  
 2135 bigger claws which can have more number of receptors (Houston et al. 2017). This can be mapped to an inter-  
 2136 granule cells variability in their input synaptic weights where the input synaptic weights closer to the axons are

2137 stronger.

2138 In addition, the granule cells exhibit diversity in their sizes and shapes, which has affected the latency till  
2139 detecting the first action potential evoked by mossy fibres stimulation. It was shown that bigger granule cells with  
2140 longer axons and/or dendrites (and more dendrite branching) are slower in integrating their mossy fibre inputs  
2141 because they will have higher surface area and hence higher membrane capacitance (Houston et al. 2017; Sultan  
2142 2001).

2143 The previous circuit models of the mossy fibres-granule cells has assumed that a granule cell receive its inputs  
2144 from independent mossy fibres rosettes (Sultan 2001; Huang et al. 2013). The morphological treatment in (Houston  
2145 et al. 2017) though has disagreed with this overlooked assumption in the cases of granule cells with branched  
2146 dendrites, they found that the inter-claw distances are smaller in branched dendrites which could allow for multiple  
2147 connections from the granule cell to the same mossy fibre rosette.

2148 These morphological differences can be mapped onto parameters to describe the inter-granule cell variability  
2149 in a similar manner to modelling the inter-KCs variability in Chapter 2. The dendrites branching in granule cells  
2150 could dial up/down the number of independent input mossy fibers, and dendrites complexity can contribute to  
2151 the variability in the inputs synaptic strength and generally to the granule cells intrinsic variability. It will be  
2152 intriguing to apply my computational framework to see the effect of these morphological differences on the memory  
2153 performance in the cerebellum. However, the translation of my framework over to the cerebellum circuit would  
2154 not be straightforward due to many reasons. First, granule cells were recently found to respond less sparsely  
2155 (Giovannucci et al. 2017) than what was envisioned by (Albus 1971) and (Marr 1969). Albus, and Marr before him,  
2156 hypothesized that ideally 1% of granule cells would respond in a given input context, which will increase the learning  
2157 capacity as different granule cells will respond to fine changes in the inputs contexts (Albus 1971; Marr 1969). New  
2158 findings in (Giovannucci et al. 2017; Jörntell and Ekerot 2006; Knogler et al. 2017) have challenged this notion  
2159 about granule cells sparsity, as they found that almost two thirds of the granule cells responded while presenting  
2160 the conditioned stimulus (CS). Thus, it remains controversial whether we can, using finer spatial resolutions, assume  
2161 that granule cells also respond sparsely like KCs.

2162 Second, granule cells were found to encode the conditioned stimulus (CS), as well as predicting the conditioned  
2163 response (CR) (Giovannucci et al. 2017), which is very different from the role of the KCs in the MB. This can be due  
2164 to the closed feedback loop between the cerebellum output and input, the cerebellar nuclei and cortex respectively;  
2165 it is also referred to as nucleocortical feedback (Raymond and Medina 2018; Brandi et al. 2013). Or it can be due  
2166 to descending action-related information from corticopontine pathways converging back onto the granule cells as  
2167 sensory pathways (Huang et al. 2013).

2168 The full neural circuitry required in memory formation differs significantly between the cerebellum and MB,  
2169 which may discourage investigating my models predictions in the cerebellar granular layer. However, various previ-  
2170 ous findings could remarkably support my hypotheses and boost the potential of its happening in the mammalian  
2171 cerebellum. For example, activity dependent plasticity has been observed in the granule cells, where they can alter



2172 their intrinsic excitability properties and hence their mean firing rates in response to tetanic stimulation (Aizenman  
2173 and Linden 2000). In addition, evidences of long term potentiation has been reported between the mossy fibres  
2174 and granule cells in the rat cerebellum (D'Angelo et al. 1999). These findings reveal a range of ongoing plasticity  
2175 mechanisms in the cerebellar granular layer, which can have some computational implications related to enhancing  
2176 the information encoding and memory performance. Indeed, the implications of these plasticity mechanisms has  
2177 been studied in a few computational models. In (Schweighofer, Doya, and Lay 2001), they suggested that plasticity  
2178 in the granular layer is relevant to an unsupervised learning, or a gated activity dependent tuning, of the synaptic  
2179 weights between the mossy fibers and granule cells neurons, which reduced the error in a motor control task. Also in  
2180 (Litwin-Kumar et al. 2017), the classification error of input patterns was minimised when the weights between the  
2181 mossy fibres and the granule cells were modified in an activity dependent manner. In particular, an input weight  
2182 between a given mossy fibre and granule cell was set equal to the inverse square root of the variance in the inputs  
2183 of this mossy fibre.

2184 Although these models used activity dependent methods to model the plasticity between mossy fibres and  
2185 granule cells, they have assumed that these mechanisms are gated by neuromodulatory signals (e.g. serotonin)  
2186 which have not been tested to see their effects on weakening the LTP between mossy fibres and granule cells.

2187 In contrast, in my models I have suggested simpler 'non-gated' activity dependent tuning models, whose roles  
2188 were to reduce the inter-neuronal variability among the expansion layer neurons. The previous modeling studies in  
2189 (Schweighofer, Doya, and Lay 2001) and (Litwin-Kumar et al. 2017) are still of good relevance to my work as they  
2190 provide some computational implications for plasticity in the granular layer which was always thought to have fixed  
2191 (hardwired) input weights. Viewing these plasticity mechanisms in light of the granule cells morphological diversity  
2192 will be the novel contribution suggested from my work.

#### 2193 4.1.2 Compensatory variability in other models

2194 Compensatory variability has been implicitly assumed in previous modelling work without showing its computational  
2195 implications nor explaining the rationale of using it.

2196 (Litwin-Kumar et al. 2017) presented a theoretical proof for the optimality of having 7 PNs inputs for every  
2197 KC, which is the number of average PNs inputs reported in the MB experimental data. They showed that the  
2198 dimensionality (degree of decorrelation) in the KCs responses is maximised at this number. In their work they  
2199 modelled the non-linear KCs outputs using Rectified Linear Unit (ReLU) functions and had picked their spiking  
2200 thresholds such that all KCs will have the same probability of firing,  $f_i$ . This is somewhat similar to equalising  
2201 the KCs lifetime sparsity levels, however we achieve this result as a product of tuning KCs spiking thresholds for  
2202 activity equalisation. Also, I showed that the performance levels after tuning the KCs spiking thresholds for activity  
2203 equalization were higher than when tuned for equalising their firing probability. Interestingly they also used the  
2204 same KCs model neurons to show that the optimal degree of synaptic connectivity between the granule-Purkinje  
2205 cells was indeed equal to 4 which is also the observed value in the cerebellum. This might encourage the chance of

2206 observing compensatory variability also in the cerebellar granule cells.

2207 In another interesting study by (Barak, Rigotti, and Fusi 2013), they studied the trade off between discrim-  
2208 inability and generalization between patterns. In particular they studied how randomly connected neurons (RCNs)  
2209 layer can improve the classification of  $p$  input patterns into their correct output classes. The goal of these RCNs is  
2210 to transform the input patterns such that they become linearly separable. In their model they tuned the spiking  
2211 thresholds in the RCNs transfer functions (Rectified Linear Unit, ReLu) such that all RCNs have the same coding  
2212 level. Their tuning method indeed has achieved the best balance between generalisation and discriminability at  
2213 coding level =10%; like the sparsity levels reported in the KCs. In addition, they have not distinguished between  
2214 population sparseness and lifetime sparseness; rather, they used the term coding level interchangeably to refer to  
2215 the fraction of inputs that activate a given RCN and the fraction of RCNs active on average per input pattern. In  
2216 contrast, in my models I draw a clear discrimination between both quantities. For example, I showed how they  
2217 were very distinct in the random model; some KCs were even silent as they had zero lifetime sparsity whilst the  
2218 population coding level was 10%. I also found that the network performance was better when I tuned the spiking  
2219 thresholds to equalise activity levels than when they were tuned to equalise their firing probabilities, or coding  
2220 levels as defined in [(Barak, Rigotti, and Fusi 2013)].

2221 The computational study in (Barak, Rigotti, and Fusi 2013) and my work has reached similar conclusions but  
2222 from different angles. In their work they found that a coding level of 10% (i.e. every RCN is activated by 10% of the  
2223 input patterns and 10% of the RCNS are active for an input) has reduced the classification error, this coding level  
2224 has optimised the balance between the network discrimination and generalization. In my work though I addressed  
2225 this the other way around. Given the coding level (population sparseness) in the MB network equal 10%, I found  
2226 that equalizing the KC activity levels will enhance the memory performance and reduce, as a by product, the  
2227 variability in the KCs lifetime sparsity levels. The collective conclusions from this study and my work will be that:  
2228 the classification error will be the minimum when the RCNs or the KCs (in my model) have an equalised sparsity  
2229 or activity levels, respectively, and in particular when the network coding level or the population sparseness is 10%.

### 2230 4.1.3 Inter-neuronal variability is beneficial in dense coding regimes

2231 Compensatory variability no longer rescues the memory performance as the coding level increases, see Fig.2.2 and  
2232 2.3. This is because as coding level increases, the random model will have more specialised neurons than the  
2233 homogeneous model and the models with equalised KCs activity levels.

2234 Indeed, some studies have suggested the benefits of randomness and variability among neurons which are known  
2235 to respond densely to input stimuli. In (Tsai et al. 2018), diversity and variability among the lateral inter-neurons  
2236 (LNs) of the fly's antennal lobe were found beneficial to optimise the network's encoding capacity and reliability.  
2237 They found that different sources of variability in the LNs had complementary effects on the network encoding  
2238 capacity and reliability. This study provided an insight into why the local inter-neurons in the antennal lobe  
2239 display morphological variability and randomness (or irregularity) in their connections to other LNs, and why

2240 there are different types (excitatory and inhibitory) of LNs (Tsai et al. 2018). Similarly an intermediate level of  
2241 inter-neuronal variability in the mitral cells (the vertebrate analog to fly PNs) enhanced the encoding capacity and  
2242 reduced the representations redundancy (Tripathy et al. 2013; Padmanabhan and Urban 2010).

## 4.2 Significance of learning by long term depression in the mushroom body and its alike circuits

Understanding the mechanisms underlying motor learning in the mammalian cerebellum has been at the center of interest in many neuroscience studies for decades. Although the cerebellum and the fruit fly MB vary significantly in their sizes and in the nature of their mechanisms for memory formation (Ito 1989; Modi, Shuai, and Turner 2020), memory formation in both circuits happen by long term depression (LTD) (Aso et al. 2014b; Ito 1989). The conservation of the learning rule in the simple MB circuit and the far-ahead developed cerebellum poses an interesting question to ask: Has the LTD learning rule been conserved across species to optimise for some aspects in the learning and cognition abilities? Or has this just happened by chance?

In the cerebellum, an activity burst in the climbing fibres causes an intense depolarization in the postsynaptic Purkinje cell. In particular, the Purkinje cell will respond by a single spike and then pause for around 15-30 msec before it can restore its spontaneous firing rate (Granit and Phillips 1956).

Since then numerous investigations underwent to reveal the synaptic plasticity rule at the Purkinje cells' input synapses. Purkinje cells receive excitatory inputs from the granule cells (parallel fibres) as well the inputs from the climbing fibres. Ahead of any experimental evidences, Marr and Albus have hypothesized the theory behind cerebellar learning, they suggested that activity in the climbing fibre, which carry an error signal in some movement, along with the activation of parallel fibres will induce synaptic plasticity in the active parallel fibres inputs to the Purkinje cells. Whilst Marr suggested this learning to happen by potentiation (Marr 1969), Albus has argued it would rather be by long term depression (LTD) (Albus 1971), which was indeed confirmed decades later in the experiments by Ito (Ito 1989).

In his seminal theoretical treatment, Albus has provided some reasons to support why learning in the granule-Purkinje cells synapses would happen by depression (Albus 1971): "if the Purkinje cell learns to pause when the climbing fibres are active then this can only happen by 'weakening' the excitatory inputs from the granule cells". In addition, he suggested that learning by depression is better for the circuit stability, because if the granule-Purkinje cells synaptic weights are changed by potentiation then they will grow exponentially and eventually reach saturation.

In chapter 4, I accumulate on the evidences that support why learning can happen by depression. In a fly model, I showed that learning by depression lead to a higher memory performance than in potentiation, in specific when the fly's estimation about an odour's value, either good or bad, endures noise of multiplicative nature.

First I show that learning by depression is better than potentiation if the bias in the fly behaviour to approach or avoid an odour depends on the normalised difference between the MBONs firing rates, which encode the respective behaviours, see Fig.3.9.

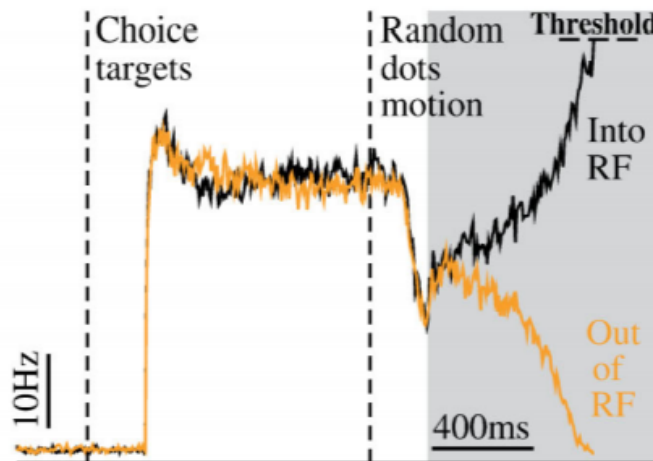
Divisive normalisation has been used to model the responses in visual and sensory neurons. In some studies it was used to explain why neurons behave in a way that is best described by divisive normalisation and how is it necessary for the efficient coding of sensory stimuli and inputs gain control (Schwartz and Simoncelli 2001b; Olsen

2277 and Wilson 2008; Olsen, Bhandawat, and Wilson 2010; Carandini, Heeger, and Movshon 1997; Carandini and  
2278 Heeger 2012). In contrast, I used it here to describe how the fly might be choosing between approaching or avoiding  
2279 a conditioned stimulus. It is important to clarify the difference between how I used divisive normalisation here and  
2280 how it was used before in other models; the divisive normalisation modality I described here does not describe the  
2281 MBONs responses, rather it describes the decision making function that could be implemented via another circuit  
2282 downstream to the MBONs.

2283 Moreover, I showed that a WTA circuit, i.e. neural competition, between the MBONs outputs can be a bio-  
2284 plausible model that will replicate the results from the divisive normalisation model. Compare Fig.3.9 and 3.16,  
2285 only if the noise in the MBONs responses is multiplicative; i.e. under a given stimulus condition, the variance in  
2286 the MBON responses is proportional to its mean firing rate.

2287 WTA circuits have been used previously to model neural correlates of perceptual decision making in the lateral  
2288 intraparietal area in monkeys (KF et al. 2007; Wang 2002; Roitman and Shadlen 2002a). In these models the slow  
2289 NMDA mediated recurrent self excitation and a faster lateral inhibition were key features of the neural dynamics.  
2290 We can use electrophysiological recordings to test the existence of the slow accumulator component of the WTA  
2291 behaviour in a circuit downstream to the MBONs. We will need to record the conditioned odour drive simultaneously  
2292 in opposite valence MBONs or perhaps in downstream neurons, like the Descending Neurons (DNs); (Li et al.,2020)  
2293 have identified some DNs which receive direct connections from the MBONs and can be potential candidates to  
2294 study their responses dynamics. We can then see whether these opposite valence neurons (neurons pools) would  
2295 show slow ramping activities, as well as one of them (or one neurons pool) will have a growing activity while the  
2296 other neuron (or another neurons pool) with the opposite valence will have a diminishing activity, as it is inhibited  
2297 by the winner neuron/s, similar to the behaviour observed in the neurons involved in perceptual decision making  
2298 in mammals, for e.g. the LIP neurons (as shown in Fig.4.2).

2299 Previous drift-diffusion and WTA models have focused on using the neurons' spiking activity, i.e., supra-threshold  
2300 activity, to represent evidence accumulation in competing neurons pools (Gold and Shadlen 2007b; Shadlen and  
2301 Kiani 2013; Kriener, Chaudhuri, and Fiete 2020b). The evidences favoring one decision over the other will be  
2302 accumulated in each opponent pool until one of the opponent neurons/pools will pass the decision making threshold,  
2303 or silence the other neuron and win the competition, resulting in taking this decision. In contrast, some new  
2304 studies(Groschner et al. 2018; Vrontou et al. 2021) have drawn the attention to the existence of sensory evidence  
2305 accumulation in the KCs subthreshold activities. They showed that neurons integrate evidences in their analog  
2306 graded potentials not in their digital spiking activities. For example, (Groschner et al. 2018) found that subsets  
2307 of the  $\alpha\beta_c$  KCs preferentially respond to the increase and decrease in the input odour concentration. These were  
2308 referred to as up/on and down/off cells respectively. The analog evidence accumulation in the on and off  $\alpha\beta_c$  KCs  
2309 was then studied in (Vrontou et al. 2021), where they showed that these opponent KCs pools resemble the perceptual  
2310 neuron-antineuron pools in the drift-diffusion models, with the APL serving the role of lateral inhibition between the  
2311 opponent neurons. In addition, these KCs have all-to-all feedforward connections to downstream neuron-antineuron



**Figure 4.2:** Activity of the winner (loser) LIP neuron - shown in black and orange traces respectively- ramps up (down) as the evidences accumulate to decide the dots motion direction in a visual discrimination task. The duration of the dots motion stimulus presentation is indicated by the gray box. Decision is made by saccade eye movements when the wining neuron reaches the decision threshold shown by the dashed line. Image courtesy of (KF et al. 2007)

2312 pool of MBONs which drive the fly’s behavioural intent. The  $\alpha\beta_c$  KC pool that will spike first then will instruct  
 2313 the fly’s decision by activating its MBON partners.

2314 These findings reveal that flies have neural pathways that share many aspects with the theoretical models for  
 2315 mammalian decision making (Gold and Shadlen 2007b; Roitman and Shadlen 2002b; Latimer et al. 2015). This  
 2316 is motivating for us to look further for other aspects of the mammalian decision making models inside the flies’  
 2317 MB. For instance, it will be interesting to study the evidence integration in the sub-threshold activity of opponent  
 2318 MBONs, or their downstream DNs, or the WTA dynamics in circuits downstream to the MBONs. A few MBONs  
 2319 indeed can inhibit each other (Aso et al. 2014a; Li et al. 2020), however not all the MBONs do; they do not have  
 2320 a candidate to serve the role of global neuron-antineuron inhibitory pool like the APL for the KCs. But we can  
 2321 still investigate the sub-threshold activity in the MBONs that synapse and inhibit one another through direct and  
 2322 indirect feedback loops in (Scheffer et al. 2020; Li et al. 2020).

2323 One of classical models for evidence accumulation is the drift diffusion model (DDM) by (Ratcliff 1978). DDM  
 2324 has successfully explained the observed distribution of choice reaction times in humans and other animals (Wald  
 2325 2004; Gold and Shadlen 2002; Hanes and Schall 1996; Shadlen and Newsome 2001). In this model the evidence  
 2326 that support one choice over another is accumulated at a constant rate till it reaches a fixed threshold value where  
 2327 the decision will be made (Bogacz et al. 2006; Ratcliff 1978). Although, this model did not define the value of  
 2328 this threshold, other work have recently suggested models of how animals can vary these threshold levels optimally

2329 to minimise their energy costs (for e.g. attentional efforts) and make accurate decisions (Drugowitsch et al. 2012;  
2330 Milosavljevic et al. 2010). In addition, other physiologically inspired models were suggested to implement the  
2331 DDM using two or more evidence integrators (Wang 2002; Usher and McClelland 2001; Bogacz et al. 2006). One  
2332 can test the existence of the evidence accumulation in the MBONs or in their downstream neurons by measuring  
2333 the spiking thresholds in these neurons to see if they all reach an almost equal threshold level before spiking or  
2334 not. Another interesting test would be to look for the mutual inhibition or pooled inhibition circuit motifs in  
2335 the MBONs downstream areas. Recent work in (Keung, Hagen, and Wilson 2020) have suggested that humans  
2336 might accumulate evidences in a perceptual auditory task using a divisive normalisation circuit motif. They showed  
2337 that the integration kernel emerging from the divisive normalisation circuit explains well the uneven weighing of  
2338 evidences in the data. This study would encourage a new testable prediction, that is one could record the MBONs  
2339 sub-threshold activities to see if they exhibit this uneven weighing (a bump kernel) of the KCs inputs spikes, such  
2340 that the later and early input spikes from the KCs weigh less than the mid ones.

2341 Last but not least, in my model I find that learning by depression is better than potentiation if the MBONs  
2342 responses have multiplicative noise, i.e. the variance in the MBONs responses is proportional to their mean firing  
2343 rates. This multiplicative relationship between the responses variances and means has been shown before in the  
2344 cortical neurons (Tolhurst, Movshon, and Dean 1983; Carandini, Heeger, and Movshon 1997). The responses  
2345 variability in these neurons were of a Poisson like nature. They found that the variances in the visual cortex  
2346 neurons responses (in cats and monkeys) were proportional to their means; a least squares regression on this data  
2347 yielded a slope  $\approx 1$  (Tolhurst, Movshon, and Dean 1983). However, no similar studies exist to the date of this thesis  
2348 which confirm such relationship, neither in the MBONs nor in its analogs (like the Purkinje cells in the cerebellum).  
2349 This could be another potential testable avenue, if the variance in the MBONs noise will be found multiplicative,  
2350 then this work will be a another evidence to support the case of learning by depression.

### 2351 **4.3 Possible extensions and improvements to the models**

2352 The variability among KCs in their number of PN inputs vary between the different KCs sub-types (Caron et al.  
2353 2013), also see (Fig.2.10F). For simplicity I ignored modeling the KCs sub-types (and their associated distributions  
2354 of the number of input PNs). It will be interesting to see the effect of adding this extra level of detail on the  
2355 performance results across the different models, and if a certain compensatory model will outperform the others.

2356 In my models I predicted stronger levels of correlations between the KCs intrinsic parameters than these observed  
2357 in the connectome, yet the nature of the correlations were the same both in my models and the connectome as  
2358 depicted in (Fig.2.10 B-I). One possible explanation is that there might be more than one mode of compensation  
2359 occurring at the same time in real KCs, whereas I have only simulated only one at a time. It will be insightful  
2360 to either simulate all the compensatory models simultaneously, or feedback the parameters correlations strengths  
2361 observed in the connectome data into my models and see how the results will change.

2362 In addition to the inter-KCs variability in connectivity parameters, KCs also vary in their somas shapes, neu-

2363 rites lengths and dendrites branching patterns. Similar to (Houston et al. 2017), it is compelling to add these  
2364 morphological differences to the sources of inter-KCs variability that I used in my models. On that account, we can  
2365 build joint probability distributions for the KCs connectivity parameters (number of their inputs PNs, their spiking  
2366 thresholds and input synaptic weights) and their morphological features (neurite lengths, somas size, degree of the  
2367 dendritic branching). This can give a better understanding of which feature among these morphological differences  
2368 that can be the most detrimental for the memory performance.

2369 In this work, the fly behaviour was determined either by the difference or the normalised difference between  
2370 the MBONs firing rates, where they encode opposite behaviours. But how can the neurons downstream to the  
2371 MBONs affect the memory performance? MBONs send their axons to 4 major neuropils and to the Lateral Horn  
2372 (LH), which is hardwired for innate olfactory associated behaviours. In addition, some MBONs send feed forward  
2373 signals inside the MB lobes (see Fig.1.6) and to some DANs too. In extensions of this work, we can model the  
2374 downstream circuits involved in decision making and some of the motor control neurons, Descending Neurons (DNs)  
2375 which receive direct inputs from the MBONs. In addition, we can add the MBON-MBON and MBON-DAN feed  
2376 forward and feedback connections to see if they provide the neural pathways for the normalisation or inhibitory  
2377 pool in the WTA and drift-diffusion models.



# Bibliography

- 2379 Achard, Pablo and Erik DeSchutter (2006). “Complex Parameter Landscape for a Complex Neuron Model”. In:  
2380 *PLOS Computational Biology*. DOI: <https://doi.org/10.1371/journal.pcbi.0020094>.
- 2381 Aiello, Leslie C and Peter Wheeler (1995). “The expensive-tissue hypothesis: the brain and the digestive system in  
2382 human and primate evolution”. In: *Current anthropology* 36.2, pp. 199–221.
- 2383 Aizenman, Carlos D. and David J. Linden (2000). “Rapid, synaptically driven increases in the intrinsic excitability  
2384 of cerebellar deep nuclear neurons”. In: *Nature Neuroscience* 3, pp. 109–111.
- 2385 Albrecht, Duane G and Wilson S Geisler (1991). “Motion selectivity and the contrast-response function of simple  
2386 cells in the visual cortex”. In: *Visual neuroscience* 7.6, pp. 531–546.
- 2387 Albus, James S (Feb. 1971). “A theory of cerebellar function”. In: *Mathematical Biosciences* 10.1-2, pp. 25–61.
- 2388 Amaral, David G and Menno P Witter (1989). “The three-dimensional organization of the hippocampal formation:  
2389 a review of anatomical data”. In: *Neuroscience* 31.3, pp. 571–591.
- 2390 Amin, Hoger and Andrew C. Lin (2019). “Neuronal mechanisms underlying innate and learned olfactory processing  
2391 in *Drosophila*”. In: *Current Opinion in Insect Science* 36, pp. 9–17. DOI: [https://doi.org/10.1016/j.cois.](https://doi.org/10.1016/j.cois.2019.06.003)  
2392 2019.06.003.
- 2393 Amin, Hoger et al. (2020). “Localized inhibition in the *Drosophila* mushroom body”. In: *eLife* 9, e56954.
- 2394 Apostolopoulou, Anthi A and Andrew C. Lin (July 2020). “Mechanisms underlying homeostatic plasticity in the  
2395 *Drosophila* mushroom body in vivo.” In: *Proceedings of the National Academy of Sciences of the United States*  
2396 *of America* 117.28, pp. 16606–16615.
- 2397 Ardin, Paul et al. (Feb. 2016). “Using an Insect Mushroom Body Circuit to Encode Route Memory in Complex  
2398 Natural Environments”. In: *PLOS Computational Biology* 12.2, pp. 1–22. DOI: 10.1371/journal.pcbi.1004683.  
2399 URL: <https://doi.org/10.1371/journal.pcbi.1004683>.
- 2400 Aso, Yoshinori et al. (2010). “Specific dopaminergic neurons for the formation of labile aversive memory”. In:  
2401 *Current biology* 20.16, pp. 1445–1451.
- 2402 Aso, Yoshinori et al. (2014a). “Mushroom body output neurons encode valence and guide memory-based action  
2403 selection in *Drosophila*”. In: *eLife* 3:e04580. DOI: <https://doi.org/10.7554/eLife.04580.001>.
- 2404 Aso, Yoshinori et al. (2014b). “The neuronal architecture of the mushroom body provides a logic for associative  
2405 learning”. In: *elife* 3, e04577.

- 2406 Aso, Yoshinori et al. (Nov. 2019). “Nitric oxide acts as a cotransmitter in a subset of dopaminergic neurons to  
2407 diversify memory dynamics”. In: *eLife* 8, e49257.
- 2408 Awata, Hiroko et al. (2019). “The neural circuit linking mushroom body parallel circuits induces memory consoli-  
2409 dation in *Drosophila*”. In: *Proceedings of the National Academy of Sciences* 116.32, pp. 16080–16085.
- 2410 Babadi, Baktash and Haim Sompolinsky (2014). “Sparseness and Expansion in Sensory Representations”. In: *Neuron*  
2411 83.5, pp. 1213–1226. ISSN: 0896-6273. DOI: <https://doi.org/10.1016/j.neuron.2014.07.035>. URL: <https://www.sciencedirect.com/science/article/pii/S0896627314006461>.
- 2412
- 2413 Barak, Omri, Mattia Rigotti, and Stefano Fusi (2013). “The Sparseness of Mixed Selectivity Neurons Controls the  
2414 Generalization–Discrimination Trade-Off”. In: *Journal of Neuroscience* 33.9, pp. 3844–3856. ISSN: 0270-6474.  
2415 DOI: 10.1523/JNEUROSCI.2753-12.2013. eprint: <https://www.jneurosci.org/content/33/9/3844.full.pdf>.  
2416 URL: <https://www.jneurosci.org/content/33/9/3844>.
- 2417 Barnes, Christopher L, Daniel Bonnerly, and Albert Cardona (Oct. 2020). “Synaptic counts approximate synaptic  
2418 contact area in *Drosophila*”. In: *bioRxiv*, p. 2020.10.09.333187.
- 2419 Barton, Robert A. and Chris Venditti (2014). “Rapid Evolution of the Cerebellum in Humans and Other Great  
2420 Apes”. In: *Current Biology* 24.20, pp. 2440–2444. ISSN: 0960-9822. DOI: <https://doi.org/10.1016/j.cub.2014.08.056>.  
2421 URL: <https://www.sciencedirect.com/science/article/pii/S0960982214010690>.
- 2422 Bauman, L.A. and A.B. Bonds (Feb. 1991). “Inhibitory refinement of spatial frequency selectivity in single cells of  
2423 the cat striate cortex”. In: *Vision research* 31, pp. 933–44. DOI: 10.1016/0042-6989(91)90201-F.
- 2424 Bhandawat, Vikas et al. (2007). “Sensory Processing in the *Drosophila* Antennal Lobe Increases the Reliability and  
2425 Separability of Ensemble Odor Representations”. In: *Nature Neuroscience* 10.11, pp. 1474–1482. DOI: 10.1038/  
2426 nn1976.
- 2427 Bliss, Tim VP and Graham L Collingridge (1993). “A synaptic model of memory: long-term potentiation in the  
2428 hippocampus”. In: *Nature* 361.6407, pp. 31–39.
- 2429 Bogacz, Rafal and Kevin Gurney (2007). “The basal ganglia and cortex implement optimal decision making between  
2430 alternative actions”. In: *Neural computation* 19.2, pp. 442–477.
- 2431 Bogacz, Rafal et al. (2006). “The physics of optimal decision making: a formal analysis of models of performance  
2432 in two-alternative forced-choice tasks.” In: *Psychological review* 113.4, p. 700.
- 2433 Bolshakov, Vadim Y. and Steven A. Siegelbaum (1995). “Regulation of Hippocampal Transmitter Release During  
2434 Development and Long-Term Potentiation”. In: *Science* 269.5231, pp. 1730–1734. DOI: 10.1126/science.  
2435 7569903. eprint: <https://www.science.org/doi/pdf/10.1126/science.7569903>. URL: <https://www.science.org/doi/abs/10.1126/science.7569903>.
- 2436
- 2437 Boynton, Robert M. and David N. Whitten (1970). “Visual Adaptation in Monkey Cones: Recordings of Late  
2438 Receptor Potentials”. In: *Science* 170.3965, pp. 1423–1426. DOI: 10.1126/science.170.3965.1423. eprint:  
2439 <https://www.science.org/doi/pdf/10.1126/science.170.3965.1423>. URL: <https://www.science.org/doi/abs/10.1126/science.170.3965.1423>.
- 2440

- 2441 Brandi, Santiago et al. (2013). “Learning of motor sequences based on a computational model of the cerebellum”.  
2442 In: *Conference on Biomimetic and Biohybrid Systems*. Springer, pp. 356–358.
- 2443 Brunel, Nicolas et al. (2004). “Optimal information storage and the distribution of synaptic weights: perceptron  
2444 versus Purkinje cell”. In: *Neuron* 43.5, pp. 745–757.
- 2445 Burke, Christopher J et al. (2012). “Layered reward signalling through octopamine and dopamine in *Drosophila*”.  
2446 In: *Nature* 492.7429, pp. 433–437.
- 2447 Burkhardt, Dwight (Mar. 1994). “Light Adaptation and Photopigment Bleaching in Cone Photoreceptors in situ in  
2448 the Retina of the Turtle”. In: *The Journal of neuroscience : the official journal of the Society for Neuroscience*  
2449 14, pp. 1091–105. DOI: 10.1523/JNEUROSCI.14-03-01091.1994.
- 2450 Burrell, Brian D. and Christie L. Sahley (2004). “Multiple Forms of Long-Term Potentiation and Long-Term De-  
2451 pression Converge on a Single Interneuron in the Leech CNS”. In: *Journal of Neuroscience* 24.16, pp. 4011–4019.  
2452 ISSN: 0270-6474. DOI: 10.1523/JNEUROSCI.0178-04.2004. eprint: [https://www.jneurosci.org/content/24/](https://www.jneurosci.org/content/24/16/4011.full.pdf)  
2453 [16/4011.full.pdf](https://www.jneurosci.org/content/24/16/4011). URL: <https://www.jneurosci.org/content/24/16/4011>.
- 2454 Buzsáki, György and Kenji Mizuseki (Apr. 2014). “The log-dynamic brain: how skewed distributions affect network  
2455 operations”. In: *Nature Reviews Neuroscience* 15.4, pp. 264–278.
- 2456 Caporale, Natalia and Yang Dan (2008). “Spike timing–dependent plasticity: a Hebbian learning rule”. In: *Annu.*  
2457 *Rev. Neurosci.* 31, pp. 25–46.
- 2458 Carandini, Matteo and David J Heeger (2012). “Normalization as a canonical neural computation”. In: *Nature*  
2459 *Reviews Neuroscience* 13.1, pp. 51–62.
- 2460 Carandini, Matteo, David J Heeger, and J Anthony Movshon (1997). “Linearity and normalization in simple cells  
2461 of the macaque primary visual cortex”. In: *Journal of Neuroscience* 17.21, pp. 8621–8644.
- 2462 Caron, Sophie et al. (2013). “Random Convergence of Olfactory Inputs in the *Drosophila* Mushroom Body”. In:  
2463 *Nature* 497.7447, pp. 113–117. DOI: 10.1038/nature12063.
- 2464 Castellucci, V and ER Kandel (1976). “Presynaptic facilitation as a mechanism for behavioral sensitization in  
2465 *Aplysia*”. In: *Science* 194.4270, pp. 1176–1178.
- 2466 Castellucci, Vincent F and Eric R Kandel (1974). “A quantal analysis of the synaptic depression underlying ha-  
2467 bituation of the gill-withdrawal reflex in *Aplysia*”. In: *Proceedings of the National Academy of Sciences* 71.12,  
2468 pp. 5004–5008.
- 2469 Chialvo, D R and P Bak (1999). “Learning From Mistakes”. In: *Neuroscience* 90.4, pp. 1137–48. DOI: 10.1016/  
2470 [s0306-4522\(98\)00472-2](https://doi.org/10.1016/S0306-4522(98)00472-2).
- 2471 Churchland, Anne K and Jochen Ditterich (2012). “New advances in understanding decisions among multiple  
2472 alternatives”. In: *Current opinion in neurobiology* 22.6, pp. 920–926.
- 2473 Claridge-Chang, Adam et al. (2009). “Writing memories with light-addressable reinforcement circuitry”. In: *Cell*  
2474 139.2, pp. 405–415.

- 2475 Cohn, Raphael, Ianessa Morantte, and Vanessa Ruta (2015). “Coordinated and Compartmentalized Neuromodu-  
2476 lation Shapes Sensory Processing in *Drosophila*”. In: *Cell* 163.7, pp. 1742–1755. DOI: [https://doi.org/10.](https://doi.org/10.1016/j.cell.2015.11.019)  
2477 [1016/j.cell.2015.11.019](https://doi.org/10.1016/j.cell.2015.11.019).
- 2478 Costa, Rui Ponte et al. (2017). “Synaptic Transmission Optimization Predicts Expression Loci of Long-Term Plas-  
2479 ticity”. In: *Neuron* 96.1, 177–189.e7. ISSN: 0896-6273. DOI: <https://doi.org/10.1016/j.neuron.2017.09.021>.  
2480 URL: <https://www.sciencedirect.com/science/article/pii/S0896627317308619>.
- 2481 Coultrip, Robert, Richard Granger, and Gary Lynch (1992). “A cortical model of winner-take-all competition via  
2482 lateral inhibition”. In: *Neural networks* 5.1, pp. 47–54.
- 2483 Couto, Africa, Mattias Alenius, and Barry J Dickson (2005). “Molecular, anatomical, and functional organization  
2484 of the *Drosophila* olfactory system”. In: *Current Biology* 15.17, pp. 1535–1547.
- 2485 Crittenden, Jill R et al. (1998). “Tripartite mushroom body architecture revealed by antigenic markers”. In: *Learning*  
2486 *& memory* 5.1, pp. 38–51.
- 2487 Dan, Yang and Mu-Ming Poo (2006). “Spike Timing-Dependent Plasticity: From Synapse to Perception”. In: *Phys-*  
2488 *iological Reviews* 86.3. PMID: 16816145, pp. 1033–1048. DOI: [10.1152/physrev.00030.2005](https://doi.org/10.1152/physrev.00030.2005). eprint: <https://doi.org/10.1152/physrev.00030.2005>. URL: <https://doi.org/10.1152/physrev.00030.2005>.
- 2490 D’Angelo, Egidio et al. (1999). “Evidence for NMDA and mGlu receptor-dependent long-term potentiation of mossy  
2491 fiber–granule cell transmission in rat cerebellum”. In: *Journal of neurophysiology* 81.1, pp. 277–287.
- 2492 Delvendahl, Igor, Katarzyna Kita, and Martin Müller (Nov. 2019). “Rapid and sustained homeostatic control of  
2493 presynaptic exocytosis at a central synapse.” In: *Proceedings of the National Academy of Sciences of the United*  
2494 *States of America* 116.47, pp. 23783–23789.
- 2495 Desimone, Robert and John Duncan (1995). “Neural Mechanisms of Selective Visual Attention”. In: *Annual Review*  
2496 *of Neuroscience* 18.1. PMID: 7605061, pp. 193–222. DOI: [10.1146/annurev.ne.18.030195.001205](https://doi.org/10.1146/annurev.ne.18.030195.001205). eprint:  
2497 <https://doi.org/10.1146/annurev.ne.18.030195.001205>. URL: [https://doi.org/10.1146/annurev.ne.](https://doi.org/10.1146/annurev.ne.18.030195.001205)  
2498 [18.030195.001205](https://doi.org/10.1146/annurev.ne.18.030195.001205).
- 2499 Drugowitsch, Jan et al. (2012). “The cost of accumulating evidence in perceptual decision making”. In: *Journal of*  
2500 *Neuroscience* 32.11, pp. 3612–3628.
- 2501 Dufresne, Daniel (2008). “Sums of lognormals”. In: *Actuarial Research Conference*, pp. 1–6.
- 2502 Eccles, John C (1967). “Circuits in the cerebellar control of movement.” In: *Proceedings of the National Academy*  
2503 *of Sciences of the United States of America* 58.1, p. 336.
- 2504 Eichler, Katharina et al. (Aug. 2017). “The complete connectome of a learning and memory centre in an insect  
2505 brain.” In: *Nature* 548.7666, pp. 175–182.
- 2506 Farris, Sarah M (July 2011). “Are mushroom bodies cerebellum-like structures?” In: *Arthropod Structure and*  
2507 *Development* 40.4, pp. 368–379.
- 2508 Fenton, L.F. (1960). “The Sum of Log-Normal Probability Distributions In Scattered Transmission Systems”. In:  
2509 *IRE Trans. Commun. Systems* 8, pp. 57–67.

- 2510 Fishilevich, Elane and Leslie B Vosshall (2005). “Genetic and functional subdivision of the *Drosophila* antennal  
2511 lobe”. In: *Current Biology* 15.17, pp. 1548–1553.
- 2512 Földiák, Peter and Malcolm P Young (1995). “Sparse coding in the primate cortex”. In:
- 2513 Freeman, Tobe et al. (Sept. 2002). “Suppression without Inhibition in Visual Cortex”. In: *Neuron* 35, pp. 759–71.  
2514 DOI: 10.1016/S0896-6273(02)00819-X.
- 2515 Gao, Qian, Bingbing Yuan, and Andrew Chess (2000). “Convergent projections of *Drosophila* olfactory neurons to  
2516 specific glomeruli in the antennal lobe”. In: *Nature neuroscience* 3.8, pp. 780–785.
- 2517 Gaudry, Quentin et al. (Jan. 2013). “Asymmetric neurotransmitter release enables rapid odour lateralization in  
2518 *Drosophila*.” In: *Nature* 493.7432, pp. 424–428.
- 2519 Gillespie, Daniel T (1996). “Exact numerical simulation of the Ornstein-Uhlenbeck process and its integral”. In:  
2520 *Physical review E* 54.2, p. 2084.
- 2521 Giovannucci, Andrea et al. (2017). “Cerebellar granule cells acquire a widespread predictive feedback signal during  
2522 motor learning”. In: *Nature neuroscience* 20.5, pp. 727–734.
- 2523 Gjorgjieva, Julijana, Guillaume Drion, and Eve Marder (Apr. 2016). “Computational implications of biophysical  
2524 diversity and multiple timescales in neurons and synapses for circuit performance”. In: *Current Opinion in*  
2525 *Neurobiology* 37, pp. 44–52.
- 2526 Glanzman, David L (2010). “Common mechanisms of synaptic plasticity in vertebrates and invertebrates”. In:  
2527 *Current Biology* 20.1, R31–R36.
- 2528 Gold, Joshua I. and Michael N. Shadlen (2002). “Banburismus and the Brain: Decoding the Relationship between  
2529 Sensory Stimuli, Decisions, and Reward”. In: *Neuron* 36.2, pp. 299–308. ISSN: 0896-6273. DOI: [https://doi.org/10.1016/S0896-6273\(02\)00971-6](https://doi.org/10.1016/S0896-6273(02)00971-6). URL: <https://www.sciencedirect.com/science/article/pii/S0896627302009716>.
- 2530  
2531
- 2532 Gold, Joshua I and Michael N Shadlen (2007a). “The neural basis of decision making”. In: *Annu. Rev. Neurosci.*  
2533 30, pp. 535–574.
- 2534 Gold, Joshua I. and Michael N. Shadlen (July 2007b). “the neural basis of decision making”. In: *Annual Review of*  
2535 *Neuroscience* 30. DOI: <https://doi.org/10.1146/annurev.neuro.29.051605.113038>.
- 2536 Golowasch, Jorge et al. (2002). “Failure of Averaging in the Construction of a Conductance-Based Neuron Model”.  
2537 In: *Journal of Neurophysiology* 87.2, pp. 1129–1131. DOI: <https://doi.org/10.1152/jn.00412.2001>.
- 2538 Grabe, Veit et al. (2016). “Elucidating the Neuronal Architecture of Olfactory Glomeruli in the *Drosophila* Antennal  
2539 Lobe”. In: *Cell Reports* 16.12, pp. 3401–3413. ISSN: 2211-1247. DOI: <https://doi.org/10.1016/j.celrep.2016.08.063>. URL: <https://www.sciencedirect.com/science/article/pii/S2211124716311445>.
- 2540
- 2541 Granit, Ragnar and CG Phillips (1956). “Excitatory and inhibitory processes acting upon individual Purkinje cells  
2542 of the cerebellum in cats”. In: *The Journal of Physiology* 133.3, p. 520.
- 2543 Grashow, Rachel, Ted Brookings, and Eve Marder (July 2010). “Compensation for variable intrinsic neuronal  
2544 excitability by circuit-synaptic interactions.” In: *Journal of Neuroscience* 30.27, pp. 9145–9156.

- 2545 Groschner, Lukas N. et al. (2018). “Dendritic Integration of Sensory Evidence in Perceptual Decision-Making”.  
2546 In: *Cell* 173.4, 894–905.e13. ISSN: 0092-8674. DOI: <https://doi.org/10.1016/j.cell.2018.03.075>. URL:  
2547 <https://www.sciencedirect.com/science/article/pii/S0092867418304471>.
- 2548 Grubb, Matthew S and Juan Burrone (June 2010). “Activity-dependent relocation of the axon initial segment  
2549 fine-tunes neuronal excitability.” In: *Nature* 465.7301, pp. 1070–1074.
- 2550 Gruntman, Eyal and Glenn C. Turner (2013). “Integration of the olfactory code across dendritic claws of single  
2551 mushroom body neurons”. In: *Nature Neuroscience* 16.12, pp. 1821–1829. DOI: 10.1038/nn.3547.
- 2552 Hallem, Elissa A. and John R. Carlson (2006). “Coding of Odors by a Receptor Repertoire”. In: *Cell* 125.1, pp. 143–  
2553 160. DOI: <https://doi.org/10.1016/j.cell.2006.01.050>.
- 2554 Han, Xiaoying and Peter E. Kloeden (2020). “Sigmoidal approximations of Heaviside functions in neural lattice  
2555 models”. In: *Journal of Differential Equations* 268.9, pp. 5283–5300. DOI: <https://doi.org/10.1016/j.jde.2019.11.010>.
- 2557 Handler, Annie et al. (2019). “Distinct Dopamine Receptor Pathways Underlie the Temporal Sensitivity of Asso-  
2558 ciative Learning”. In: *Cell* 178.1, pp. 60–75. DOI: <https://doi.org/10.1016/j.cell.2019.05.040>.
- 2559 Hanes, Doug P and Jeffrey D Schall (1996). “Neural control of voluntary movement initiation”. In: *Science* 274.5286,  
2560 pp. 427–430.
- 2561 Hebb, Donald Olding (2005). *The organization of behavior: A neuropsychological theory*. Psychology Press.
- 2562 Heeger, David J (1992). “Normalization of cell responses in cat striate cortex”. In: *Visual neuroscience* 9.2, pp. 181–  
2563 197.
- 2564 Heeger, DJ, M Landy, and JA Movshon (1991). *Computational models of visual processing*.
- 2565 Hige, Toshihide et al. (2015). “Heterosynaptic Plasticity Underlies Aversive Olfactory Learning in *Drosophila*”. In:  
2566 *Neuron* 88.5, pp. 985–998. DOI: 10.1016/j.neuron.2015.11.003.
- 2567 Hiratani, Naoki and Peter E. Latham (2022). “Developmental and evolutionary constraints on olfactory circuit  
2568 selection”. In: *Proceedings of the National Academy of Sciences* 119.11, e2100600119. DOI: 10.1073/pnas.  
2569 2100600119. eprint: <https://www.pnas.org/doi/pdf/10.1073/pnas.2100600119>. URL: <https://www.pnas.org/doi/abs/10.1073/pnas.2100600119>.
- 2571 Hofman, Michel (2014). “Evolution of the human brain: when bigger is better”. In: *Frontiers in Neuroanatomy* 8.  
2572 ISSN: 1662-5129. DOI: 10.3389/fnana.2014.00015. URL: <https://www.frontiersin.org/article/10.3389/fnana.2014.00015>.
- 2574 Holler, Simone et al. (Jan. 2021). “Structure and function of a neocortical synapse”. In: *Nature* 314, pp. 1–6.
- 2575 Holt, Gary R and Christof Koch (1997). “Shunting inhibition does not have a divisive effect on firing rates”. In:  
2576 *Neural computation* 9.5, pp. 1001–1013.
- 2577 Honegger, Kyle S., Robert A. A. Campbell, and Glenn C. Turner (2011). “Cellular-Resolution Population Imag-  
2578 ing Reveals Robust Sparse Coding in the *Drosophila* Mushroom Body”. In: *Journal of Neuroscience* 31.33,  
2579 pp. 11772–11785. DOI: <https://doi.org/10.1523/JNEUROSCI.1099-11.2011>.

- 2580 Houston, Catriona M et al. (Apr. 2017). “Exploring the significance of morphological diversity for cerebellar granule  
2581 cell excitability”. In: *Scientific Reports* 7.1, pp. 1–16.
- 2582 Huang, Cheng-Chiu et al. (2013). “Convergence of pontine and proprioceptive streams onto multimodal cerebellar  
2583 granule cells”. In: *Elife* 2, e00400.
- 2584 Ichinose, Toshiharu et al. (2015). “Reward signal in a recurrent circuit drives appetitive long-term memory for-  
2585 mation”. In: *eLife* 4. Ed. by Leslie C Griffith, e10719. ISSN: 2050-084X. DOI: 10.7554/eLife.10719. URL:  
2586 <https://doi.org/10.7554/eLife.10719>.
- 2587 Inada, Kengo, Yoshiko Tsuchimoto, and Hokto Kazama (July 2017). “Origins of Cell-Type-Specific Olfactory Pro-  
2588 cessing in the Drosophila Mushroom Body Circuit.” In: *Neuron* 95.2, 357–367.e4.
- 2589 Ito, Kei et al. (1997). “The Drosophila mushroom body is a quadruple structure of clonal units each of which  
2590 contains a virtually identical set of neurones and glial cells”. In: *Development* 124.4, pp. 761–771.
- 2591 Ito, M (1989). “Long-Term Depression”. In: *Annual Review of Neuroscience* 12.1. PMID: 2648961, pp. 85–102. DOI:  
2592 10.1146/annurev.ne.12.030189.000505. eprint: <https://doi.org/10.1146/annurev.ne.12.030189.000505>. URL: <https://doi.org/10.1146/annurev.ne.12.030189.000505>.
- 2593
- 2594 Ito, Masao (1972). “Neural design of the cerebellar motor control system”. In: *Brain Research* 40.1, pp. 81–84. ISSN:  
2595 0006-8993. DOI: [https://doi.org/10.1016/0006-8993\(72\)90110-2](https://doi.org/10.1016/0006-8993(72)90110-2). URL: <https://www.sciencedirect.com/science/article/pii/0006899372901102>.
- 2596
- 2597 — (2006). “Cerebellar circuitry as a neuronal machine”. In: *Progress in Neurobiology* 78.3. The Contributions of  
2598 John Carew Eccles to Contemporary Neuroscience, pp. 272–303. ISSN: 0301-0082. DOI: <https://doi.org/10.1016/j.pneurobio.2006.02.006>. URL: <https://www.sciencedirect.com/science/article/pii/S0301008206000232>.
- 2599
- 2600
- 2601 Jacob, Pedro F and Scott Waddell (2020). “Spaced training forms complementary long-term memories of opposite  
2602 valence in Drosophila”. In: *Neuron* 106.6, pp. 977–991.
- 2603 Johansen, Joshua P. et al. (2014). “Hebbian and neuromodulatory mechanisms interact to trigger associative memory  
2604 formation”. In: *Proceedings of the National Academy of Sciences* 111.51, E5584–E5592. DOI: 10.1073/pnas.  
2605 1421304111. eprint: <https://www.pnas.org/doi/pdf/10.1073/pnas.1421304111>. URL: <https://www.pnas.org/doi/abs/10.1073/pnas.1421304111>.
- 2606
- 2607 Jörntell, Henrik and Carl-Fredrik Ekerot (2006). “Properties of somatosensory synaptic integration in cerebellar  
2608 granule cells in vivo”. In: *Journal of Neuroscience* 26.45, pp. 11786–11797.
- 2609 Kawato, Mitsuo et al. (2021). “50 Years Since the Marr, Ito, and Albus Models of the Cerebellum”. In: *Neuroscience*  
2610 462. In Memoriam: Masao Ito—A Visionary Neuroscientist with a Passion for the Cerebellum, pp. 151–174.  
2611 ISSN: 0306-4522. DOI: <https://doi.org/10.1016/j.neuroscience.2020.06.019>. URL: <https://www.sciencedirect.com/science/article/pii/S0306452220303961>.
- 2612
- 2613 Kazama, Hokto and Rachel I. Wilson (May 2008). “Homeostatic matching and nonlinear amplification at identified  
2614 central synapses.” In: *Neuron* 58.3, pp. 401–413.

- 2615 Kennedy, Ann (2019). “Learning with naturalistic odor representations in a dynamic model of the *Drosophila*  
2616 olfactory system”. In: *bioRxiv*. DOI: <https://doi.org/10.1101/783191>.
- 2617 Keung, Waitsang, Todd A Hagen, and Robert C Wilson (2020). “A divisive model of evidence accumulation explains  
2618 uneven weighting of evidence over time”. In: *Nature communications* 11.1, pp. 1–9.
- 2619 KF, Wong et al. (2007). “Neural circuit dynamics underlying accumulation of time-varying evidence during percep-  
2620 tual decision making”. In: *Frontiers in computational neuroscience* 1. DOI: [doi:10.3389/neuro.10.006.2007](https://doi.org/10.3389/neuro.10.006.2007).
- 2621 Knogler, Laura D et al. (2017). “Sensorimotor representations in cerebellar granule cells in larval zebrafish are  
2622 dense, spatially organized, and non-temporally patterned”. In: *Current Biology* 27.9, pp. 1288–1302.
- 2623 Kriener, Birgit, Rishidev Chaudhuri, and Ila R. Fiete (2020a). “Robust parallel decision-making in neural circuits  
2624 with nonlinear inhibition”. In: *Proceedings of the National Academy of Sciences* 117.41, pp. 25505–25516. DOI:  
2625 [10.1073/pnas.1917551117](https://doi.org/10.1073/pnas.1917551117). eprint: <https://www.pnas.org/doi/pdf/10.1073/pnas.1917551117>. URL:  
2626 <https://www.pnas.org/doi/abs/10.1073/pnas.1917551117>.
- 2627 — (2020b). “Robust parallel decision-making in neural circuits with nonlinear inhibition”. In: *Proceedings of the*  
2628 *National Academy of Sciences* 117.41, pp. 25505–25516. DOI: [10.1073/pnas.1917551117](https://doi.org/10.1073/pnas.1917551117). eprint: <https://www.pnas.org/doi/pdf/10.1073/pnas.1917551117>. URL: <https://www.pnas.org/doi/abs/10.1073/pnas.1917551117>.
- 2631 Krishnamurthy, K, A M Hermundstad, and T Mora (2017). “Disorder and the neural representation of complex  
2632 odors: smelling in the real world”. In: *arXiv.org*. scholar: C257FC5C-A5A2-4C43-8284-E9BA402F1DAA.
- 2633 Latimer, Kenneth W. et al. (2015). “Single-trial spike trains in parietal cortex reveal discrete steps during decision-  
2634 making”. In: *Science* 349.6244, pp. 184–187. DOI: [10.1126/science.aaa4056](https://doi.org/10.1126/science.aaa4056). eprint: <https://www.science.org/doi/pdf/10.1126/science.aaa4056>. URL: <https://www.science.org/doi/abs/10.1126/science.aaa4056>.
- 2637 Lee, Tzumin, Arthur Lee, and Liqun Luo (1999). “Development of the *Drosophila* mushroom bodies: sequential  
2638 generation of three distinct types of neurons from a neuroblast”. In: *Development* 126.18, pp. 4065–4076.
- 2639 Leiss, Florian et al. (2009). “Synaptic organization in the adult *Drosophila* mushroom body calyx”. In: *Journal of*  
2640 *Comparative Neurology* 517.6, pp. 808–824.
- 2641 Lewis, Laurence PC et al. (2015). “A higher brain circuit for immediate integration of conflicting sensory information  
2642 in *Drosophila*”. In: *Current biology* 25.17, pp. 2203–2214.
- 2643 Li, Feng et al. (Dec. 2020). “The connectome of the adult *Drosophila* mushroom body provides insights into func-  
2644 tion”. In: *eLife* 9.
- 2645 Li, Hao et al. (July 2013). “Transformation of odor selectivity from projection neurons to single mushroom body  
2646 neurons mapped with dual-color calcium imaging.” In: *Proceedings of the National Academy of Sciences of the*  
2647 *United States of America*.



- 2648 Li, Quan, Adam C Roberts, and David L Glanzman (2005). “Synaptic facilitation and behavioral dishabituation  
2649 in *Aplysia*: dependence on release of Ca<sup>2+</sup> from postsynaptic intracellular stores, postsynaptic exocytosis, and  
2650 modulation of postsynaptic AMPA receptor efficacy”. In: *Journal of Neuroscience* 25.23, pp. 5623–5637.
- 2651 Lin, Andrew C. et al. (2014). “Sparse, Decorrelated Odor Coding in the Mushroom Body Enhances Learned Odor  
2652 Discrimination”. In: *Nature Neuroscience* 17.4, pp. 559–568. DOI: 10.1038/nn.3660.
- 2653 Lin, Xiang Y and David L Glanzman (1994). “Long-term potentiation of *Aplysia* sensorimotor synapses in cell  
2654 culture: regulation by postsynaptic voltage”. In: *Proceedings of the Royal Society of London. Series B: Biological  
2655 Sciences* 255.1343, pp. 113–118.
- 2656 Lisman, John and Sridhar Raghavachari (2006). “A unified model of the presynaptic and postsynaptic changes  
2657 during LTP at CA1 synapses”. In: *Science’s STKE* 2006.356, re11–re11.
- 2658 Litwin-Kumar, Ashok et al. (2017). “Optimal Degrees of Synaptic Connectivity”. In: *Neuron* 93, pp. 1153–1164.  
2659 DOI: <https://doi.org/10.1016/j.neuron.2017.01.030>.
- 2660 Liu, Chang et al. (2012). “A subset of dopamine neurons signals reward for odour memory in *Drosophila*”. In:  
2661 *Nature* 488.7412, pp. 512–516.
- 2662 Louie, Kenway, Lauren E Grattan, and Paul W Glimcher (2011). “Reward value-based gain control: divisive nor-  
2663 malization in parietal cortex”. In: *Journal of Neuroscience* 31.29, pp. 10627–10639.
- 2664 Luo, Sean X., Richard Axel, and L. F. Abbott (2010). “Generating sparse and selective third-order responses in the  
2665 olfactory system of the fly”. In: *PNAS* 107.23, pp. 10713–10718. DOI: 10.1073/pnas.1005635107.
- 2666 Lyu, Siwei and Eero Simoncelli (Mar. 2009). “Nonlinear Extraction of Independent Components of Natural Images  
2667 Using Radial Gaussianization”. In: *Neural computation* 21, pp. 1485–519. DOI: 10.1162/neco.2009.04-08-773.
- 2668 MacLean, Jason N et al. (Jan. 2003). “Activity-Independent Homeostasis in Rhythmically Active Neurons”. In:  
2669 *Neuron* 37.1, pp. 109–120.
- 2670 MacLean, Jason N et al. (Nov. 2005). “Activity-independent coregulation of IA and Ih in rhythmically active  
2671 neurons.” In: *Journal of Neurophysiology* 94.5, pp. 3601–3617.
- 2672 Malenka, Robert C. (1991). “Postsynaptic factors control the duration of synaptic enhancement in area CA1 of  
2673 the hippocampus”. In: *Neuron* 6.1, pp. 53–60. ISSN: 0896-6273. DOI: [https://doi.org/10.1016/0896-](https://doi.org/10.1016/0896-6273(91)90121-F)  
2674 [6273\(91\)90121-F](https://doi.org/10.1016/0896-6273(91)90121-F). URL: <https://www.sciencedirect.com/science/article/pii/089662739190121F>.
- 2675 Manneschi, Luca, Andrew C. Lin, and Eleni Vasilaki (in press). *SpaRCe: Improved Learning of Reservoir Computing  
2676 Systems through Sparse Representations*.
- 2677 Manneschi, Luca et al. (2021). *Exploiting Multiple Timescales in Hierarchical Echo State Networks*. arXiv: 2101.  
2678 04223 [cs.LG].
- 2679 Mao, Zhengmei and Ronald L Davis (2009). “Eight different types of dopaminergic neurons innervate the *Drosophila*  
2680 mushroom body neuropil: anatomical and physiological heterogeneity”. In: *Frontiers in neural circuits* 3, p. 5.
- 2681 Marder, Eve and Jean-Marc Goaillard (2006). “Variability, compensation and homeostasis in neuron and network  
2682 function”. In: *Nature reviews Neuroscience* 7.7, pp. 563–74. DOI: 10.1038/nrn1949.

- 2683 Markram, Henry et al. (1997). “Regulation of Synaptic Efficacy by Coincidence of Postsynaptic APs and EPSPs”.  
2684 In: *Science* 275.5297, pp. 213–215. DOI: 10.1126/science.275.5297.213. eprint: <https://www.science.org/doi/pdf/10.1126/science.275.5297.213>. URL: <https://www.science.org/doi/abs/10.1126/science.275.5297.213>.  
2685  
2686
- 2687 Marr, D (June 1969). “A theory of cerebellar cortex.” In: *Journal of physiology* 202.2, pp. 437–470.
- 2688 Marsat, Gary and Leonard Maler (Nov. 2010). “Neural Heterogeneity and Efficient Population Codes for Commu-  
2689 nication Signals”. In: *Journal of Neurophysiology* 104.5, pp. 2543–2555.
- 2690 Masse, Nicolas Y, Glenn C Turner, and Gregory SXE Jefferis (2009). “Olfactory information processing in *Drosophila*”.  
2691 In: *Current Biology* 19.16, R700–R713.
- 2692 Memmesheimer, Raoul-Martin et al. (2014). “Learning Precisely Timed Spikes”. In: *Neuron* 82.4, pp. 925–938. ISSN:  
2693 0896-6273. DOI: <https://doi.org/10.1016/j.neuron.2014.03.026>. URL: <https://www.sciencedirect.com/science/article/pii/S0896627314002566>.  
2694
- 2695 Menzel, Randolf and Gisela Manz (2005). “Neural plasticity of mushroom body-extrinsic neurons in the honeybee  
2696 brain”. In: *Journal of Experimental Biology* 208.22, pp. 4317–4332.
- 2697 Milosavljevic, Milica et al. (2010). “The drift diffusion model can account for value-based choice response times  
2698 under high and low time pressure”. In: *Judgment and Decision making* 5.6, pp. 437–449.
- 2699 Mink, Jonathan W (1996). “The basal ganglia: focused selection and inhibition of competing motor programs”. In:  
2700 *Progress in neurobiology* 50.4, pp. 381–425.
- 2701 Mock, William B. T. (2011). “Pareto Optimality”. In: *Encyclopedia of Global Justice*. Ed. by Deen K. Chatterjee.  
2702 Dordrecht: Springer Netherlands, pp. 808–809. ISBN: 978-1-4020-9160-5. DOI: 10.1007/978-1-4020-9160-  
2703 5\_341. URL: [https://doi.org/10.1007/978-1-4020-9160-5\\_341](https://doi.org/10.1007/978-1-4020-9160-5_341).
- 2704 Modi, Mehrab N., Yichun Shuai, and Glenn C. Turner (July 2020). “The *Drosophila* Mushroom Body: From  
2705 Architecture to Algorithm in a Learning Circuit.” In: *Annual Review of Neuroscience* 43.1, pp. 465–484.
- 2706 Namiki, Shigehiro et al. (2018). “The functional organization of descending sensory-motor pathways in *Drosophila*”.  
2707 In: *Elife* 7, e34272.
- 2708 Normann, R A and I Perlman (1979). “The effects of background illumination on the photoresponses of red and green  
2709 cones.” In: *The Journal of Physiology* 286.1, pp. 491–507. DOI: <https://doi.org/10.1113/jphysiol.1979.sp012633>.  
2710 eprint: <https://physoc.onlinelibrary.wiley.com/doi/pdf/10.1113/jphysiol.1979.sp012633>.  
2711 URL: <https://physoc.onlinelibrary.wiley.com/doi/abs/10.1113/jphysiol.1979.sp012633>.
- 2712 O’Leary, Timothy and Eve Marder (Nov. 2016). “Temperature-Robust Neural Function from Activity-Dependent  
2713 Ion Channel Regulation.” In: 26.21, pp. 2935–2941.
- 2714 Olsen, Shawn R., Vikas Bhandawat, and Rachel I. Wilson (2010). “Divisive normalization in olfactory population  
2715 codes”. In: *Neuron* 66.2, pp. 287–299. DOI: <https://doi.org/10.1016/j.neuron.2010.04.009>.
- 2716 Olsen, Shawn R and Rachel I Wilson (2008). “Lateral presynaptic inhibition mediates gain control in an olfactory  
2717 circuit”. In: *Nature* 452.7190, pp. 956–960.

- 2718 Oswald, David et al. (Apr. 2015). “Activity of defined mushroom body output neurons underlies learned olfactory  
2719 behavior in *Drosophila*.” In: *Neuron* 86.2, pp. 417–427.
- 2720 Padmanabhan, Krishnan and Nathaniel N Urban (Aug. 2010). “Intrinsic biophysical diversity decorrelates neuronal  
2721 firing while increasing information content”. In: *Nature Neuroscience* 13.10, pp. 1276–1282.
- 2722 — (Sept. 2014). “Disrupting information coding via block of 4-AP-sensitive potassium channels”. In: *Journal of*  
2723 *Neurophysiology* 112.5, pp. 1054–1066.
- 2724 Parnas, Moshe et al. (2013). “Odor Discrimination in *Drosophila*: From Neural Population Codes to Behavior”. In:  
2725 *Neuron* 79.5, pp. 932–44. DOI: 10.1016/j.neuron.2013.08.006..
- 2726 Parrish, Jay Z et al. (May 2014). “Krüppel Mediates the Selective Rebalancing of Ion Channel Expression”. In:  
2727 *Neuron* 82.3, pp. 537–544.
- 2728 Peng, Fei and Lars Chittka (2017). “A Simple Computational Model of the Bee Mushroom Body Can Explain  
2729 Seemingly Complex Forms of Olfactory Learning and Memory”. In: *Current Biology* 27.2, pp. 224–230. ISSN:  
2730 0960-9822. DOI: <https://doi.org/10.1016/j.cub.2016.10.054>. URL: <https://www.sciencedirect.com/science/article/pii/S096098221631288X>.
- 2732 Perez-Nieves, Nicolas et al. (Dec. 2020). “Neural heterogeneity promotes robust learning”. In: *bioRxiv* 15.4, p. 2020.12.18.423468.
- 2733 Platt, Michael L and Paul W Glimcher (1999). “Neural correlates of decision variables in parietal cortex”. In: *Nature*  
2734 400.6741, pp. 233–238.
- 2735 Powell, Kate et al. (June 2015). “Synaptic representation of locomotion in single cerebellar granule cells”. In: *eLife*  
2736 4, p. 977.
- 2737 Rakic, Pasko (2009). “Evolution of the neocortex: a perspective from developmental biology”. In: *Nature Reviews*  
2738 *Neuroscience* 10.10, pp. 724–735.
- 2739 Rapp, Hannes and Martin Paul Nawrot (2020). “A spiking neural program for sensorimotor control during foraging  
2740 in flying insects”. In: *Proceedings of the National Academy of Sciences* 117.45, pp. 28412–28421. DOI: 10.1073/  
2741 pnas.2009821117. eprint: <https://www.pnas.org/doi/pdf/10.1073/pnas.2009821117>. URL: <https://www.pnas.org/doi/abs/10.1073/pnas.2009821117>.
- 2743 Ratcliff, Roger (1978). “A theory of memory retrieval.” In: *Psychological review* 85.2, p. 59.
- 2744 Ravenscroft, Thomas A et al. (Sept. 2020). “*Drosophila* voltage-gated sodium channels are only expressed in active  
2745 neurons and are localized to distal axonal initial segment-like domains”. In: *Journal of Neuroscience*, JN–RM–  
2746 0142–20.
- 2747 Raymond, Jennifer L and Javier F Medina (2018). “Computational principles of supervised learning in the cerebel-  
2748 lum”. In: *Annual review of neuroscience* 41, pp. 233–253.
- 2749 Redgrave, Peter, Tony J Prescott, and Kevin Gurney (1999). “The basal ganglia: a vertebrate solution to the  
2750 selection problem?” In: *Neuroscience* 89.4, pp. 1009–1023.

- 2751 Roitman, Jamie D. and Michael N. Shadlen (2002a). “Response of neurons in the lateral intraparietal area during  
2752 a combined visual discrimination reaction time task”. In: *journal of neuroscience*. DOI: <https://doi.org/10.1523/JNEUROSCI.22-21-09475.2002>.
- 2753 — (2002b). “Response of Neurons in the Lateral Intraparietal Area during a Combined Visual Discrimination  
2754 Reaction Time Task”. In: *Journal of Neuroscience* 22.21, pp. 9475–9489. ISSN: 0270-6474. DOI: 10.1523/  
2755 JNEUROSCI.22-21-09475.2002. eprint: <https://www.jneurosci.org/content/22/21/9475.full.pdf>. URL:  
2756 <https://www.jneurosci.org/content/22/21/9475>.
- 2757 Rosenblatt, Frank (1958). “The perceptron: a probabilistic model for information storage and organization in the  
2758 brain.” In: *Psychological review* 65.6, p. 386.
- 2759 Scheffer, Louis K et al. (Apr. 2020). “A Connectome and Analysis of the Adult *Drosophila* Central Brain”. In:  
2760 *bioRxiv* 12.4, p. 2020.04.07.030213.
- 2761 Schulz, David J., Jean-Marc Goillard, and Eve Marder (2006). “Variable channel expression in identified single  
2762 and electrically coupled neurons in different animals”. In: *Nature Neuroscience* 9, 356–362. DOI: <https://doi.org/10.1038/nn1639>.
- 2763 — (2007). “Quantitative expression profiling of identified neurons reveals cell-specific constraints on highly variable  
2764 levels of gene expression”. In: *PNAS* 32 (104), pp. 13187–91. DOI: 10.1073/pnas.0705827104.
- 2765 Schwaerzel, Martin et al. (2003). “Dopamine and octopamine differentiate between aversive and appetitive olfactory  
2766 memories in *Drosophila*”. In: *Journal of Neuroscience* 23.33, pp. 10495–10502.
- 2767 Schwartz, Odelia and Eero Simoncelli (Sept. 2001a). “Natural signal statistics and sensory gain control”. In: *Nature*  
2768 *neuroscience* 4, pp. 819–25. DOI: 10.1038/90526.
- 2769 Schwartz, Odelia and Eero P Simoncelli (2001b). “Natural signal statistics and sensory gain control”. In: *Nature*  
2770 *neuroscience* 4.8, pp. 819–825.
- 2771 Schwartz, S. and Y.S. Yeh (1982). “The Distribution Function and Moments Of Power Sums With Log-normal  
2772 Components”. In: *Bell Syst. Tech. J.* 61, pp. 1441–1462.
- 2773 Schweighofer, N., K. Doya, and F. Lay (2001). “Unsupervised learning of granule cell sparse codes enhances cerebellar  
2774 adaptive control”. In: *Neuroscience* 103.1, pp. 35–50. ISSN: 0306-4522. DOI: [https://doi.org/10.1016/S0306-4522\(00\)00548-0](https://doi.org/10.1016/S0306-4522(00)00548-0). URL: <https://www.sciencedirect.com/science/article/pii/S0306452200005480>.
- 2775 Séjourné, Julien et al. (2011). “Mushroom body efferent neurons responsible for aversive olfactory memory retrieval  
2776 in *Drosophila*”. In: *Nature neuroscience* 14.7, pp. 903–910.
- 2777 Shadlen, Michael N. and William T. Newsome (2001). “Neural Basis of a Perceptual Decision in the Parietal  
2778 Cortex (Area LIP) of the Rhesus Monkey”. In: *Journal of Neurophysiology* 86.4. PMID: 11600651, pp. 1916–  
2779 1936. DOI: 10.1152/jn.2001.86.4.1916. eprint: <https://doi.org/10.1152/jn.2001.86.4.1916>. URL:  
2780 <https://doi.org/10.1152/jn.2001.86.4.1916>.

- 2784 Shadlen, Michael N. and Roozbeh Kiani (2013). “Decision Making as a Window on Cognition”. In: *Neuron* 80.3,  
2785 pp. 791–806. ISSN: 0896-6273. DOI: <https://doi.org/10.1016/j.neuron.2013.10.047>. URL: <https://www.sciencedirect.com/science/article/pii/S0896627313009999>.  
2786
- 2787 Shafir, Sharoni, Tom A Waite, and Brian H Smith (2002). “Context-dependent violations of rational choice in  
2788 honeybees (*Apis mellifera*) and gray jays (*Perisoreus canadensis*)”. In: *Behavioral Ecology and Sociobiology* 51.2,  
2789 pp. 180–187.
- 2790 Shen, Yang, Sanjoy Dasgupta, and Saket Navlakha (2020). “Habituation as a neural algorithm for online odor  
2791 discrimination”. In: *Proceedings of the National Academy of Sciences* 117.22, pp. 12402–12410.
- 2792 — (2021). “Algorithmic insights on continual learning from fruit flies”. In: *arXiv:2107.07617*. DOI: <https://doi.org/10.48550/arXiv.2107.07617>.  
2793
- 2794 Shen, Yang, Julia Wang, and Saket Navlakha (Nov. 2021). “A Correspondence Between Normalization Strategies  
2795 in Artificial and Biological Neural Networks”. In: *Neural Computation* 33.12, pp. 3179–3203. ISSN: 0899-7667.  
2796 DOI: 10.1162/neco\_a\_01439. eprint: [https://direct.mit.edu/neco/article-pdf/33/12/3179/1971993/  
2797 neco\\_a\\_01439.pdf](https://direct.mit.edu/neco/article-pdf/33/12/3179/1971993/neco_a_01439.pdf). URL: [https://doi.org/10.1162/neco\\_a\\_01439](https://doi.org/10.1162/neco_a_01439).
- 2798 Song, Sen et al. (Mar. 2005). “Highly nonrandom features of synaptic connectivity in local cortical circuits.” In:  
2799 *PLoS biology* 3.3, e68.
- 2800 Sreenivasan, Sameet and Ila Fiete (2011). “Grid cells generate an analog error-correcting code for singularly precise  
2801 neural computation”. In: *Nature neuroscience* 14.10, pp. 1330–1337.
- 2802 Stevens, Charles F (2015). “What the fly’s nose tells the fly’s brain”. In: *Proceedings of the National Academy of  
2803 Sciences* 112.30, pp. 9460–9465.
- 2804 Sultan, Fahad (2001). “Distribution of mossy fibre rosettes in the cerebellum of cat and mice: evidence for a  
2805 parasagittal organization at the single fibre level”. In: *European Journal of Neuroscience* 13.11, pp. 2123–2130.
- 2806 Sutton, Richard S and Andrew G Barto (2018). *Reinforcement learning: An introduction*. MIT press.
- 2807 Tabone, Christopher J and J Steven de Belle (2011). “Second-order conditioning in *Drosophila*”. In: *Learning &  
2808 Memory* 18.4, pp. 250–253.
- 2809 Tanaka, Nobuaki K, Hiromu Tanimoto, and Kei Ito (2008). “Neuronal assemblies of the *Drosophila* mushroom  
2810 body”. In: *Journal of Comparative Neurology* 508.5, pp. 711–755.
- 2811 Taylor, Adam L., Jean-Marc Goaillard, and Eve Marder (2009). “How Multiple Conductances Determine Electro-  
2812 physiological Properties in a Multicompartment Model”. In: *Journal of Neuroscience* 29.17, pp. 5573–5586. DOI:  
2813 10.1523/JNEUROSCI.4438-08.2009.
- 2814 Tobin, Anne-Elise and Ronald L. Calabrese (2006). “Endogenous and half-center bursting in morphologically-  
2815 inspired models of leech heart interneurons”. In: *Journal of Neurophysiology* 96.4. DOI: 10.1152/jn.00025.2006.
- 2816 Tobin, William F., Rachel I. Wilson, and Wei-Chung Allen Lee (May 2017). “Wiring variations that enable and  
2817 constrain neural computation in a sensory microcircuit.” In: *eLife* 6, e24838.

- 2818 Tolhurst, David J, J Anthony Movshon, and Andrew F Dean (1983). “The statistical reliability of signals in single  
2819 neurons in cat and monkey visual cortex”. In: *Vision research* 23.8, pp. 775–785.
- 2820 Tripathy, Shreejoy J et al. (May 2013). “Intermediate intrinsic diversity enhances neural population coding”. In:  
2821 *PNAS* 110.20, pp. 8248–8253.
- 2822 Trunova, Svetlana, Brian Baek, and Edward Giniger (July 2011). “Cdk5 regulates the size of an axon initial segment-  
2823 like compartment in mushroom body neurons of the Drosophila central brain.” In: *Journal of Neuroscience* 31.29,  
2824 pp. 10451–10462.
- 2825 Tsai, Kuo-Ting et al. (2018). “Circuit variability interacts with excitatory-inhibitory diversity of interneurons to  
2826 regulate network encoding capacity”. In: *scientific reports* 8. DOI: [https://doi.org/10.1038/s41598-018-](https://doi.org/10.1038/s41598-018-26286-8)  
2827 [26286-8](https://doi.org/10.1038/s41598-018-26286-8).
- 2828 Turner, Glenn C., Maxim Bazhenov, and Gilles Laurent (2008). “Olfactory Representations by Drosophila Mush-  
2829 room Body Neurons”. In: *Journal of Neurophysiology* 99.2, pp. 734–746. DOI: [https://doi.org/10.1152/jn.](https://doi.org/10.1152/jn.01283.2007)  
2830 [01283.2007](https://doi.org/10.1152/jn.01283.2007).
- 2831 Tversky, Amos and Itamar Simonson (1993). “Context-Dependent Preferences”. In: *Management Science* 39.10,  
2832 pp. 1179–1189. ISSN: 00251909, 15265501. URL: <http://www.jstor.org/stable/2632953> (visited on 04/15/2022).
- 2833 Usher, M and Jay Mcclelland (July 2001). “On the time course of perceptual choice: The leaky competing accu-  
2834 mulator model”. English. In: *Psychological Review* 108, pp. 550–592. ISSN: 0033-295X. DOI: [10.1037/0033-](https://doi.org/10.1037/0033-295X.108.3.550)  
2835 [295X.108.3.550](https://doi.org/10.1037/0033-295X.108.3.550).
- 2836 Vrontou, Eleftheria et al. (2021). “Response competition between neurons and antineurons in the mushroom body”.  
2837 In: *Current Biology* 31.22, 4911–4922.e4. ISSN: 0960-9822. DOI: [https://doi.org/10.1016/j.cub.2021.09.](https://doi.org/10.1016/j.cub.2021.09.008)  
2838 [008](https://doi.org/10.1016/j.cub.2021.09.008). URL: <https://www.sciencedirect.com/science/article/pii/S0960982221012483>.
- 2839 Wald, Abraham (2004). *Sequential analysis*. Courier Corporation.
- 2840 Wang, Xiao-Jing (2002). “Probabilistic Decision Making by Slow Reverberation in Cortical Circuits”. In: *Neuron*  
2841 36.5, pp. 955–968. ISSN: 0896-6273. DOI: [https://doi.org/10.1016/S0896-6273\(02\)01092-9](https://doi.org/10.1016/S0896-6273(02)01092-9). URL: [https:](https://www.sciencedirect.com/science/article/pii/S0896627302010929)  
2842 [//www.sciencedirect.com/science/article/pii/S0896627302010929](https://www.sciencedirect.com/science/article/pii/S0896627302010929).
- 2843 Wessnitzer, Jan et al. (2012). “A model of non-elemental olfactory learning in Drosophila”. In: *Journal of compu-*  
2844 *tational neuroscience* 32.2, pp. 197–212.
- 2845 Williams, Stephen R and Greg J Stuart (Mar. 2003). “Role of dendritic synapse location in the control of action  
2846 potential output”. In: *Trends in Neurosciences* 26.3, pp. 147–154.
- 2847 Willmore, B and D J Tolhurst (Aug. 2001). “Characterizing the sparseness of neural codes.” In: *Network: Compu-*  
2848 *tation in Neural Systems* 12.3, pp. 255–270.
- 2849 Xia, Shouzhen et al. (2005). “NMDA Receptors Mediate Olfactory Learning and Memory in Drosophila”. In:  
2850 *Current Biology* 15.7, pp. 603–615. ISSN: 0960-9822. DOI: <https://doi.org/10.1016/j.cub.2005.02.059>.  
2851 URL: <https://www.sciencedirect.com/science/article/pii/S0960982205002320>.

- 2852 Xie, Xiaohui, Richard HR Hahnloser, and H Sebastian Seung (2002). “Selectively grouping neurons in recurrent  
2853 networks of lateral inhibition”. In: *Neural computation* 14.11, pp. 2627–2646.
- 2854 Yasuyama, Kouji, Ian A Meinertzhagen, and Friedrich-Wilhelm Schürmann (2002). “Synaptic organization of the  
2855 mushroom body calyx in *Drosophila melanogaster*”. In: *Journal of Comparative Neurology* 445.3, pp. 211–226.
- 2856 Zakharenko, Zablow, and Siegelbaum (2001). “Visualization of changes in presynaptic function during long-term  
2857 synaptic plasticity”. In: *Nature Neuroscience* 4.7. Cited by: 259, 711 – 717. DOI: 10.1038/89498. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-0034955245&doi=10.1038%2f89498&partnerID=40&md5=b30455582b37f25fba095bbb10aad13f>.  
2858  
2859
- 2860 Zavitz, Daniel et al. (2021). “Connectivity patterns that shape olfactory representation in a mushroom body network  
2861 model”. In: *bioRxiv*. DOI: 10.1101/2021.02.10.430647. eprint: <https://www.biorxiv.org/content/early/2021/08/16/2021.02.10.430647.full.pdf>. URL: <https://www.biorxiv.org/content/early/2021/08/16/2021.02.10.430647>.  
2862  
2863
- 2864 Zeldenrust, Fleur, Boris Gutkin, and Sophie Denève (Oct. 2019). “Efficient and robust coding in heterogeneous  
2865 recurrent networks”. In: *bioRxiv* 10.6, p. 804864.
- 2866 Zhao, Xiaoliang et al. (2018). “Persistent activity in a recurrent circuit underlies courtship memory in *Drosophila*”.  
2867 In: *eLife* 7. Ed. by Kristin Scott and Catherine Dulac, e31425. ISSN: 2050-084X. DOI: 10.7554/eLife.31425.  
2868 URL: <https://doi.org/10.7554/eLife.31425>.
- 2869 Zheng, Zhihao et al. (Apr. 2020). “Structured sampling of olfactory input by the fly mushroom body”. In: *bioRxiv*,  
2870 p. 2020.04.17.047167.
- 2871 Zhu, Sijun, Ann-Shyn Chiang, and Tzumin Lee (2003). “Development of the *Drosophila* mushroom bodies: elabo-  
2872 ration, remodeling and spatial organization of dendrites in the calyx”. In: