Biol Cybern (2008) 99:89–103
DOI 10.1007/s00422-008-0241-1

ORIGINAL PAPER

Biological Cybernetics

A model of associative learning in the mushroom body

Darren Smith · Jan Wessnitzer · Barbara Webb

Received: 16 August 2007 / Accepted: 5 May 2008 / Published online: 8 July 2008

Abstract The mushroom body is a prominent invertebrate neuropil strongly associated with learning and memory. We built a high-level computational model of this structure using simplified but realistic models of neurons and synapses, and developed a learning rule based on activity dependent pre-synaptic facilitation. We show that our model, which is consistent with mushroom body *Drosophila* data and incorporates *Aplysia* learning, is able to both acquire and later recall CS–US associations. We demonstrate that a highly divergent input connectivity to the mushroom body and strong periodic inhibition both serve to improve overall learning performance. We also examine the problem of how synaptic conductance, driven by successive training events, obtains a value appropriate for the stimulus being learnt. We employ two feedback mechanisms: one stabilises strength at an initial level appropriate for an association; another prevents strength increase for established associations.

Keywords Mushroom body · Drosophila · Neural network · ADPF

1 Introduction

Mushroom bodies are paired neuropils found in invertebrate brains, characterised by prominent structures and linked to diverse and interesting roles, including sensory integration, place memory, motor control, visual navigation and certain types of learning and memory, with an emphasis on olfactory stimuli (Farris 2005; Zars 2000; Martin et al. 1998; Strausfeld et al. 1998; Ferveur et al. 1995; de Belle and Heisenberg 1994). They are one of the best studied regions of the insect brain, and a framework for explaining learning and memory processes is beginning to emerge (Margulies et al. 2005; Davis 2005; Heisenberg 2003; Waddell and Quinn 2001; Roman and Davis 2001).

In this paper, we present a concrete model of learning and memory in the mushroom body. Building such a model implies completing any gaps in existing knowledge, and to do so we incorporate models from other invertebrate studies (so that overall our model is constrained by biology). A particular gap is the nature of synapse modification underlying learning and memory, which we address by including the synaptic mechanisms underlying learning and memory in *Aplysia*. We use simplified but realistic models of neurons and synapses as model components.

A second constraint is functionality; the model had to successfully perform a learning task. This requirement drove the selection and modification of both connections, mechanisms and parameters. We selected a learning task by beginning with an hypothesis of mushroom body functionality, and then devised a simplified and representative behaviour.

In the sections that follow we begin by presenting our hypothesis of mushroom body functionality. This in turn sets the learning problem the network has to solve, and the later sections describe the model in detail and present data to illustrate it performing the task and how it does so.

2 Mushroom body function

While there is compelling evidence for mushroom body involvement in learning and memory behaviours (Heisenberg et al. 1985; de Belle and Heisenberg 1994; Dubnau et al. 2001) there is much less certainty of how, at a detailed level, it actually contributes to these behaviours. We adopt a

D. Smith (✉) · J. Wessnitzer · B. Webb
IPAB, School of Informatics, University of Edinburgh,
James Clerk Maxwell Building, The King's Buildings,
Mayfield Road, Edinburgh EH9 3JZ, UK
e-mail: mail@darrenjs.net

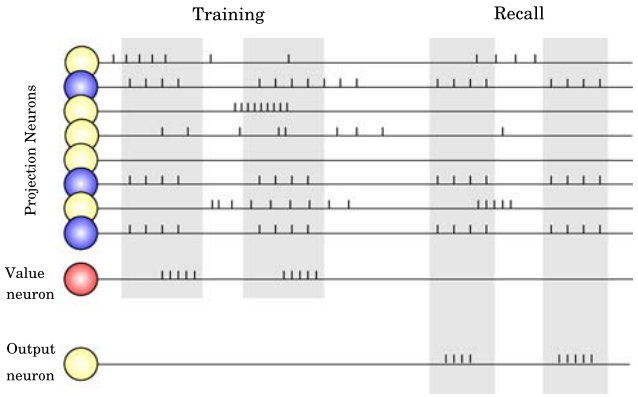

Fig. 1 The mushroom body (here represented by the single "*output*" neuron) functions to learn and later identify biologically salient patterns. During an initial training phase a particular pattern (*shaded neurons*) is learnt. This process is driven by a value signal which indicates when learning should occur (and is activated by the occurrence of a salient event in the environment). During these times the representative mushroom body output neuron does not fire; the pattern has not yet been acquired. Later, following sufficient learning, the pattern is internalised. Further presentations of the pattern result in activity of the output neuron, indicating the pattern has been detected

modelling approach whereby we first assume a specific role provided by the mushroom body. A model is then developed to meet this functional goal, while also being constrainted by mushroom body and other invertebrate data.

We suggest the mushroom body serves to acquire, and later detect, the patterns of synchronized activity in upstream sensory neurons that reliably precede (and therefore predict) the occurrence of a value signal (which itself typically signifies the occurrence of a *biologically salient* event). This is illustrated in Fig. 1. A particular set of projection neurons (which convey stimulus input to the mushroom body) fire in synchrony just prior to activity of the value signal; this set is the pattern which is acquired during learning. Once learnt, subsequent occurrences of this pattern lead to activity of an "output neuron", (the *lobe neurons*), which permits downstream neuropils to initiate appropriate reflexes. The mushroom body is the neuropil that performs this pattern learning and subsequent activation of the output neuron. A naïve mushroom body will detect no patterns; only through learning and memory processes does it acquire patterns to detect. The existence of value signals, or reward signals, in the insect brain is supported by Schwaerzel et al. (2003) who found certain neuro-modulators play an important role in *Drosophila* learning and memory. Similar neuro-modulators also exist in the honeybee brain, including prominent projections to the mushroom body (Hammer 1993; Hammer and Menzel 1998).

Such an overall function makes sense from a behavioural perspective. Any animal undergoing the daily struggle to survive and reproduce will be strongly interested in salient events that either endanger or support these aims. Animals that can predict when such events are about to occur will gain a clear advantage over competitors that are less able to do so. Liu et al. (1999) also consider the problem faced by animals living in the "real world" which use learning as an adaptive mechanism to survive: "animals need to extract from the universe of sensory signals the actual predictors of salient events by separating them from non-predictive stimuli (context)". Their experiments indicate that the mushroom body may support this process.

3 Model description

Figure 2 illustrates the complete mushroom body model. The core components are the calyx (excitatory synapses), Kenyon cells (KCs), and the synapses of the lobes. Also shown are the nearby networks for input and output. The model has a two layer architecture. The first layer—the calyx and KCs—transforms the spatial code of projection neuron (PN) activity into a sparser spatial code at the Kenyon cells. The second layer—the lobe synapses and lobe neurons—act as detectors for particular spatial codes among Kenyon cells.

The model was implemented using "simplified" models of neurons and synapses; a spiking integrate and fire model is used for neurons, and a conductance based model for synapses (Koch 1999). Such simplifications are necessary when computing large networks of neurons and serve to highlight the minimal features necessary to support network properties (Trappenberg 2002 p. 13). This approach is commonly termed "biologically inspired", which places emphasis on replicating general principles rather than detailed physiology.

Although the architecture and learning mechanisms used in the model were strongly inspired by *Drosophila* data, the objective of the simulation was to demonstrate the conceptual model of layer interactions and synapse learning rule, rather than attempting to replicate the real fly-brain. Model size therefore reflected what was minimally necessary for interesting scenarios to be investigated (e.g., multiple input stimuli), and what was tolerable in terms of the duration of simulation execution. The later point is particularly important because some parameter values were found by a manual process; we needed simulation runs to complete in hours rather than days. Simulation execution time depends on the total network size, and here our architecture compounds the problem of scaling; if instead of 630 KCs we matched *Drosophila's* 2500 (Heisenberg et al. 1995), the number of KC to lobe neuron synapses would increase from ≈16,000 to ≈65,000. It is also important to note that the size of some layers result from our interlayer connectivity rule; e.g., the number of KCs is determined by the number of upstream PNs. Our present focus is to investigate the behaviour that follows from such rules alone, rather than how these rules

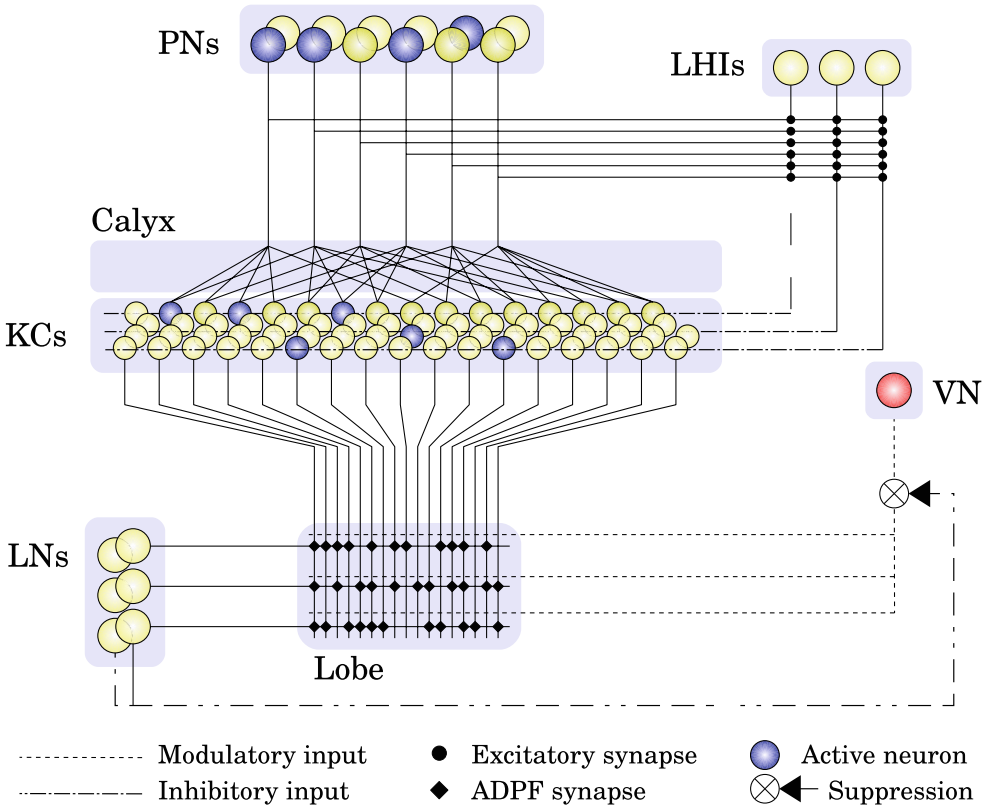

Fig. 2 Overview of the mushroom body model and the input/output networks it connects to. Input is provided by 36 projection neurons (*PNs*). These connect to 630 Kenyon cells (*KCs*) in a *pair-wise* fashion: for every combination of two PNs, there is a KC which receives input from solely those two PNs. Each KC acts as a coincidence detector; it spikes only if its upstream PN pair fire in loose synchrony. This leads to a sparse activity pattern at KCs, which aides learning and memory. Learning and memory processes occur at the ADPF lobe synapses which connect KCs to 40 lobe neurons (*LNs*). ADPF is a form of value driven learning; synapse strengthening occurs when the synapse and value neuron (VN) are simultaneously active. Each KC–LN pair is connected with probability 0.65, leading to ≈16,000 lobe synapses

would need to be adjusted to fit within more realistic network sizes.

3.1 Input transformation

The ability of the network to sparsen an earlier spatial activity pattern rests on the PN-to-KC connectivity and the role of individual Kenyon cells. Actual *Drosophila* mushroom bodies exhibit a highly divergent connectivity; 50 antennal lobe glomeruli project to 2,500 Kenyon cells (Heisenberg et al. 1995). Our model uses an idealised connectivity scheme which serves to highlight the *principle* of high divergence. PN-to-KC connectivity is based on each Kenyon cell acting as a simple coincidence detector for a pair of PN inputs; a KC will only fire if both inputs are synchronously active. A straightforward method to connect a population of PNs to a downstream population of KCs is to mandate that each combination of two PNs projects to a unique Kenyon cell. This *pair-wise* connectivity leads to a hugely divergent connection, and also determines the number of KCs required to support a given number of PNs (36 PNs requires 630 Kenyon cells and 1,260 synapses). A further consequence of this pair-wise connectivity and KC coincidence detection role is that activity transformation is entirely deterministic.

Figure 2 illustrates sparsening occurring. The active PNs (shaded blue) account for 33% of their population, however the pattern they affect at the KC layer accounts for only ≈9% of the population. In general a sparseness of S at the PN layer will result in sparseness of S^2 at the KC layer.

3.2 Coincidence detection

Kenyon cells function as coincident detectors by way of appropriate parameter adjustment for both neuron and excitatory PN-to-KC synapses. Parameters for each KC (listed in the appendix) were chosen to reflect other simulation studies (Wüstenberg et al. 2004; Gingrich and Byrne 1987, 1985). Parameters for the excitatory PN-to-KC synapses were found empirically. This was done by tuning their values to achieve robust coincidence detection for a pair of single spikes, and also for a pair of loosely synchronous spike trains arriving at

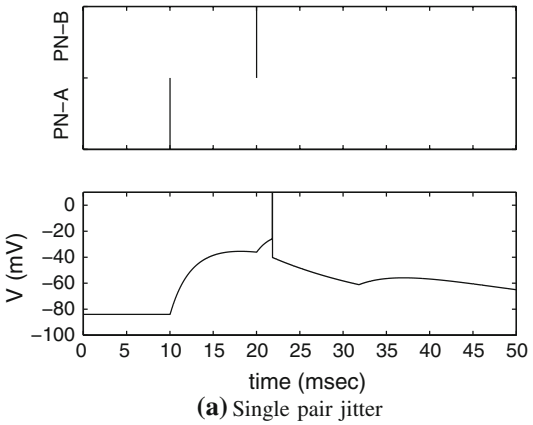

(a) Single pair jitter

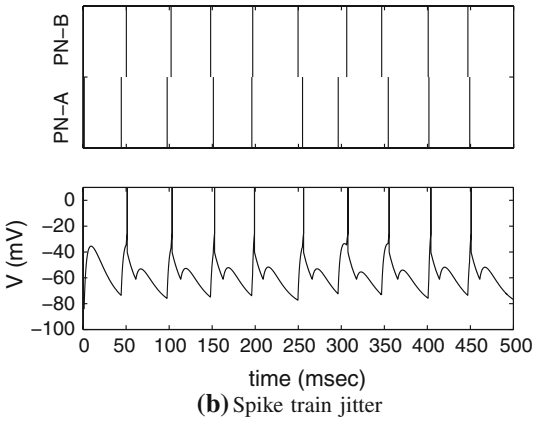

(b) Spike train jitter

Fig. 3 a Input spikes (*top*) for a pair of PNs, PN-A and PN-B, do not have to be perfectly synchronised to trigger KC $V(t)$ (*bottom*) to threshold. Because of temporal processes in both the synapse and neuron models, coincidence detection is robust to some spike jitter, i.e., a separation of several milliseconds. This also shows IF neuron behaviour post spiking. $V(t)$ is immediately reset to V_{recov}, which is set just below V_{thresh} (see Table 1). Then follows a brief refractory period (10 ms) during which end plate currents from pre-synaptic spikes are ignored. When normal behaviour resumes "after potentials" are observed; these are due to residual levels of end plate currents. b Response to two loosely synchronous spike trains. The KC is still able to perform coincidence detection

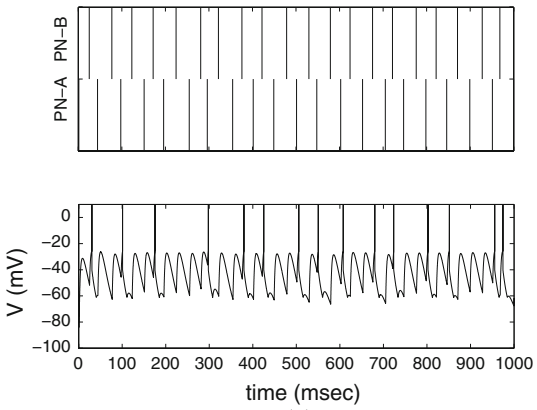

(a)

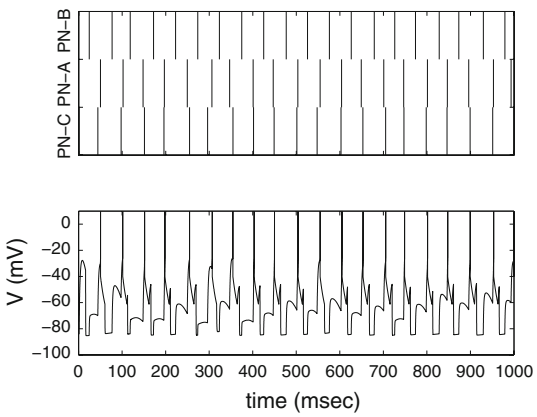

(b)

Fig. 4 a KC mis-firing arising from two out-of-phase spike trains. There are no occurrences of loosely synchronous spike pairs, but the KC coincidence detector still (wrongly) fires. b KC mis-firing is eliminated by synchronized inhibition from LHI neurons. Inhibition arrives shortly following a global population response (here signified by PN-C and PN-A synchrony). A feature of this inhibition is that it is very powerful. KCs are strongly polarised, preventing the PN-B out-of-phase spike train raising $V(t)$ to threshold. The KC only spikes for PN-C and PN-A synchrony

20 Hz (Fig. 3), which is based on the frequency of projection neuron spiking in the Locust (Wehr and Laurent 1996). The end result is an upper and lower bound for synapse intrinsic conductance g_{syn}. Values for individual synapses were drawn from this range (with $g_{\text{mean}} = 1.3$ nS and $g_{\sigma} = 0.1$ nS). These synapses were also fast acting, with a conductance decay half-life of 5 mS.

Robust KC performance could not be achieved solely via parameter selection. Problems arise when spike trains arrive completely out-of-phase. In such cases we observed occasional KC firing, even though spike pairs were not synchronous (Fig. 4a). To solve this problem the functionality of lateral horn inhibitory (LHI) network was introduced to provide regular bursts of inhibition to the KC layer. This inhibition served to reset the KCs following the most recent detection episode (Fig. 4b), and is implemented through powerful inhibitory synapses which are activated whenever a PN population response occurs (i.e., when a large proportion of the PNs begin to fire in sychrony). A delay of 15 ms is introduced between PN population response and inhibition occurrence, which gives the KCs time to integrate PN spikes before inhibition arrives.

Building a biologically plausible LHI network from neuron and synapse models to detect and respond to PN population responses is a non-trivial task. This was not attempted for the present study, although work on antennal lobe modelling by Bazhenov et al. (2001a,b) suggests an approach for solving this kind of problem. Instead, we replicated the *functionality* of such an LHI network. This was possible because the input stimuli (i.e., patterns of PN firing) were under experimental control, and so the output of an LHI network could

be replaced by an LHI signal that was coupled (and delayed, as noted above) to the pre-determined PN patterns.

As noted earlier, the combined effect of KCs acting as coincidence detectors and a hugely divergent connectivity from PNs to KCs serves to sparsen patterns of activity. Sparse coding is typically associated with greater memory and recall efficiency (Olshausen and Field 2004), and in the current case, such coding leads to improvements in learning and memory when similar patterns must be discriminated, which can be illustrated with a numerical example. Consider two patterns both of size 10 with an overlap of 70%. With such a high overlap an animal could easily mis-identify one pattern for another, leading to potentially life threatening wrong decisions. However, via the pair-wise transformation, this overlap is reduced to ≈47% at the KC layer. Thus the animal will be better able to discriminate KC patterns, and it is the patterns of this layer that feed to the next layer, where learning, memory and recall take place.

3.3 Associative learning

The second layer of the model is representative of the mushroom body lobe. Its chief component are the lobe synapses, which convey input from the upstream KCs to 40 lobe neurons. Any KC is connected to any lobe neuron with a probability of 0.65. The lobe also receives input from the external value neuron. Collectively the lobe provides the learning and memory ability of the model.

The lobe serves two roles. The first is to detect patterns of KC activity. Each lobe neuron may become associated with one or more KC spatial activity patterns; if one of these pattern arises, the associated lobe neuron will start to spike. These lobe neuron spikes in turn trigger a behaviour that is appropriate to the CS being presented, i.e., we assume the lobe neurons are connected to "command neurons" (Carew 2000) that can initiate reflex actions (not shown in Fig. 2). Through these roles the mushroom body is able to associate a CS (KC patterns) with a conditioned response (lobe neuron spiking).

The second role of the lobe is the formation of these associations. Initially each lobe neuron will not detect any KC patterns. It is only through learning and memory processes that lobe neurons are recruited as detectors. The basis of these processes is purposeful change to lobe synapse strength. Memory acquisition corresponds to strength increase; sufficiently strong lobe synapses permit KC spikes to depolarise lobe neurons to supra-threshold levels, inducing spiking (and in turn an appropriate reflex action).

In choosing a learning rule for the lobe synapses we faced a synapse *weight dilemma*. Because arbitrary KC patterns could be of any size, there is no uniform value that synapse strengths should tend toward during learning. Rather, for a useful association to be made, the final lobe synapse strength must correlate closely with the size of the KC pattern; too low and the association will not be recalled, but too high and the association will be recalled for partial presentations of the pattern.

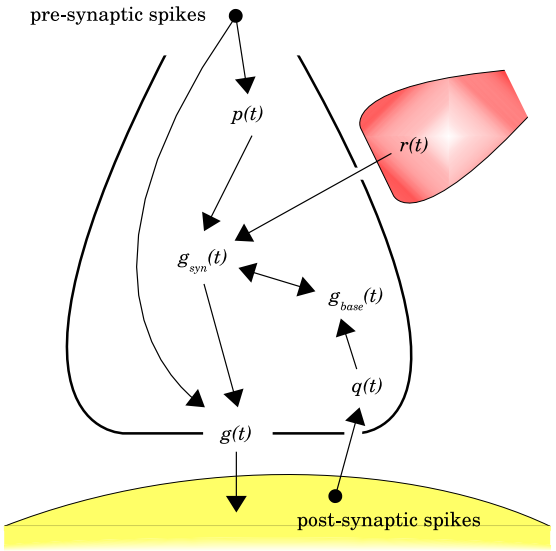

Fig. 5 The variables and interactions of the ADPF based synapse model. A pre-synaptic spike train drives the pre-synaptic history trace $p(t)$, which together with modulatory input $r(t)$ drives growth of intrinsic conductance $g_{\text{syn}}(t)$. Each spike increases the actual time varying conductance occurring at the target neuron, $g(t)$; these increments are equal to the value of $g_{\text{syn}}(t)$ at the spike times. Spiking activity at the post-synaptic neuron results, via a retrograde signal, in a second history trace $q(t)$ which in turn drives $g_{\text{base}}(t)$ to quickly approach $g_{\text{syn}}(t)$. Outside of learning episodes $g_{\text{base}}(t)$ decays very slowly to zero and $g_{\text{syn}}(t)$ decays quickly toward $g_{\text{base}}(t)$; through this relationship $g_{\text{base}}(t)$ indicates longer term memory. The quantity $p(t)$ is loosely analogous to an intra-cellular Ca^{2+}trace elevated by CS activity, while $r(t)$ relates to US induced serotonin input

3.3.1 Synapse model

A synapse model was built with this problem in mind, and also with the motivation to remain consistent with *Drosophila* and other invertebrate biology. This model is depicted in Fig. 5, which shows the time varying quantities and interactions underlying learning, memory and signal transmission.

The synapse uses a conductance based model of signal transmission. Pre-synaptic spikes drive transient increments to an instantaneous conductance $g(t)$ according to,

$$\frac{dg}{dt} = \frac{-\ln 2}{\lambda_{\text{syn}}} g + \sum_{t_{\text{p}}} \delta(t - t_{\text{p}}) g_{\text{syn}}(t)$$

where λ_{syn} is $g(t)$ decay half-life, and t_{p} are the pre-synaptic spike times. This equation (and others like it encountered below) was simulated using an update rule,

$$g(t+\Delta t) = g(t)e^{\frac{-\ln 2}{\lambda_{\text{syn}}}\Delta t} + n_{\text{p}}(t,\Delta t)g_{\text{syn}}(t)$$

where $n_{\text{p}}(t,\Delta t)$ is the number of pre-synaptic spikes between t and $t+\Delta t$.

Whereas $g(t)$ is a measure of the present activity of the synapse, $g_{\text{syn}}(t)$ corresponds to intrinsic strength. Learning and memory processes are introduced through equations involving this latter quantity.

No complete mechanistic account of synaptic plasticity in the *Drosophila* mushroom body exists (i.e., a description of how multiple events in different neurons qualitatively affect strength). To overcome this gap, we modelled learning and memory processes on ADPF—*activity dependent pre-synaptic facilitation* (Hawkins et al. 1983; Walters and Byrne 1983). This is the mechanism underlying sensitisation and classical conditioning of defensive reflexes in *Aplysia*. Briefly, a sensitising stimulus (the US) causes release of the modulatory neurotransmitter serotonin onto sensory neuron terminals, where it binds to G-protein-linked receptors. This initiates an intra-cellular biochemical cascade: first the production of second messenger cyclic AMP (cAMP) is stimulated; increased levels of cAMP next activate protein kinase A (PKA); PKA promotes the phosphorylation of other proteins, and those it targets are thought to include K^+channels, the phosphorylation of which can lead to their closure. Closed K^+channels result in broadened action potentials in the sensory neuron, which in turn leads to greater neurotransmitter release per spike (i.e., an enhanced synapse strength). The net consequence is that an enhanced excitatory post synaptic potential (EPSP) is invoked in neurons post-synaptic to the sensory neuron.

This mechanism accounts only for facilitation. Activity dependent facilitation describes US interaction with a preceding CS. The main concept is that activity in sensory neurons (the CS) immediately before the arrival of serotonin neurotransmitter (the US) results in an amplification of the sensitisation effect that serotonin alone induces. Successive sensory neuron action potentials transiently elevate intra-cellular Ca^{2+}, which in turn activates a calmodulin dependent "priming" of a serotonin-sensitive adenylyl cyclase (AC). This primed cyclase causes a greater production of intra-cellular cAMP upon subsequent coupling of serotonin to sensory neuron receptors. Increased cAMP levels amplify the remaining steps in the sensitisation cascade, finally resulting in enhanced EPSPs. The key difference between the cellular mechanisms for sensitisation and classical conditioning is this "priming" of cAMP production by preceding sensory neuron (CS) activity. The pre-synaptic neuron is the site of CS and US convergence and where initial learning-associated changes occur.

Core concepts of this mechanism motivate the synapse model of Fig. 5. Intra-cellular Ca^{2+}allows early occurring CS related pre-synaptic spike activity to influence a later occurring US. This is modeled as a *history trace* of spike activity,

$$\frac{dp}{dt} = \frac{-\ln 2}{\lambda_{\text{p}}}p + \sum_{t_{\text{p}}}\delta(t-t_{\text{p}})$$

where λ_{p} is decay half-life and t_{p} are the times of pre-synaptic spikes.

According to ADPF, enhanced strengthening occurs when this spike activity interacts with modulatory input. In Fig. 5 such input is represented by $r(t)$, and this interacts with $p(t)$ to modify intrinsic strength according to,

$$\frac{dg_{\text{syn}}(t)}{dt} = r(t)p(t)\eta_{\text{syn}} - \left(g_{\text{syn}}(t) - g_{\text{base}}(t)\right)\kappa_{\text{syn}} \qquad (1)$$

where η_{syn} is the growth rate, κ_{syn} the decay rate, and $g_{base}(t)$ the decay target.

Learning events (i.e., when $r(t) > 0$ and $p(t) > 0$) cause immediate change in $g_{\text{syn}}(t)$. However, this initial learning is not long lasting; $g_{\text{syn}}(t)$ quickly decays to $g_{\text{base}}(t)$, which itself represents longer lasting memory.

Consolidating initial learning into memory is based on further *Aplysia* neurobiology. Learning and memory in *Aplysia* is not an entirely pre-synaptic process. Antonov et al. (2003) provide direct evidence that together with ADPF, there is a longer term post-synaptic component to learning, and that these interact with each other. A revised "hybrid" model of associative learning (Antonov et al. 2003; Lechner and Byrne 1998) has post-synaptic Hebbian LTP interacting with the pre-synaptic ADPF process via a postulated retrograde signal; this signal arises from post-synaptic increases in Ca^{2+}, which in turn results from activation of post-synaptic NMDA receptors. Roberts and Glanzman (2003) suggest that learning associated changes in sensory neurons could be made more persistent by this trans-synaptic signal, and are essential for longer term memory.

We incorporate this concept of a retrograde signal being critical for consolidation into the synapse model. The retrograde signal is $q(t)$ in Fig. 5, and is just a history trace of post-synaptic activity,

$$\frac{dq}{dt} = \frac{-\ln 2}{\lambda_{\text{q}}}q + \sum_{t_{\text{q}}}\delta(t-t_{\text{q}})$$

where λ_{q} is decay half-life, and t_{q} are the times of post-synaptic spikes.

The memory term $g_{\text{base}}(t)$ evolves according to,

$$\frac{dg_{\text{base}}}{dt} = \left(g_{\text{syn}}(t)\alpha_{\text{syn}} - g_{\text{base}}(t)\right)q(t)\eta_{\text{base}} - g_{\text{base}}(t)\kappa_{\text{base}}$$

where η_{base} is the growth rate, κ_{base} the decay rate and α_{syn} influences the target value that $g_{\text{base}}(t)$ will grow toward. The driving process is the difference between short term conductance $g_{\text{syn}(t)}$ and longer term conductance $g_{\text{base}}(t)$, and such

differences occur shortly after learning events at the synapse. Before initial learning is forgotten (i.e., before $g_{\text{syn}}(t)$ decays to base levels), the above equation allows $g_{\text{syn}}(t)$ to "pull" $g_{\text{base}}(t)$ toward it. This pulling-up process is the consolidation of learning into longer term memory, since changes to $g_{\text{base}}(t)$ will decay more slowly than changes to $g_{\text{syn}}(t)$. For this volatile to persisted transfer to happen there must be ongoing activity of the retrograde signal $q(t)$. In the simulation runs $\alpha_{\text{syn}} = 1.05$ was used. This allowed $g_{\text{base}}(t)$ to very slightly overshoot $g_{\text{syn}}(t)$, leading to more robust memory formation. However this has a problem; later activation of the strengthened synapse that leads to post-synaptic firing will result in further slight increments of its strength. In the present model this can continue without limit. One solution is to take the approach of other synapse models, where the weight variable [in this case $g_{\text{base}}(t)$] is restricted to hard limits (Damper et al. 2001).

3.3.2 *Feedback mechanisms*

The above model solves one aspect of the weight dilemma. Due to memory consolidation requiring the retrograde signal and because this signal is only active once lobe synapses have been sufficiently strengthened to trigger lobe neuron spiking, consolidation will always tend toward a strength that is sufficient for activating lobe neurons.

The retrograde signal is one of two feedback mechanisms. The second form of feedback prevents synapses from becoming over-strengthened. It takes the form of an inhibitory pathway from lobe neurons to a relay on the value neuron (VN) pathway (illustrated in Fig. 2). Once lobe synapses possess sufficient strength in order to trigger lobe neuron spikes, the value signal is suppressed (via the feedback). This causes the modulatory input $r(t)$ to become zero, and so no further strengthening takes place; lobe synapses remain at a strength appropriate for the size of the KC pattern just acquired.

The overall performance of $r(t)$ has similarity with the idea that expectation of a reward/punishment influences learning. If a US is *predicted* (i.e., if a lobe neuron is firing), then its immediate occurrence does not cause learning. However if the US occurs unexpectedly, learning takes place.

Aside from achieving an appropriate synapse strength, the overall learning network has the advantage of being very selective in the synapses that undergo consolidated enhancement. Only synapses that conveyed the CS pattern *and* connect to a lobe neuron that fired, undergo consolidated strengthening. This helps to minimise the introduction of spurious associations (i.e., accidental associations between KC patterns and lobe neuron firing that are not based on experience).

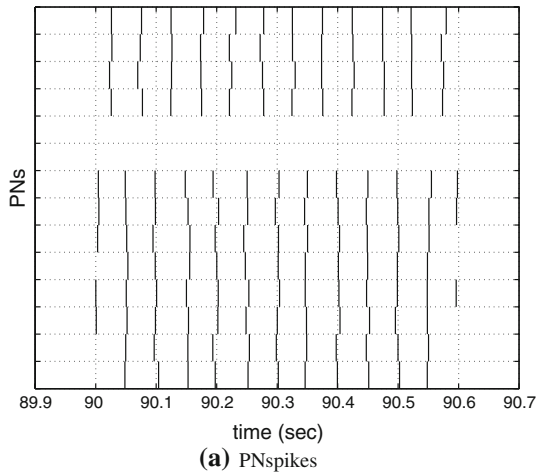

(a) PNspikes

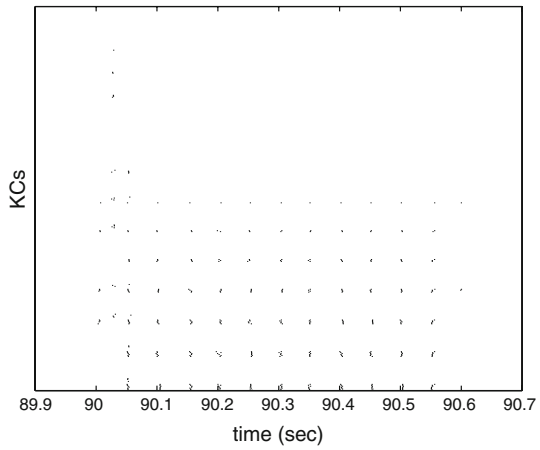

(b) KCspikes

Fig. 6 a Spiking activity of 14 out of 36 PNs. Each PN is *plotted* on a different row, and a short *vertical line* signifies a single spike. The spike trains are loosely synchronous, and one set of 4 PNs is firing out-of-phase to another set of 8. b Corresponding response of the 630 Kenyon cells. Each KC is plotted on a separate row, although most are inactive and so produce no spikes. Due to the strong LHI inhibition the KC firing quickly reflects only activity of the 8 PNs, and not the 4 that are out-of-phase. Very few of the Kenyon cell population fire, illustrating the sparse response

4 Results

The model was repeatedly evaluated as a computer simulation in order to refine values for selected neuron and synapse parameters until coincidence detection and learning behaviours were reliably established (other parameters were based on values taken from biological and simulation literature, as explained in the tables in the appendix). Data from selected runs are next presented to demonstrate these behaviours.

4.1 Coincidence detection

The response of KCs to synchronised input from PNs is shown in Fig. 6. Input comes from two active subsets of

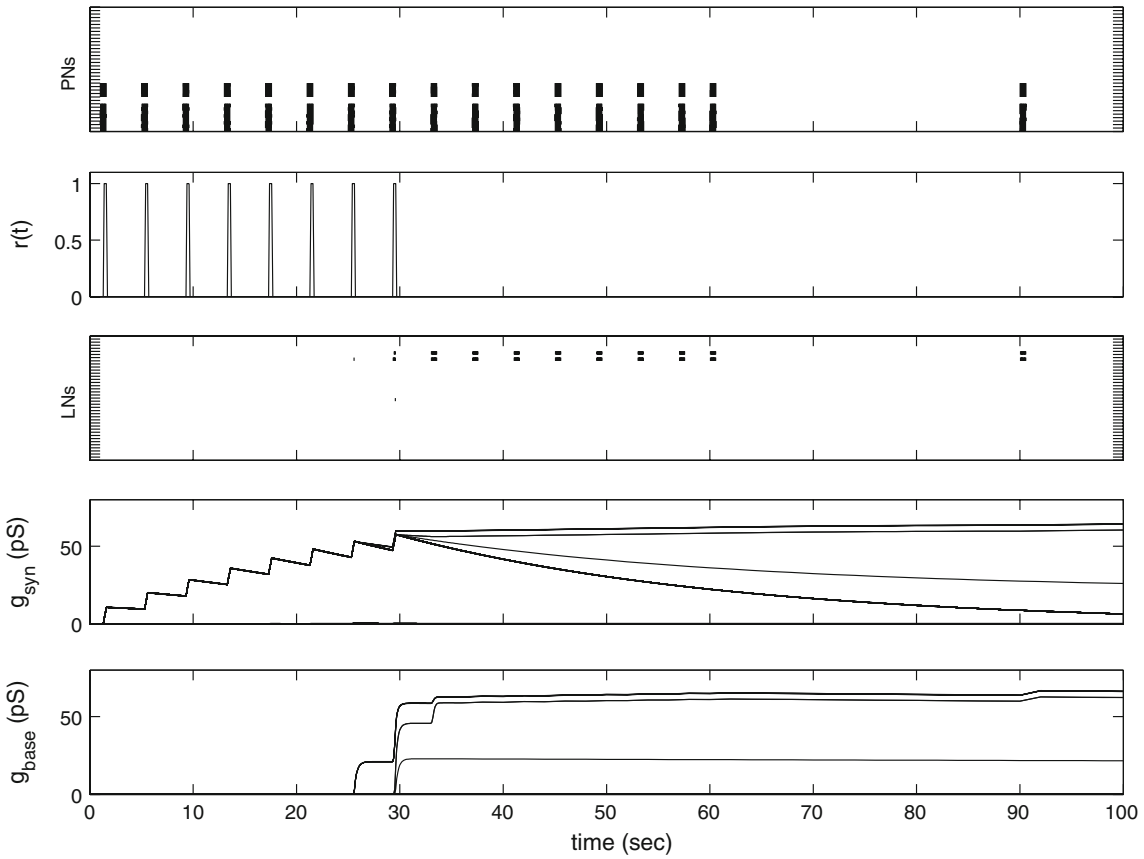

Fig. 7 Learning demonstration. The *upper plot* shows 12 projection neurons producing a series of short spike trains which are accompanied by US input, $r(t)$. The *middle graph* shows lobe neuron activity; at $t \approx 30$ s several start reliably firing, which indicates an association has formed. This is caused by increasing lobe synapse strengths reaching a level sufficient to trigger lobe neuron firing (*bottom graphs*)

PNs. Out of a population of 36 (only 14 are shown in 6a), a first subset of 8 (22% of population) is synchronously active, while a second subset of 4 is also synchronously active but out-of-phase to the first. Activity takes the form of a brief spike train with duration 600 ms and frequency ≈20 Hz. All other PNs are inactive.

Individual KCs emit single spikes at times strongly coupled to PN spikes. A brief deviation from this regularity occurs at the start of the response (at ≈90.025 s), which is due to the LHI inhibition not yet being effective. Moving from individual KCs to the KC population as a whole, the spareness of the response is clearly evident. Very few KC fire; the Y-axis of 6(b) represents all 630 KCs, but only ≈28 KCs are spiking, which is less than 5% of the population.

4.2 Learning

Figure 7 demonstrates learning. Synchronous PN activity is delivered together with a value signal $r(t)$, which consists of a sequence of pulses (brief periods of $r(t) = 1$, while at all other times $r(t) = 0$) each 300 ms in duration and following CS onset by 300 ms. Lobe neurons do not initially spike (demonstrating that initially there is no association based on the PN patterns), but a pair start to do so at ≈30 s. This change from no-response to spiking occurs because of learning-driven synaptic strengthening, which is shown in the plot of synapse strength $g_{\text{syn}}(t)$.

The presence of a learning process is evident in the nature of $g_{\text{syn}}(t)$ changes. These consist of abrupt increases that occur with each training event [i.e., occurrence of $r(t)$ pulses]. Each increase is followed by a period of quick decay as $g_{\text{syn}}(t)$ falls back to the base level $g_{\text{base}}(t)$. However $g_{\text{syn}}(t)$ does not fully decay to zero; training events occur in close enough succession to raise $g_{\text{syn}}(t)$ to increasingly higher levels. A threshold effectively exists at ≈60 pS; this is the level at which a particular group of active lobe synapses are sufficiently strong to induce spikes in lobe neurons. This threshold is not a fixed parameter of the model; it rather is determined by the number of lobe synapses that convey a particular KC pattern to a lobe neuron, and is hence dependent on the size of the KC pattern, which is illustrated in Fig. 8.

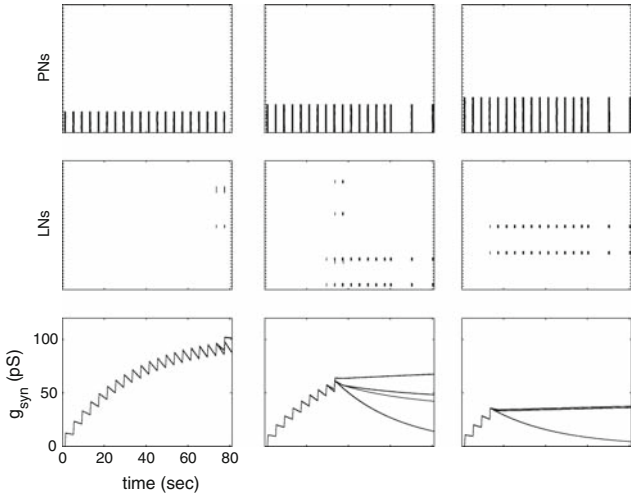

Fig. 8 Relationship between pattern size and synapse strength. For larger PN patterns, the lobe synapse strength required to form an association is lower (*right hand column*); for smaller patterns, the required synapse strength is higher (*left hand column*). The speed of acquisition is also influenced by pattern size; lobe neurons begin spiking earlier for larger PN patterns

The start of regular spiking by 2 lobe neurons indicates that an association has been acquired (and the feedback mechanism from lobe neurons to VN means that this activity shutsdown the modulatory input). Following this point two main pathways exist for $g_{\text{syn}}(t)$. The synapses connected to the 2 spiking lobe neurons receive the retrograde signal which induces longer term memory, and subsequent $g_{\text{syn}}(t)$ decay is greatly reduced (indeed it is not even apparent from the figure), because the longer term conductance $g_{\text{base}}(t)$ has been rapidly "pulled up" to around 60 pS. The association between the CS pattern of 8 PNs and activity in these 2 lobe neurons has been consolidated into longer term memory via changes in these synapses. Persistence is illustrated by the later presentation of the CS pattern, which is able to activate the lobe neuron pair.

Nearly all other lobe synapses—including those strengthened but not connected to the active lobe neurons—undergo $g_{\text{syn}}(t)$ decay toward base level of zero, illustrated in Fig. 9a. These decay to zero because that is the initial value of $g_{\text{base}}(t)$, and for these synapses $g_{\text{base}}(t)$ was never increased by consolidation events (i.e., spiking of their post-synaptic lobe neuron). The $g_{\text{base}}(t)$ trace in Fig. 7 shows that a small group of synapses stabilized at ≈20 pS; this indicates partial consolidation, and is due to their post-synaptic neuron not firing a sufficient number of times in order to induce full consolidation.

$V(t)$ for one of the lobe neurons that became active is plotted in Fig. 9b, and illustrates another effect of increasing synapse strength. In response to successive volleys of KC spikes (4 s period) the lobe neuron becomes increasingly depolarised, which is due to the increasing strength of the lobe synapses. Eventually synapses are strong enough to cause depolarisation to exceed threshold, and the neuron fires.

4.3 Partial CS

As noted earlier, the value of $g_{\text{syn}}(t)$ attained during consolidation is determined by the size of the input pattern presented during learning. A consequence is that the implicated lobe neuron will not be activated for partial presentations of the learnt pattern. This is a useful property, since an association is between the learnt stimulus and the US, i.e., between the *whole set* of PNs activated by the stimulus and the US; an

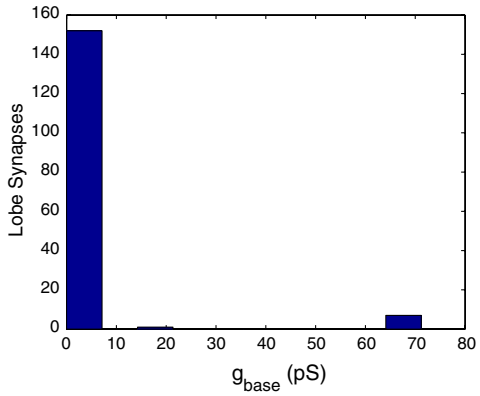

(a) $g_{base}(t)$ distribution

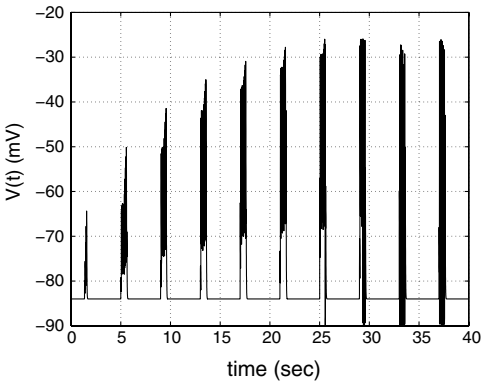

(b) Lobe neuron $V(t)$

Fig. 9 a Histogram of lobe synapse $g_{\text{base}}(t)$ at $t = 100$ s. Very few synapses underwent consolidation, indicating the highly selective nature of synapse learning (only 10% of lobe synapses linked to KCs that conveyed the two PN patterns in Fig. 7 were recorded.) b $V(t)$ for a single lobe neuron during the acquisition of an association. As successive training events drive synapse strength growth, KC activity (occurring at 4 s intervals) causes increasing depolarisation. Threshold is first exceeded at ≈25 s [indicated by $V(t)$ reseting to the recovery potential of −90 mV], and regular spiking begins shortly afterwards

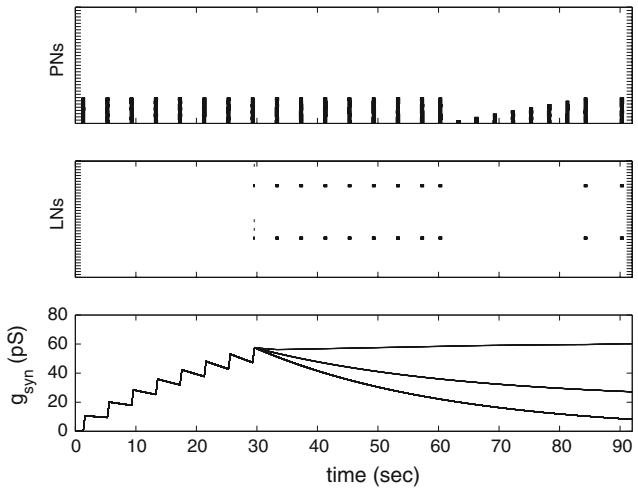

Fig. 10 Presentation of partial stimuli to test recall of an association. After an association has been acquired for a PN pattern of eight neurons, partial patterns are presented (between 60 and 90 s). None are able to cause firing of lobe neurons; partial patterns fail to recall the association

association should not be recalled by some arbitrary subset of the CS.

Figure 10 demonstrates this feature of the model. During an initial training phase an association develops between a pattern of 8 PNs and two lobe neurons; acquisition occurs at ≈30 s, and the memory still exists when tested at 90 s. Synapse strength stabilises at ≈60 pS, the level suitable for an input size of 8 PNs. The test phase comprises the successive presentation of partial PN patterns. The first such stimulus, at $t = 63$ s, consists of only a single PN (12.5% of the CS), and subsequent presentations are steadily increased until the whole pattern is presented at $t = 84$ s. The figure shows that lobe neurons do not fire for any partial-CS. Even when 7 out of 8 PNs are firing (87.5% of the CS) the association is still not recalled. This is because synapse strengths are at the *minimum* value for a PN pattern of size 8 (corresponding to a KC pattern of ≈28 neurons) to trigger a lobe neuron. Even for large subsets of the CS, the corresponding KC pattern is significantly reduced; e.g., when the PN pattern is reduced to 87.5% of the CS, only ≈21 KCs are activated (or 75% of the original 28). This large reduction in the number of active KCs results in insufficient synaptic input to induce lobe neuron spiking.

4.4 Suppressing mushroom body output

We next introduced specific mushroom body impairments. Actual conductance $g(t)$ was selectively held at zero for particular time intervals. This impairment is motivated by *Drosophila* mushroom bodies studies in which output from mushroom body neurons was transiently blocked (Dubnau et al. 2001; McGuire et al. 2001).

Figure 11a shows the situation when mushroom body output is blocked during training; i.e., there is no lobe synapse output for the first 60 s of the experiment. During this period the PNs and the US are being presented as per the established training routine, and lobe synapse strengthening still occurs as is demonstrated by the increasing $g_{syn}(t)$. Due to $g(t) = 0$ for this period, lobe neurons receive no stimulation; there is no lobe neuron spiking. Consequently the feedback control mechanism—for limiting the increasing strength—is broken, and learning becomes distorted in the sense that $g_{syn}(t)$ exceeds the approximate value of 60 pS attained during normal operation of the model.

When synapse output is permitted again, at $t = 60$ s, many lobe neurons instantly begin firing, more than would occur under normal conditions. This extra volume of response is due to the enhanced synapse strengthening that occurred during training; i.e., even for lobe neurons receiving a smaller number of synaptic connections, the overall effect, when combined with their excessive strength, is sufficient to cause firing. This lobe neuron activity in turn causes retrograde signalling, which consolidates the distorted synapse strength into longer term memory.

The second run has lobe synapse output blocked between 60 and 90 s, i.e., during the test period (Fig. 11b). During training synapse behaviour is normal, and consequently an association develops at $t \approx 25$ s, which is consolidated during the remainder of the period.

During the test period there is no lobe neuron spiking; i.e., there is no recall of the association that developed during training. This is not very surprising, since lobe neuron stimulation is driven by lobe synapse $g(t)$, which for this period is fixed at zero. This period also continues to present the US (not shown) along with the stimulus, which results in further strengthening of synapses. The resulting effect can be seen when lobe synapse output is once more permitted at $t = 90$ s; the number of lobe neurons responding has increased.

Taken together these demonstrate the requirement of lobe synapse signalling during learning: output is required during recall, but not during learning, which matches *Drosophila* observations. However once the messy details of a neural network are considered, this conclusion has some caveats. For instance, the suppression of output may lead to distorted learning (which might not be immediately detectable at the behavioural level). Further, while synaptic transmission is not required for acquisition of an association, in our model it is required for consolidation (the retrograde signal). This prediction could be examined in actual *Drosophila* by training flies during a period of KC transmission restriction, and later testing for longer term (24 h) memory. We would expect such memory to be significantly impaired due to the absence of any prior consolidation events.

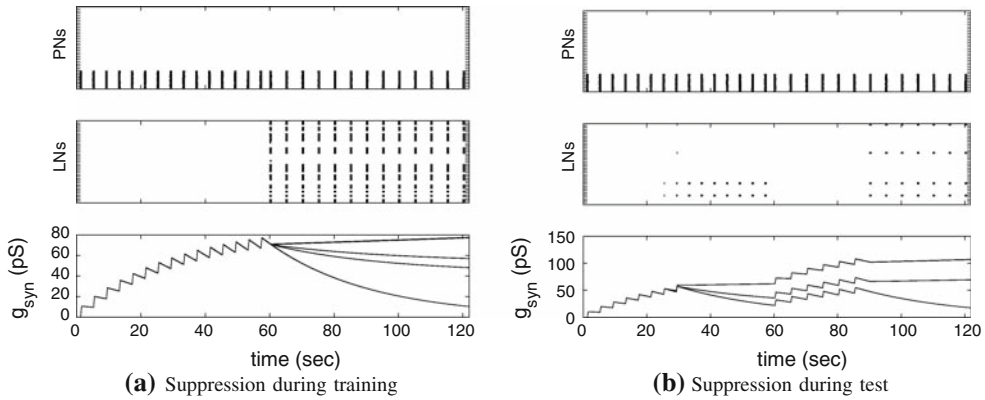

(a) Suppression during training **(b)** Suppression during test

Fig. 11 a Suppression of mushroom body output during training. An association still develops but it is now distorted; synapses become overly strengthened, so that many lobe neurons now fire for later presentations of the CS. b Suppression during test (60 − 90 s). During the preceding training phase an association successfully develops; however association is not recalled while KC output is suppressed

5 Discussion

In this paper, we have aimed to give a mechanistic account of learning and memory in the *Drosophila* mushroom body. A simplified model was built to emphasise the higher level processes involved. In developing the model we had the aim of making it achieve a particular behaviour. This behaviour is our hypothesis for mushroom body function: it serves to acquire associations based on stimuli patterns that predict biologically salient events, and later detects for such learnt patterns. Our computer simulations demonstrates the model is able to perform such learning.

Our model has a clear functional separation between the KC layer and the lobe neuron layer. We assume a relatively simple role for individual Kenyon cells; each acts as a coincidence detector for a pair of inputs. Although the nature of this connectivity is somewhat idealised (i.e., every possible set of two PNs is represented by a unique Kenyon cell), the hypothesis supported here is that actual Kenyon cells do offer a very similar unit of functionality, i.e., they detect for coincidence across a small number of inputs. We address the specific problem of KC mis-firing from extended spike trains, or out-of-phase spike trains. Appropriate parameter adjust solves the first problem, while the strong inhibition provided by the lateral horn solves the second.

This coincidence detector role builds upon current ideas in the field. A model developed by Huerta et al. (2004) similarly has "individual KCs" serving as "coincidence detectors for the synchronous PN input." In his review of mushroom body function Heisenberg (2003) discusses an olfactory role in which he assumes "Kenyon cells are coincidence detectors" that respond to activity in "three projection neurons simultaneously". The idea that the LHI provides a gain control mechanism also extends existing ideas. Nowotny et al. (2005) employ "gain control through the known feed-forward inhibition of lateral horn inter-neurons," concluding that it "increases the capacity of the system but is not essential for its general function". The present work assigns a similar but more specific role; the LHI input is precisely timed, and strongly weighted, to disrupt the influence of PN spike trains that are out-of-phase to the population response. The inhibition serves to effectively reset KCs between coincidence detection episodes. If LHI neurons were disrupted in actual *Drosophila*, during either learning or recall phases, then we would expect learning performance to be significantly impaired.

Our model demonstrates the sparsening of a PN activity pattern to a KC pattern, which is due to our pair-wise connectivity between layers. Other mushroom body models suggest this sparsening enhances learning; we too support this conclusion, but also offer qualitative and quantitative support. It was shown a sparse representation aids in the situation of preventing partial patterns from triggering an association through the large reduction in KC overlap between full and partial patterns. This is advantageous for the mushroom body role of identifying predictors, because the predictor will be the *whole* set of stimuli that strongly precede the US and not just some arbitrarily small sub-part. A further aspect of our pair-wise connectivity is the deterministic nature of PN to KC pattern mapping; a defined PN pattern will robustly establish a specific KC pattern. This is supported by Wang et al. (2001), who found that odour quality (and also quantity) can be identified with particular distributions of activity in the *Drosophila* mushroom body. Our combinatorial connectivity predicts that odours which share molecular components should yield different but partially overlapping patterns.

Our model also suggests a multi-modal integration role for the mushroom body, for two reasons. First, our assumed role—detector for patterns predicting salient events—is inherently a multi-modal sensory processing task; predictors will likely have simultaneous aspects in several modalities. Second, we assume a homogeneous connectivity and

operation of Kenyon cells, supported by *Drosophila* data (Yusuyama et al. 2002; Ito et al. 1997), that is not restricted to any single modality. Signals from a range of sensory modalities could project to the mushroom body, providing they satisfy the two constraints the model places on input signal: they must be roughly synchronous and fire at the same frequency (20 Hz); and more significantly, in terms of *how* they represent information, input signals must convey an invariant representation of stimuli. Early sensory pre-processing networks could affect such transformations to raw sensory output (Krichmar and Edelman 2002; Sporns and Alexander 2002).

Learning and memory processes are isolated to the lobe synapses, supporting the hypothesis that changes in these underly such abilities. Our synapse model is based on *Aplysia* research; a neuro-modulatory transmitter acts on recently active synapses to selectively increase their strengths (ADPF). Other related models exist. Gingrich and Byrne (1987) developed a synapse model of much greater detail for their simulation of the mechanism in *Aplysia*, and Damper et al. (2000) used a simpler model to incorporate ADPF ideas into a robot capable of performing learning behaviours. However, in terms of detail—at the level of simplified but realistic models (with counterparts such as IF-neurons and conductance based synapses)—the model developed here is a new contribution. Our model is extended to incorporate longer term memory by using two interacting internal "weight" variables that decay at different rates. This is a promising approach for building simplified/descriptive synapse models that display a range of learning dynamics because it is straightforward to introduce additional memory phases. We solved part of the weight stabilisation problem (i.e., when to trigger the consolidation process that converts short term memory into longer term memory) by incorporating more biology, in this case the hypothesised feedback mechanism from a spiking neuron onto its afferent synapse. Because learning is driven by a modulatory neuro-transmitter, we would expect that the US could be replaced by the micro-injection of a suitable neuro-transmitter directly into the *Drosophila* brain, as similarly occurs in the honeybee (Hammer and Menzel 1998).

The possibility of an *Aplysia* learning mechanism existing within *Drosophila* is supported by biochemical processes common to each animal. ADPF begins with the binding of serotonin to G-protein-linked receptors. Partly disrupting this form of signalling in *Drosophila* completely abolishes olfactory learning (Connolly et al. 1996). ADPF also involves the activation of a Ca^{2+}dependent adenylyl cyclase (AC). Commonality to *Drosophila* is found in *rutabaga* mutants, which are flies that lack a particular form of AC, and which show significant learning impairment (McGuire et al. 2003; Zars et al. 2000). The retrograde signalling is more speculative. However, there is a possible post-synaptic contribution to learning in *Drosophila* (Glanzman 2005). Xia et al. (2005) found that flies in which NMDA receptors (suggested to exist at neurons that project to the mushroom body) had been disrupted scored significantly less in odour-shock conditioning and abolished formation of long term memory.

Our model specifically addressed the need to prevent synapse strengths becoming too strong. We hypothesised a feedback mechanism between lobe neurons and modulatory input. The essential nature of this connection is that when an association has been invoked by the occurrence of a certain CS pattern, any reinforcement events that might quickly follow do not result in the over strengthening of synapses underlying the association. In our model this mechanism was implemented as inhibition from lobe neurons to a value neuron. If feedback does take this form (rather than being a biochemical process internal to KCs) then we would expect that learning in *Drosophila* that occurs when KC output is suppressed during training would result in distorted memories; such an association based on a first odour would be recalled if a different but similar odour were presented. Although speculative, this feedback mechanism relates to ideas in formal learning theories where the discrepancy ("prediction error") between expected reward/punishment and actual reward/punishment influences stimulus–response learning (Sutton and Barto 1981; Rescorla and Wagner 1972). In our model the firing of a lobe neuron corresponds to US expectation and $r(t)$ to (partly) prediction error.

The mushroom body is an increasingly attractive target system for computational modelers. It is manageable in size, relatively well understood, is implicated in diverse and interesting behaviours, including some that arguably constitute primitive forms of cognition (Giurfa 2003; Brembs and Heisenberg 2001; Menzel and Giurfa 1999). However, it is only recently that computational models have appeared (Wessnitzer et al. 2007; Nowotny et al. 2005; Huerta et al. 2004; Nowotny et al. 2003). A notable difference between these and ours is the choice of synapse learning rule. In other work there is a preference for using some form of STDP (Sjöström and Nelson 2002), rather than ADPF. Another important difference is overall function. While the present work focuses on the problems with simple associative learning, Wessnitzer et al. (2007) consider non-elemental learning, Nowotny et al. (2005) and Huerta et al. (2004) study classification of odours and Nowotny et al. (2003) examine how interconnections between Kenyon cells might assist the decoding of spatio-temporal patterns. These differences in model function and composition will decrease as the mushroom body is increasingly understood at behavioural, cellular and genetic levels. The challenges facing modelers is to determine which data is most significant, and to incorporate this into increasingly detailed but also insightful models that give an account of how mushroom body processes interact with other invertebrate brain regions to produce animal behaviour.

Table 1 Parameters for an IF neuron model of a Kenyon cell

Parameter	Value
V_{rest}	−84 mV
g_{leak}	0.26 nS
C	4.0 pF
V_{thresh}	−25.8 mV
V_{spike}	9.5 mV
V_{recov}	−40.2 mV
t_{refract}	10 ms

See text for details

Table 2 Parameter values used for calyx synapses

Parameter	value
E	0 mV
λ_{syn}	5 ms
g_{min}	1.1 nS
g_{max}	1.5 nS
g_{mean}	1.3 nS
g_{σ}	0.1 nS

See text for details

Table 3 Parameters for the calyx inhibitory synapses conveying LHI input

Parameter	value
E	−86 mV
λ_{syn}	4.3 ms
g_{syn}	25 nS
axon delay	15 ms

See text for details

Appendix

The parameters and values for the simulation components that were modelled as neurons or synapses are listed below. All equations were numerically solved user the (forward) Euler method, with time-step of $dt = 0.01$ ms (chosen so that it is much smaller than other time constants and so maintains accuracy of simulation).

Keynon cell parameters

Table 1 lists parameter values for the IF neuron model of a Kenyon cell. V_{rest}, g_{leak}, C and V_{thresh} are copied from physiological recordings (Wüstenberg et al. 2004 Table 1). V_{spike}, the voltage the membrane is set to when a spike occurs, is taken from a simulated Kenyon cell (Pelz et al. 1999 p. 1754). t_{refract} was initially set so that the maximum spiking frequency is approximately 60 Hz [To match the observed maximum firing frequency of a Kenyon cell Wüstenberg et al. 2004 p. 2596]; this required a value of 16 ms. However due to the time pressure of fitting coincidence detection into 50 ms windows (due to the 20 Hz spiking frequency of projection neurons) the refractory period was reduced to 10 ms, which in turn gives the calyx inhibition (if present) greater time to act. V_{recov} is chosen to be 15 mV less than V_{thresh}; although this is an artificial parameter of the IF neuron, the value was chosen to reflect the approximately 15 mV dip in $V(t)$ that follows an action potential, as observed in empirical traces (Wüstenberg et al. 2004). Lobe neurons used the same parameter values except for V_{recov}, which was set to −90 mV.

Calyx synapse parameters

Table 2 lists parameter values for the calyx synapses. The reversal conductance was chosen to be above the resting potential of the KC (which makes the synapse excitatory). The decay half-life was set low, particularly to keep it significantly under the refractory period of the KC. This is done to prevent the synaptic input resulting from one spike from contributing to more that one KC spike. g_{min} and g_{max} are the limits of intrinsic conductance which meet the behavioural requirements outlined in the main text. Actual g_{syn} values were drawn from a Gaussian with the specified g_{mean} and g_{σ}. The latter was set so that 95% of g_{syn} value fell within the guide limits

LHI Synapse Parameters

Table 3 lists parameter values for the LHI synapses. The reversal potential is set very close to V_{rest} of the Kenyon cells, so that inhibition is effective at resetting a neuron to its rest state. Synapse strength, g_{syn}, has to be very strong; it only has a very brief time to force $V(t)$ to the resting level. λ_{syn}, the decay half-life is set sufficiently large so that inhibition persists long enough to prevent out-of-phase spikes from triggering a KC. If set smaller the inhibition would be unsuccessful in resetting the KC; if set too high inhibition would leak into the next cycle of spikes travelling from the PNs to the KCs. The axon delay is the amount of time it takes inhibition to arrive at the KC, relative to a PN population synchronous response. i.e., if PNs are active at time $t = 0$ their spikes will arrive at the KCs almost immediately, while the LHI inhibition will follow at 15 ms. This value provides a hard limit to the level of loose-synchrony KCs can accommodate.

ADPF synapse parameters

Table 4 lists parameter values for the ADPF synapses that underpin learning and memory in the model. All synapses were configured identically. Most values were found through

Table 4 Parameters for ADPF based lobe synapses

Parameter	value
E	0 mV
λ_{syn}	8 ms
g_{syn}	0 pS
λ_p	50 ms
λ_q	200 ms
η_{syn}	2.5×10^{-10}
κ_{syn}	3.125×10^{-2}
α_{syn}	1.05
η_{base}	1.66
κ_{base}	8.33 x 10^{-4}

Refer to the text for details

a lengthy try-and-test manual process (i.e., the model was repeatedly "run" with slightly different parameter sets until learning and memory behaviour was reliably established). However some insights exist. λ_p, the pre-synaptic trace half-life decay, was chosen to reflect the approximate KC ISI of 50 ms. λ_q, the half-life of the retrograde signal decay, was given the larger value of 200 ms so that strengthening events were prolonged. Reversal potential was set above KC threshold in order to make the interaction excitatory; a relatively long decay half-life was chosen to support temporal integration of many active lobe synapses; and the 0 pS initial strength reflects the initial latent state of all synapses. The remaining parameters control change in $g_{syn}(t)$ and $g_{base}(t)$. The equations for these are first order linear differential equations. As such these parameters have, in a sense, standard roles. η controls rate of increase and κ is the leak, and limits the value the variable can obtain. It may appear odd that η_{syn} is eight orders of magnitude *less* than κ_{syn}. However in Eq. (1) κ_{syn} is multiplied with a conductance term having magnitude 10^{-12}, while η_{syn} combines with $r(t)$ and $p(t)$ having magnitude 1. The net effect is that the growth term is three to four orders of magnitude larger than decay; this results in the strong but brief spurts of g_{syn} growth visible in Fig. 7. The remaining parameter α_{syn} was chosen to be just above zero; this ensures that during consolidation $g_{base}(t)$ slightly overshoots $g_{syn}(t)$, which supports robust learning.

References

Antonov I, Antonova I, Kandel E, Hawkins R (2003) Activity-dependent presynaptical facilitation and hebbian ltp are both required and interact during classical conditioning in Aplysia. Neuron 37:135–147

Bazhenov M, Stopfer M, Rabinovich M, Abarbanel H, Sejnowski T, Laurent G (2001a) Model of cellular and network mechanisms for odor-evoked temporal patterning in the locust antennal lobe. Neuron 30:569–581

Bazhenov M, Stopfer M, Rabinovich M, Huerta R, Abarbanel H, Sejnowski T, Laurent G (2001b) Model of transient osciallatory synchronization in the locust antennal lobe. Neuron 30:553–567

Brembs B, Heisenberg M (2001) Conditioning with compound stimuli in Drosophila melanogaster in the flight simulator. J Exp Biol 204:2849–2859

Carew T (2000) Behavioral neurobiology: the cellular organization of natural behaviour. Sinauer Associates, Massachusetts

Connolly J, Roberts I, Armstrong J, Kaiser K, Forte M, Tully T, O'ane C (1996) Associative learning disrupted by impaired gs signaling in Drosophila mushroom bodies. Science 274:2104–2107

Damper R, French R, Scutt TW (2000) Arbib: an autonomous robot based on inspirations from biology. Robotics Auton Syst 31:247–274

Damper R, French R, Scutt T (2001) The hi-noon neural simulator and its applications. Microelectron Reliab 41(12):2051–2065

de Belle J, Heisenberg M (1994) Associative odor learning in Drosophila abolished by chemical ablation of mushroom bodies. Science 263:692–695

Davis R (2005) Olfactory memory formation in Drosophila: from molecular to systems neuroscience. Annu Rev Neurosci 28:275–302

Dubnau J, Grady L, Kitamoto T, Tully T (2001) Disruption of neurotransmission in Drosophila mushroom body blocks retrieval but not acquisition of memory. Nature 411:476–480

Farris S (2005) Evolution of insect mushroom bodies: old clues, new insights. Arthropod Struct Dev 34:211–234

Ferveur J, Strtkuhl K, Stocker R, Greenspan R (1995) Genetic feminization of brain structures and changed sexual orientation in male Drosophila. Science 267:902–905

Gingrich K, Byrne J (1985) Simulation of synaptic depression, posttetanic potentiation, and presynaptic facilitation of synaptic potentials from sensory neurons mediating gill-withdrawal reflex in Aplysia. J Neurophysiol 53(3):652–669

Gingrich K, Byrne J (1987) Single-cell neuronal model for associative learning. J Neurophysiol 57(6):1705–1715

Giurfa M (2003) Cognitive neuroethology: dissecting non-elemental learning in a honeybee brain. Curr Opin Neurobiol 13:726–735

Glanzman D (2005) Associative learning: Hebbian flies. Curr Biol 15:R416

Hammer M (1993) An identified neuron mediates the unconditioned stimulus in associative olfactory learning in honeybees. Nature 366:59–63

Hammer M, Menzel R (1998) Multiple sites of associative odor learning as revealed by local brain microinjections of octopamine in honeybees. Learn Mem 5:146–156

Hawkins R, Abrams T, Carew T, Kandel E (1983) A cellular mechanism of classical conditioning in Aplysia: activity-dependent amplification of presynaptic facilitation. Science 219:400–405

Heisenberg M (2003) Mushroom body memoir: from maps to models. Nature Revi Neurosci 4:266–275

Heisenberg M, Borst A, Wagner S, Byers D (1985) Drosophila mushroom body mutants are deficient in olfactory learning. J Neurogenet 2(1):1–30

Heisenberg M, Heusipp M, Wanke C (1995) Structural plasticity in the Drosophila brain. J Neurosci 15(3):1951–1960

Huerta R, Nowotny T, Garcia-Sanchez M, Abarbanel H, Rabinovish M (2004) Learning classification in the olfactory system of insects. Neural Comput 16:1601–1640

Ito K, Awano W, Suzuki K, Hiromi Y, Yamamoto D (1997) The Drosophila mushroom body is a quadruple structure of clonal units each of which contains a virtually identical set of neurones and glial cells. Development 124:761–771

Koch C (1999) Biophysics of computation. Oxford University Press, Oxford

Krichmar J, Edelman G (2002) Machine psychology: autonomous behavior, perceptual categorization and conditioning in a brain-based device. Cerebral Cortex 12:818–830

Lechner H, Byrne J (1998) New perspectives on classical conditioning: a synthesis of hebbian and non-hebbian mechanisms. Neuron 20:355–358

Liu L, Wolf R, Ernst R, Heisenberg M (1999) Context generalization in Drosophila visual learning requires the mushroombodies. Nature 400:753–756

Margulies C, Tully T, Dubnau J (2005) Deconstructing memory in Drosophila. Curr Biol 15:R700–R713

Martin JR, Ernst R, Heisenberg M (1998) Mushroom bodies suppress locomotor activity in Drosophila melanogaster. Learn Mem 5:179–191

McGuire S, Le P, Davis R (2001) The role of Drosophila mushroom body signalling in olfactory memory. Science 293:1330–1333

McGuire S, Le P, Osborn A, Matsumoto K, Davis R (2003) Spatiotemporal rescue of memory dysfunction in Drosophila. Science 302:1765–1768

Menzel R, Giurfa M (1999) Cognition by a mini brain. Nature 400:718–719

Nowotny T, Rabinovish M, Huerta R, Abarbanel H (2003) Decoding temporal information through slow lateral excitation in the olfactory system of insects. J Comput Neurosci 15:271–281

Nowotny T, Huerta R, Abarbanel H, Rabinovish M (2005) Self-organisation in the olfactory system: one shot odor recognition in insects. Biol Cybern 93:436–446

Olshausen B, Field D (2004) Sparse coding of sensory inputs. Curr Opin Neurobiol 14:481–487

Pelz C, Jander J, Rosenboom H, Hammer M, Menzel R (1999) I_A in kenyon cells of the mushroom body of honeybees resembles shaker currents: kinetics, modulation by k^+, and simulation. J Neurophysiol 81:1749–1759

Rescorla R, Wagner A (1972) A theory of pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. In: Black A, Prokasy W (eds) Classical conditioning II, Appleton Century Crofts, pp 64–99

Roberts A, Glanzman D (2003) Learning in Aplysia: looking at synaptic plasticity from both sides. Trends Neurosci 26(12):662–670

Roman G, Davis R (2001) Molecular biology and anatomy of *Drosophila olfactory* associative learning. BioEssays 23:571–581

Schwaerzel M, Monastirioti M, Scholz H, Friggi-Grelin F, Birman S, Heisenberg M (2003) Dopamine and octopamine differentiate between aversive and appetitive olfactory memories in Drosophila. J Neurosci 23(33):10,495–10,502

Sjöström P, Nelson S (2002) Spike timing, calcium signals and synaptic plasticity. Curr Opin Neurobiol 12:305–314

Sporns O, Alexander W (2002) Neuromodulation and plasticity in an autonomous robot. Neural Netw 15:761–774

Strausfeld N, Hansen L, Li Y, Gomez R, Ito K (1998) Evolution, discovery, and interpretations of arthropod mushroom bodies. Learn Mem 5:11–37

Sutton R, Barto A (1981) Toward a modern theory of adaptive networks: expectation and prediction. Psychol Rev 88:135–170

Trappenberg T (2002) Fundamentals of Computational Neuroscience. Oxford University Press, Oxford

Waddell S, Quinn W (2001) What can we teach Drosophila? what can they teach us?. Trends Genet 17:719–726

Walters E, Byrne J (1983) Associative conditioning of single sensory neurons suggests a cellular mechanism for learning. Science 219:405–408

Wang Y, Wright N, Guo H, Xie Z, Svoboda K, Malinow R, Smith D, Zhong Y (2001) Genetic manipulation of the odor-evoked distributed neural activity in the Drosophila mushroom body. Neuron 29:267–276

Wehr M, Laurent G (1996) Odour encoding by temporal sequences of firing in oscillating neural assemblies. Nature 384:162–166

Wessnitzer J, Webb B, Smith D (2007) A model of non-elemental associative learning in the mushroom body neuropil of the insect brain. In: Beliczynski B, Dzielinski A, Iwanowski M, Ribeiro B (eds) Proceedings of the international conference on adaptive and natural computing algorithms, Lecture Notes in Computer Science, vol 4431. Springer, Heidelberg

Wüstenberg D, Boytcheva M, Grünewald B, Byrne J, Menzel R, Baxter D (2004) Current- and volatage-clamp recordings and computer simulations of kenyon cells in the honeybee. J Neurophysiol 92:2589–2603

Xia S, Miyashita T, Fu TF, Lin WY, Wu CL, Pyzocha L, Lin IR, Saitoe M, Tully T, Chiang AS (2005) Nmda receptors mediate olfactory learning and memory in Drosophila. Curr Biol 15:603–615

Yusuyama K, Meinertzhagen I, Schurmann FW (2002) Synaptic organization of the mushroom body calyx in *Drosophila melanogaster*. J Comp Neurol 445:211–226

Zars T (2000) Behavioral functions of the insect mushroom bodies. Curr Opin Neurobiol 10:790–795

Zars T, Wolf R, Davis R, Heisenberg M (2000) Tissue-specific expression of a type I adenylyl cyclase rescues the rutabaga mutant memory defect: in search of the engram. Learn Mem 7:18–31